

**A QUANTITATIVE REAL OPTIONS METHOD FOR AVIATION
TECHNOLOGY DECISION-MAKING IN THE PRESENCE OF
UNCERTAINTY**

A Dissertation
Presented to
The Academic Faculty

by

Cédric Y. Justin

In Partial Fulfillment
Of the Requirements for the Degree
Doctor of Philosophy in the
School of Aerospace Engineering

Georgia Institute of Technology
December 2015

Copyright © 2015 by Cédric Y. Justin

**A QUANTITATIVE REAL OPTIONS METHOD FOR AVIATION
TECHNOLOGY DECISION-MAKING IN THE PRESENCE OF
UNCERTAINTY**

Approved by:

Prof. Dimitri Mavris, Advisor
School of Aerospace Engineering,
Georgia Institute of Technology

Prof. Brian German
School of Aerospace Engineering,
Georgia Institute of Technology

Prof. Daniel Schrage
School of Aerospace Engineering,
Georgia Institute of Technology

Dr. Simon Briceño
School of Aerospace Engineering,
Georgia Institute of Technology

Dr. Frederic Villeneuve
Design Systems Engineering,
Siemens Energy

Dr. Ismael Fernandez
Founder and President,
Greengate Consulting LLC

Date Approved: November 13th 2015

ACKNOWLEDGEMENTS

This was a very long journey!

I would like to thank my advisor, Prof. Dimitri Mavris, for his continuous support throughout this enlightening experience as well as to all my committee members Prof. Schrage, Prof. German, Dr. Briceno, Dr. Villeneuve, and Dr. Fernandez for their help and guidance at the various stages of my research.

I would also like to thank my family, my friends, Alexia, and of course Rodolf for their unconditional support, help, and patience during this endeavor.

TABLE OF CONTENTS

| | |
|---|------|
| ACKNOWLEDGEMENTS | iii |
| LIST OF TABLES | xi |
| LIST OF FIGURES | xvii |
| LIST OF SYMBOLS AND ABBREVIATIONS | xxvi |
| SUMMARY | xxix |
| CHAPTER 1 : INTRODUCTION | 1 |
| 1.1 Aircraft manufacturing industry | 2 |
| 1.2 Aircraft design process | 3 |
| 1.2.1 Different phases of design | 5 |
| 1.2.2 Modern system design methods | 7 |
| 1.3 Traditional challenges | 15 |
| 1.3.1 Major investments and limited returns | 15 |
| 1.3.2 Technical and technological risks | 17 |
| 1.3.3 Timescale to bring new aircraft to the market | 19 |
| 1.3.4 Supply chain management | 20 |
| 1.4 Emergence of new challenges | 21 |
| 1.4.1 Longer development cycles increase risk | 21 |
| 1.4.2 New competition emerging | 23 |
| 1.5 Challenges to the aerospace industry: a summary | 26 |
| CHAPTER 2 : MOTIVATION | 27 |
| 2.1 First motivating example | 27 |

| | | |
|---|--|----|
| 2.2 | Second motivating example | 35 |
| 2.3 | Third motivating example | 39 |
| 2.4 | Research motivation and thesis organization | 42 |
| CHAPTER 3 : PROBLEM DEFINITION..... | | 46 |
| 3.1 | Valuation methods..... | 46 |
| 3.1.1 | Payback period..... | 46 |
| 3.1.2 | Discounted cash flow analysis and the net present value | 47 |
| 3.1.3 | Internal rate of return | 56 |
| 3.1.4 | Modified internal rate of return..... | 59 |
| 3.1.5 | Real options analysis..... | 60 |
| 3.1.6 | Valuation methods: a summary | 68 |
| 3.2 | Marketing and competitive analysis methods | 69 |
| 3.2.1 | SWOT analysis | 70 |
| 3.2.2 | Five forces analysis..... | 71 |
| 3.2.3 | Game theoretic analysis | 74 |
| 3.2.4 | Prospect theory..... | 77 |
| 3.2.5 | Competitive methods: a summary | 80 |
| CHAPTER 4 : REAL OPTIONS THINKING | | 82 |
| 4.1 | Genesis and challenges..... | 82 |
| 4.1.1 | Borrowing a paradigm from the finance industry..... | 82 |
| 4.1.2 | An interesting concept harder to implement in practice | 83 |
| 4.1.3 | Substantiating real options thinking: the Marketed Asset Disclaimer..... | 94 |
| 4.1.4 | What about the dynamics of the underlying real assets value? | 98 |

| | | |
|---|---|-----|
| 4.2 | Numerical recipes for real options | 101 |
| 4.2.1 | Three venues for real options evaluation | 101 |
| 4.2.2 | Summary of methodologies for real options evaluation | 113 |
| 4.3 | Probability measure and change of probability measure | 115 |
| 4.3.1 | Risk-neutral measure for some common stochastic processes | 117 |
| 4.3.2 | Esscher transform for option pricing | 118 |
| 4.3.3 | Non-parametric Esscher transform and real options pricing | 123 |
| 4.4 | Path-dependent options | 128 |
| 4.4.1 | Managerial flexibility and trigger events | 128 |
| 4.4.2 | American and Bermudan real options | 129 |
| 4.4.3 | Early-exercise boundary | 130 |
| 4.4.4 | Monte Carlo simulations for pricing American options | 133 |
| 4.4.5 | Pricing American options using Longstaff-Schwartz algorithm | 134 |
| 4.4.6 | Bootstrapping for American and Bermudan options | 138 |
| 4.4.7 | Summary | 143 |
| 4.5 | Meeting the Georgetown challenge and more | 144 |
| CHAPTER 5 : SCOPING THE PROBLEM: RESEARCH QUESTIONS AND | | |
| HYPOTHESES | | |
| 5.1 | Revisiting the problem | 147 |
| 5.2 | Research questions and hypotheses | 152 |
| 5.3 | Matching research questions and hypotheses | 159 |
| CHAPTER 6 : METHODOLOGY | | |
| 6.1 | Brief overview of an industry problem | 161 |

| | | |
|---|--|-----|
| 6.1.1 | Windows of possibilities and windows of opportunities | 161 |
| 6.1.2 | Development of the Performance Improvement Package | 164 |
| 6.1.3 | Identification of decision windows | 164 |
| 6.1.4 | Objectives to be attained and expected results | 165 |
| 6.2 | Methodology development..... | 167 |
| 6.2.1 | Traditional evaluation method | 167 |
| 6.2.2 | Proposed evaluation method | 170 |
| 6.2.3 | Detailed implementation of proposed evaluation method | 171 |
| CHAPTER 7 : EXPERIMENTAL PLAN - VERIFICATION..... | | 184 |
| 7.1 | Preparing the verification for the real options analysis..... | 184 |
| 7.1.1 | Verification process | 185 |
| 7.1.2 | Graphical tests – QQ Plots..... | 187 |
| 7.1.3 | Statistical tests – Kolmogorov-Smirnov test..... | 192 |
| 7.1.4 | Statistical tests – Testing the mean using z-tests and t-tests | 195 |
| 7.1.5 | Similarity tests – Hausdorff distance | 198 |
| 7.1.6 | Comparison tests – Solving PDE with finite-difference methods | 200 |
| 7.2 | Preliminary testing and lessons learned | 204 |
| 7.2.1 | Variability of results | 204 |
| 7.2.2 | Pooling sample of returns before bootstrap resampling | 205 |
| 7.2.3 | Antithetic variates, moment matching, and control variates..... | 209 |
| 7.2.4 | Quasi-Monte Carlo simulations | 224 |
| 7.2.5 | Scoping conditional expectation regressions | 231 |
| 7.2.6 | Projection basis selection for continuation value regressions..... | 237 |

| | |
|--|------------|
| 7.2.7 Multi-start Monte Carlo simulations..... | 240 |
| 7.2.8 Regression and filtering of critical prices to form trigger boundary | 248 |
| 7.2.9 Matching new research questions and hypotheses..... | 254 |
| 7.3 Verification of technical hypotheses | 259 |
| 7.3.1 Non-parametric Esscher transform technique..... | 259 |
| 7.3.2 Bootstrapping technique for resampling | 278 |
| 7.3.3 Combined non-parametric Esscher transform and resampling..... | 293 |
| 7.3.4 Least-squares Monte Carlo for trigger boundary generation..... | 310 |
| 7.3.5 Least-squares Monte Carlo for option valuation | 315 |
| 7.3.6 Real options for staggered development valuation..... | 330 |
| CHAPTER 8 : PROOF OF CONCEPT DESCRIPTION AND IMPLEMENTATION.. | 336 |
| 8.1 Presentation of an industry problem to be investigated | 337 |
| 8.1.1 Performance Improvement Package (PIP)..... | 337 |
| 8.1.2 Performance Improvement Package development timeline | 339 |
| 8.1.3 Technology description..... | 340 |
| 8.1.4 Competitive setting and technology development timeline..... | 344 |
| 8.2 Aircraft and engine operating economics..... | 345 |
| 8.2.1 Enabling the aircraft and engine economic evaluation | 348 |
| 8.2.2 Airline operations and network analysis..... | 349 |
| 8.2.3 Revenue analysis..... | 350 |
| 8.2.4 Cost analysis | 351 |
| 8.2.5 Calibration of the engine maintenance modules in i-CARE..... | 371 |
| 8.2.6 Verification and validation of aircraft and engine analysis method | 377 |

| | | |
|--|--|-----|
| 8.3 | Developing a market behavior model for the PIP adoption | 387 |
| 8.3.1 | Market segmentation..... | 387 |
| 8.3.2 | Simplified market behavior model..... | 390 |
| 8.3.3 | Decreasing market size and value leakages | 391 |
| 8.4 | Market uncertainties driving PIP demand..... | 395 |
| 8.4.1 | Down-selecting uncertainties..... | 396 |
| 8.4.2 | Calibration of jet-fuel price model..... | 397 |
| 8.4.3 | Calibration of carbon emission allowances price model | 400 |
| 8.4.4 | Treatment of correlations..... | 402 |
| CHAPTER 9 : EXPERIMENTAL PLAN - VALIDATION | | 404 |
| 9.1 | Real options to model managerial flexibility | 404 |
| 9.1.1 | Validation process and criteria for success | 406 |
| 9.1.2 | Preliminary testing and initial struggles | 407 |
| 9.1.3 | Performance Improvement Package baseline evaluation | 408 |
| 9.1.4 | Limited-flexibility Performance Improvement Package development..... | 412 |
| 9.1.5 | Flexible Performance Improvement Package evaluation | 414 |
| 9.1.6 | Fully flexible Performance Improvement Package evaluation..... | 417 |
| 9.1.7 | Comparisons | 420 |
| 9.2 | Sequential moves for competitive scenarios | 425 |
| 9.3 | Back to overarching research question..... | 430 |
| CHAPTER 10 : CONCLUSIONS AND CONTRIBUTIONS | | 433 |
| 10.1 | Overarching research question | 433 |
| 10.2 | Method research questions and hypotheses..... | 434 |

| | | |
|-------------|---|-----|
| 10.3 | Modeling research questions and hypotheses..... | 435 |
| 10.4 | Technical research questions and hypotheses | 437 |
| 10.5 | List of contributions..... | 438 |
| 10.5.1 | Contributions related to the analyses of staggered investments | 439 |
| 10.5.2 | Contributions related to the proof-of-concept..... | 440 |
| 10.5.3 | Other miscellaneous contributions..... | 440 |
| 10.6 | Proposed extensions | 441 |
| APPENDIX A: | Geometric Brownian Motion..... | 443 |
| APPENDIX B: | Jump Diffusion Processes..... | 456 |
| APPENDIX C: | Probability measure | 464 |
| APPENDIX D: | Esscher transform | 479 |
| APPENDIX E: | Modern portfolio theory | 486 |
| APPENDIX F: | Capital asset pricing model..... | 491 |
| APPENDIX G: | Black-Scholes-Merton option pricing model | 495 |
| APPENDIX H: | Implementation of real options method..... | 503 |
| APPENDIX I: | Solving partial differential equations with finite-difference schemes | 511 |
| APPENDIX J: | Implementation verification | 515 |
| APPENDIX K: | Pip Results | 521 |
| APPENDIX L: | Verification definitions..... | 522 |
| REFERENCES | | 526 |
| VITA | | 554 |

LIST OF TABLES

| | |
|---|----|
| Table 1: Aircraft development costs: projected and final estimate at completion..... | 16 |
| Table 2: Aircraft development times | 22 |
| Table 3: Identified challenges to the design process | 26 |
| Table 4: Main assumptions for economic estimations..... | 29 |
| Table 5: Economic estimation for the value of the A340-500/600 programs..... | 30 |
| Table 6: Four different types of jet-fuel price scenarios investigated | 33 |
| Table 7: Some jet-fuel time-series results related to the A340-500/600 programs | 34 |
| Table 8: Assumptions for the A380 business case estimations | 37 |
| Table 9: A380 project value and internal rate of return estimations for different production costs | 37 |
| Table 10: Aircraft development notional example | 57 |
| Table 11: Conflicting NPV and IRR values for the new and derivative projects..... | 58 |
| Table 12: Mapping financial options jargon and real options jargon | 63 |
| Table 13: Different types of real options | 67 |
| Table 14: Strengths and Weaknesses of a Selection of Valuation Methods..... | 69 |
| Table 15: SWOT Matrix | 71 |
| Table 16: Strengths and Weaknesses of Competitive Analysis Methods..... | 81 |
| Table 17: Main assumptions underpinning the Black-Scholes model..... | 84 |
| Table 18: Translating main assumptions for real options valuation using the Black- Scholes model | 84 |

| | |
|--|-----|
| Table 19: Required steps for Black-Scholes formula derivation using three valuation techniques | 94 |
| Table 20: Real options valuation methods, strengths and weaknesses | 114 |
| Table 21: Changing the measure of popular stochastic processes for real options | 118 |
| Table 22: Addressing the challenges facing the analysis of long-term corporate investment analyses featuring flexibility | 146 |
| Table 23: Q-Q plots with 80,000 data points for various simulated and theoretical distributions..... | 191 |
| Table 24: Kolmogorov-Smirnov statistical test implementation | 195 |
| Table 25: Hausdorff distance implementation verification | 200 |
| Table 26: Grid selection for finite-difference numerical scheme | 202 |
| Table 27: Range for input parameters..... | 203 |
| Table 28: Finite-difference scheme results for American call and put options | 204 |
| Table 29: Q-Q plots for the return distribution induced by bootstrap resampling of two cases of geometric Brownian motions Merton jump diffusion processes | 206 |
| Table 30: Q-Q plots with increasing pooling number | 208 |
| Table 31: American call option price and early-exercise premium for underlying assets following a geometric Brownian motion | 210 |
| Table 32: Comparison between original approach (left graphs) and control-variate improved approach (right graphs)..... | 221 |
| Table 33: Comparison between traditional continuation value regressions and scoped continuation value regressions..... | 236 |

| | |
|--|-----|
| Table 34: Comparison of the average error of early-exercise boundaries using traditional Monte Carlo simulations (red line) and using multi-start simulations (x) for five different call options | 245 |
| Table 35: Comparison of boundary position errors for a crude locus of critical prices (+) and for the filtered and regressed trigger boundary (x) using five call options with each test case repeated fifteen times. | 252 |
| Table 36: Summary of Q-Q plot tests with GBM..... | 263 |
| Table 37: Q-Q Plots for verification of the non-parametric Esscher transformation applied to geometric Brownian motions..... | 265 |
| Table 38: Summary of Q-Q plot tests with Merton jump diffusion model | 267 |
| Table 39: Q-Q Plots for verification of the non-parametric Esscher transformation applied to Merton jump diffusion processes..... | 269 |
| Table 40: Kolmogorov-Smirnov test for twenty combinations of parameters | 271 |
| Table 41: Kolmogorov-Smirnov statistical tests for twenty cases of Merton Jump diffusion process | 273 |
| Table 42: z -tests and t -tests for the mean returns of twenty cases of geometric Brownian motions..... | 276 |
| Table 43: z -tests and t -tests for the mean returns of twenty cases of Merton jump diffusion processes..... | 277 |
| Table 44: Summary of the bootstrap resampling tests using Q-Q plots with GBM..... | 281 |
| Table 45: Q-Q Plots for verification of the bootstrap technique applied to Esscher transformed geometric Brownian motions..... | 283 |

| | |
|---|-----|
| Table 46: Summary of the bootstrap resampling tests with Q-Q plots for Merton jump diffusion | 285 |
| Table 47: Q-Q Plots for verification of the bootstrap technique applied to Esscher transformed Jump Diffusion processes | 287 |
| Table 48: Kolmogorov-Smirnov statistical tests for twenty cases of geometric Brownian motions..... | 289 |
| Table 49: Kolmogorov-Smirnov statistical tests for twenty cases of Merton Jump Diffusion processes..... | 291 |
| Table 50: Summary of the combined approach tests using Q-Q plots with GBM | 295 |
| Table 51: Q-Q Plots for verification of the combined approach for geometric Brownian motions..... | 297 |
| Table 52: Summary of the combined approach tests with Q-Q plots for Merton jump diffusion | 299 |
| Table 53: Q-Q Plots for verification of the combined approach for Merton jump diffusion processes | 301 |
| Table 54: Kolmogorov-Smirnov statistical tests for twenty cases of geometric Brownian motions..... | 303 |
| Table 55: Kolmogorov-Smirnov statistical tests for twenty cases of Merton Jump Diffusion processes..... | 305 |
| Table 56: Correlation between final results and intermediate results..... | 307 |
| Table 57: z -tests and t -tests for the mean returns of twenty cases of geometric Brownian motions..... | 308 |

| | |
|---|-----|
| Table 58: z-tests and t-tests for the mean returns of twenty cases of Merton jump diffusion processes..... | 309 |
| Table 59: European option prices for twenty cases of geometric Brownian motions | 321 |
| Table 60: z-test and t-test for European option – Repeated cases of geometric Brownian motions..... | 322 |
| Table 61: New z-test and t-test for the European option price of the repeated experiment | 323 |
| Table 62: European option prices for twenty cases of Merton jump diffusion processes | 328 |
| Table 63: z-tests and t-tests for European options – Repeated cases of Merton jump diffusion processes..... | 330 |
| Table 64: PIP-light key metrics with respect to current baseline turbofan engine | 342 |
| Table 65: PIP-Involved key metrics with respect to current baseline turbofan engine .. | 343 |
| Table 66: Life-limited part lives and costs for typical narrowbody aircraft (<i>Aircraft Commerce</i> [208]) | 371 |
| Table 67: Engine removal workscope and costs..... | 376 |
| Table 68: Inputs for the verification test case..... | 378 |
| Table 69: Comparison of results for verification test case | 380 |
| Table 70: Network description and operating statistics | 383 |
| Table 71: Selection of the WACC for market segmentation | 389 |
| Table 72: Variance ratio tests for jet-fuel price time series | 398 |
| Table 73: Cowles-Jones (CJ) ratio test for jet-fuel price | 398 |
| Table 74: Variance ratio test for CO ₂ emissions..... | 401 |

| | |
|---|-----|
| Table 75: Cowles-Jones (CJ) ratio test for CO ₂ emission allowances..... | 402 |
| Table 76: Baseline PIP development assumptions | 409 |
| Table 77: baseline PIP statistics with customer NPV and market size..... | 411 |
| Table 78: Baseline PIP development program value..... | 412 |
| Table 79: Limited-flexibility PIP development assumptions | 413 |
| Table 80: Limited-flexibility PIP development program value | 414 |
| Table 81: Flexible PIP development program value..... | 417 |
| Table 82: Fully flexible PIP development program value | 419 |
| Table 83: Input parameters for PIP-Involved evaluation..... | 424 |
| Table 84: Market share assumptions in different competitive scenarios | 428 |
| Table 85: PIP development value in different competitive scenarios..... | 428 |
| Table 86: U.S. Gulf Coast Kerosene-type jet-fuel spot price – June 1994 - May 1997 . | 449 |
| Table 87: Probability of jet-fuel exceeding certain thresholds at certain dates | 451 |
| Table 88: Probability of jet-fuel price hitting certain thresholds by certain dates | 454 |
| Table 89: Expected time for the jet-fuel price hitting certain thresholds | 455 |
| Table 90: Selection of a time-space grid for European options with unit spot to strike ratios..... | 516 |
| Table 91: Comparison between finite-difference method and Black-Scholes for European put and call options | 517 |

LIST OF FIGURES

| | |
|--|-----|
| Figure 1: Traditional Development Process (adapted from [3])..... | 7 |
| Figure 2: Georgia Tech IPPD methodology adapted from Marx et al. [5]..... | 9 |
| Figure 3: Design space exploration..... | 10 |
| Figure 4: Generation of surrogate models | 12 |
| Figure 5: Robust Design Simulation (adapted from Mavris and Bandte [8])..... | 14 |
| Figure 6: A340-500 and A340-600 deliveries since program launch..... | 29 |
| Figure 7: US Gulf Coast kerosene-type jet-fuel spot price..... | 31 |
| Figure 8: Initial A380 yearly production forecasts | 37 |
| Figure 9: Consumption choice in a two-period model..... | 50 |
| Figure 10: Development program value for different discount rates..... | 58 |
| Figure 11: Vanilla options payoffs at expiration (accounting for the option premium)... | 65 |
| Figure 12: Normal form and extensive form representation of a strategic game | 75 |
| Figure 13: Notional value function representation | 79 |
| Figure 14: Real options pricing with binomial lattices..... | 104 |
| Figure 15: Monte Carlo simulations using some popular stochastic processes..... | 108 |
| Figure 16: Simulations and resulting business prospect value distributions at expiration under physical and risk-neutral probability measures | 109 |
| Figure 17: Non-parametric Esscher transform for change of probability measure | 127 |
| Figure 18: Early-exercise boundaries for American put options..... | 132 |
| Figure 19: American option valuation with regression..... | 136 |

| | |
|--|-----|
| Figure 20: Bootstrap method to generate trajectories | 142 |
| Figure 21: Research questions and hypotheses..... | 160 |
| Figure 22: Timeline of manufacturer development stream..... | 162 |
| Figure 23: Value leakages and their effects on the opportunity window | 163 |
| Figure 24: Development timeline and associated milestones | 164 |
| Figure 25: Deriving decision windows | 165 |
| Figure 26: Early-investment boundaries at each decision window | 166 |
| Figure 27: Traditional methodology to build and evaluate a business case | 169 |
| Figure 28: Proposed methodology to build and evaluate business cases in R&D..... | 171 |
| Figure 29: Scenario generation and uncertainty modeling | 172 |
| Figure 30: Market analysis and the determination of preferences | 175 |
| Figure 31: Market reaction and demand estimation | 176 |
| Figure 32: Simulation of the evolution of the R&D program value | 178 |
| Figure 33: Change of probability measure and risk-neutral terminal value distribution | 179 |
| Figure 34: Simulation of R&D program value under risk-neutral probability measure. | 180 |
| Figure 35: Derivation of the early-investment policy for path-dependent real options.. | 181 |
| Figure 36: Analysis of the optimal set of conditions to launch R&D programs..... | 182 |
| Figure 37: Decomposition of proposed methodology to evaluate R&D business cases | 183 |
| Figure 38: V-Model for systems engineering | 186 |
| Figure 39: Verification Process | 187 |
| Figure 40: Relative difference distribution for various test cases (finite-difference scheme compared to Black-Scholes solution) | 203 |
| Figure 41: Variance reduction factor versus of control variate correlation | 216 |

| | |
|--|-----|
| Figure 42: Sampling control variates at maturity (left graph) is less correlated with option payoffs than sampling control variates at exercise (right graph)..... | 217 |
| Figure 43: Illustration of the control variate improved regressions on two trajectories. The discounted underlying is used as control variate, sampled at exercise of the option, and projected onto the set of basis-functions..... | 219 |
| Figure 44: 20,000 uniformly distributed numbers across two dimensions: (a) represents pseudo-random numbers from a linear congruential generator while (b) represents a Sobol sequence..... | 225 |
| Figure 45: Serial correlation at various lags for a sequence of 1,000 Sobol numbers.... | 226 |
| Figure 46: Relative error as a function of the number of simulations for different pseudo-random and quasi-random number generators..... | 230 |
| Figure 47: Standard error as a function of the number of simulations for different pseudo-random and quasi-random number generators..... | 230 |
| Figure 48: Removing points inside the natural boundary (right graph) scopes down the conditional expectation regression domain and improves the estimation of critical prices for American and Bermudan real options (call option depicted)..... | 232 |
| Figure 49: Bisection algorithm is used to search for critical prices using a single set of returns | 234 |
| Figure 50: Scoping the regression domain for the conditional expectation continuation value..... | 235 |
| Figure 51: Continuation values for a call option on a geometric Brownian motion..... | 237 |
| Figure 52: Continuation value fitted with the first four Legendre polynomials for a call option on a geometric Brownian motion. | 239 |

| | |
|---|-----|
| Figure 53: Critical price search with ad-hoc set of basis functions ϕ_i exhibits fewer wiggles | 240 |
| Figure 54: Solving for the critical price at each time step involves extrapolations close to the starting point of the simulation since none of the trajectories cross the trigger boundary (left graph). This results in noisy and approximate critical price solutions (right graph) close to the starting point..... | 241 |
| Figure 55: Simulation with multiple starting points leads to fewer extrapolations and less noisy critical price estimates, even far from expiration..... | 243 |
| Figure 56: Range of values for starting points in Multi-Start Monte Carlo..... | 244 |
| Figure 57: Regression of critical price estimates (red) and true boundary (black)..... | 250 |
| Figure 58: Removing outliers using semi-Studentized residuals..... | 252 |
| Figure 59: Updated set of research questions and hypotheses..... | 258 |
| Figure 60: Distribution of p -values for 600 Kolmogorov-Smirnov tests (geometric Brownian motions)..... | 272 |
| Figure 61: Distribution of p -values for 600 Kolmogorov-Smirnov tests (Merton jump diffusion processes) | 274 |
| Figure 62: Distribution of p -values for 600 Kolmogorov-Smirnov tests (geometric Brownian motions)..... | 290 |
| Figure 63: Distribution of p -values for 600 Kolmogorov-Smirnov tests (Jump Diffusion processes)..... | 292 |
| Figure 64: Distribution of p -values for 600 Kolmogorov-Smirnov tests (geometric Brownian motions)..... | 304 |

| | |
|--|-----|
| Figure 65: Distribution of p -values for 600 Kolmogorov-Smirnov tests (Jump Diffusion processes)..... | 306 |
| Figure 66: Hausdorff distance between experimental and reference trigger boundaries for 600 cases of geometric Brownian motions..... | 313 |
| Figure 67: RMSE of the experimental trigger boundary for 600 cases of geometric Brownian motions..... | 313 |
| Figure 68: Maximum relative error of the experimental trigger boundary for 600 cases of geometric Brownian motions..... | 314 |
| Figure 69: Initial relative error of the experimental trigger boundary for 600 cases of geometric Brownian motions..... | 314 |
| Figure 70: Relative error of European call option prices with underlying following a geometric Brownian motion (reference Black Scholes formula)..... | 318 |
| Figure 71: Relative width of 95% confidence interval for European call option prices on underlying following a geometric Brownian motion..... | 318 |
| Figure 72: Relative error of American call option prices with underlying following a geometric Brownian motions (reference finite-difference solver)..... | 319 |
| Figure 73: Relative width of 95% confidence interval for American call option prices on underlying following a geometric Brownian motion..... | 319 |
| Figure 74: Relative error of European real call option prices with underlying following a Merton jump diffusion process (reference modified Merton-Black Scholes formula)... | 325 |
| Figure 75: Relative width of 95% confidence interval for European real call option prices on underlying following a Merton jump-diffusion process..... | 326 |

| | |
|--|-----|
| Figure 76: Relative width of 95% confidence interval for American call option prices on underlying following a Merton jump diffusion process | 327 |
| Figure 77: Relative error of compound European-European call option prices with underlying following a geometric Brownian motion (reference Geske formula)..... | 335 |
| Figure 78: Relative width of 95% confidence interval for compound European-European call option prices on underlying following a geometric Brownian motion | 335 |
| Figure 79: Development timeline for PIP-Light and PIP-Involved..... | 339 |
| Figure 80: PIP development strategies in a competitive environment | 345 |
| Figure 81: Aircraft evaluation methodology..... | 348 |
| Figure 82: Network Analysis | 350 |
| Figure 83: Elements of airframe and heavy component maintenance | 353 |
| Figure 84: Forecasting airframe and heavy components maintenance events..... | 354 |
| Figure 85: Effect of engine wear on the EGT margin | 356 |
| Figure 86: Effect of regular engine wash (a) on the EGT margin erosion and (b) on the fuel-burn degradation as a function of flight cycles | 357 |
| Figure 87: EGT margin erosion for engines operating with low severity (red) and high severity (black) factors..... | 358 |
| Figure 88: Operational severity factor | 359 |
| Figure 89: Environment harshness and impact on severity (adapted from [235])..... | 360 |
| Figure 90: Engine maintenance analysis..... | 362 |
| Figure 91: Forecasting engine maintenance events | 363 |
| Figure 92: Spare part analysis..... | 364 |
| Figure 93: Spare engine requirement computation..... | 365 |

| | |
|--|-----|
| Figure 94: Fuel-burn analysis | 367 |
| Figure 95: Emission taxation analysis | 369 |
| Figure 96: EGT margin erosion is dependent on the thrust rating and severity of operations..... | 372 |
| Figure 97: EGT margin erosion using Weibull probability distributions for the first and mature lives of low thrust engines (exhibits (a) and (b)) and high thrust engines (exhibits (c) and (d)) | 374 |
| Figure 98: Aging process requiring unscheduled engine removals using Weibull probability distribution..... | 376 |
| Figure 99: Comparison of published (+) and computed (\diamond) maintenance costs per flight cycle and per flight hour for different missions..... | 381 |
| Figure 100: Comparison between flight fuel-burn published (+) and computed (\diamond) for different missions..... | 382 |
| Figure 101: Fleet-wide operating cost breakdown..... | 384 |
| Figure 102: Life-cycle cash inflows and outflows..... | 384 |
| Figure 103: Fleet-wide EGT margin degradation (red, decreasing) and SFC degradation (blue, increasing)..... | 385 |
| Figure 104: Engine Shop Visit Rates and engine dispatch availability with two spare engines | 386 |
| Figure 105: Similar aircraft, widely different flight operations..... | 388 |
| Figure 106: Typical narrowbody utilization (a) and yearly utilization (source <i>Boeing</i> [270]) (b)..... | 389 |
| Figure 107: Probability of choice between two alternatives with attributes z_1 and z_2 | 391 |

| | |
|---|-----|
| Figure 108: Yearly deliveries in (a) and survival curve in (b) for the Airbus A320 family of aircraft..... | 393 |
| Figure 109: Yearly deliveries in (a) and survival curves in (b) for the 737 Classic family and MD80 family | 394 |
| Figure 110: Calibration of the target market size shrink over time | 395 |
| Figure 111: Closing price of jet fuel (left graph); Continuously compounded daily jet-fuel price returns (right graph) | 397 |
| Figure 112: Distribution of daily jet-fuel price returns..... | 397 |
| Figure 113: Jet-fuel futures quotes and stochastic process retained | 400 |
| Figure 114: Closing price of EUA (left graph); Continuously compounded daily returns (right graph) | 400 |
| Figure 115: Distribution of daily EUA price returns | 401 |
| Figure 116: E.U. Allowance futures quotes and stochastic process retained for subsequent analyses..... | 402 |
| Figure 117: Baseline PIP development timeline..... | 409 |
| Figure 118: PIP- <i>Light</i> impact on (a) fuel-burn, (b) CO ₂ emissions, (c) maintenance costs and (table) expected operating life..... | 410 |
| Figure 119: Limited-flexibility PIP development timeline..... | 413 |
| Figure 120: Flexible PIP development timeline | 415 |
| Figure 121: Fully flexible PIP development timeline..... | 418 |
| Figure 122: Trigger boundary for PIP- <i>Light</i> development program | 420 |
| Figure 123: Trigger boundaries for the PIP- <i>Light</i> with jet-fuel price following a geometric Brownian motion and a jump-diffusion process | 423 |

| | |
|--|-----|
| Figure 124: Trigger boundaries for the PIP-Light retrofit with (black) and without (red) carbon emission allowances..... | 424 |
| Figure 125: Comparison of trigger boundaries for PIP-Light(1) and PIP-Involved(2) and comparison of path-dependant option prices for the PIP-Light and PIP-Involved..... | 425 |
| Figure 126: Sequential scenarios under investigation and selection of the Nash equilibrium..... | 429 |
| Figure 127: Expected return of a portfolio with two perfectly and negatively correlated assets | 487 |
| Figure 128: Expected return of a portfolio with two and more correlated assets..... | 488 |
| Figure 129: Expected returns and volatilities for portfolios featuring a risk-free asset and many correlated assets | 490 |
| Figure 130: Capital Market Line (CML) and Security Market Line (SML) | 492 |
| Figure 131: FLAVIA interface for technical parameters specification | 506 |
| Figure 132: FLAVIA interface for uncertainties and program timeline specifications.. | 507 |
| Figure 133: Classes implementation for the real options tool | 510 |
| Figure 134: Forward (a) and backward (b) finite-difference for the explicit and implicit schemes..... | 511 |
| Figure 135: Required number of time steps to meet target accuracy | 516 |

LIST OF SYMBOLS AND ABBREVIATIONS

| | |
|------------------|---|
| CAPM | Capital Asset Pricing Model |
| CDF | Cumulative Distribution Function |
| $E_{\mathbb{P}}$ | Expectation under physical probability measure |
| $E_{\mathbb{Q}}$ | Expectation under risk-neutral probability measure |
| ETOPS | Extended operations |
| f | Probability density function |
| f_{Ess} | Esscher-transformed probability density function |
| \mathcal{F} | Sigma algebra |
| FAA | Federal Aviation Administration |
| FH:FC | Flight hours to flight cycles ratio |
| h | Esscher parameter |
| h^* | Esscher parameter for risk-neutral Esscher transformation |
| IPPD | Integrated Product and Process Development |
| IRMA | Interactive Reconfigurable Matrix of Alternatives |
| K | Strike price of option |

| | |
|--------------|---|
| M | Moment generating function |
| \hat{M} | Empirical moment generating function |
| NPV | Net Present Value |
| P | Payoff function for option |
| \mathbb{P} | Physical probability measure |
| \mathbb{Q} | Risk-neutral probability measure |
| r | Discount rate |
| r_f | Risk-free interest rate |
| r_i | Rate of return of i^{th} security in portfolio |
| r_{irr} | Internal rate of return |
| r_M | Rate of return of market portfolio |
| R&D | Research and Development |
| S | Asset price |
| \bar{S} | Long-term mean asset price |
| S_{t_k} | Asset price at time index t_k |
| t | Time |

| | |
|-----------|---|
| t_k | Time index |
| T | Option maturity |
| V | Option price |
| V_{t_k} | Option price at time index t_k |
| W | Wiener process |
| WACC | Weighted Average Cost of Capital |
| w_i | Share of i^{th} security in portfolio |
| Z | Radon-Nikodym derivative |
| β | Beta parameter |
| γ | Market price of risk |
| η | Speed of mean reversion |
| μ | Drift of stochastic process |
| σ | Volatility of stochastic process |
| ϕ_i | Basis function |
| ϕ_t | Share of asset security in replicating portfolio |
| ψ_t | Share of risk-free bond security in replicating portfolio |

SUMMARY

The developments of new technologies for commercial aviation involve significant risk for technologists as these programs are often driven by fixed assumptions regarding future airline needs, while being subject to many uncertainties at the technical and market levels. The effect of these uncertainties is further compounded by the fact that development programs are long and uncertainties continue to evolve after the aircraft and engine designs are frozen. Despite this overwhelming uncertainty, technologists must still assess the economic viability of these development programs. Unfortunately, standard methods used for capital budgeting are not well suited to handle the uncertainty surrounding such developments.

In this Ph.D. research, a novel methodology is formulated for the analysis of research and development (R&D) programs. This research is motivated by three observations: (1) integrating competitive aspects such as strategy selection in a competitive environment early in the design process ensures that development programs are robust with regards to moves by the competition; (2) disregarding managerial flexibility undervalues many long-term and uncertain research and development programs; and (3) windows of opportunities emerge and disappear, and manufacturers could derive significant value by exploiting their upside potential.

The main objective of this work is therefore to answer the following overarching question: “*Within the context of aerospace research and development optimization, how can value-based design methodologies be improved to identify precursors of technological and market opportunities?*” In addition, the improvements need to be able

to “reflect the specific challenges associated with long-term and uncertainty-plagued aircraft and engine developments, and to account for the competitive nature of the business”.

In order to overcome these challenges, a method based on real options analysis and cross-fertilizing different techniques borrowed from the fields of quantitative finance, actuarial sciences, and statistics is proposed to study the timing of staggered investments under uncertainty and competitive pressure. Real options analyses have been proposed in the past to address some of these points but the adoption has been slow, hindered by constraining frameworks and unrealistic assumptions. In a symposium held at Georgetown University, a panel of academics and practitioners has identified a set of requirements, known as the *Georgetown Challenge*, that real options analyses must meet in order to get more traction and wider acceptance amongst practitioners. In a bid to meet some of these requirements, this research aims at proposing a method to help substantiate decision making for R&D while having a wider domain of application and an improved ability to handle a complex reality compared to more traditional approaches.

The method, named *FLexible AViation Investment Analysis* (FLAVIA) aims at addressing and bridging two gaps identified in traditional capital budgeting techniques: (1) the evaluation of long-term technology development programs featuring decision tollgates in the presence of significant market uncertainty; (2) the generation of trigger boundaries at decision tollgates of a development program to help decision-makers identify trigger events of successful developments and to substantiate investment policies. Besides these analyses, various investigations may also be performed using the FLAVIA economic evaluation platform such as sensitivity studies to understand how

robust investment policies defined using the trigger boundary are with respect to competitive and technical perturbations.

The proposed method is based on (1) a Monte Carlo technique to value research and development programs with early-investment possibilities, (2) a non-parametric exponential tilting of probability distributions to express the evolution of investment revenues in a different but equivalent probability measure, and (3) a bootstrap resampling technique to generate trajectories representing the evolution of the research and development program revenues over time.

The proposed methodology builds upon real options techniques that have been proposed over the years and has been developed with practitioners in mind: it provides analysts with a clear and transparent process to perform staggered investment evaluations. First, the method uses widely accepted Monte Carlo simulations in order to handle multiple, possibly correlated, uncertainties. Monte Carlo simulations also offer a rich environment enabling analysts to use non-traditional stochastic processes that may be better suited to model a complex reality. Next, stochastic models representing the evolution of uncertainties over time are calibrated under the physical probability measure since this is the natural measure to use to calibrate models with any source of market data available. Because the simulation of multiple and possibly correlated uncertainty models may result in a development program revenue process with unknown characteristics, a non-parametric probability measure transformation is suggested using the time-honored Esscher transform. This tiling enables a simple and transparent transformation of the development program revenue process from the physical probability measure to the equivalent martingale probability measure typically used for option pricing purposes. The

development program revenue process, now simulated under the equivalent martingale measure, is resampled to generate non-weighted trajectories representing its evolution over time under the new measure. This enables the use of the Longstaff-Schwartz least-squares Monte Carlo algorithm to both generate the trigger boundary and estimate the development program value.

The FLAVIA method is first subjected to preliminary testing on a set of canonical examples so as to check the quality of the results before moving on to a more exhaustive verification and validation. Preliminary testing indicates that the estimations of real option values are accurate but that the generation of the trigger boundary suffers from excessive noise, rendering its use for decision-making questionable. As a result, several techniques are suggested to improve the method. These include (1) the use of moment matching and control variates sampled at the exercise of the path-dependent real options to reduce the variability in the results, (2) the scoping of the continuation value regressions used in the least-squares Monte Carlo algorithm to improve the quality of regressions, (3) a three-step filtering and regression of the locus of critical prices which constitutes the trigger boundary, (5) the use of Sobol's low-discrepancy sequences in lieu of pseudo-random numbers in (Quasi-) Monte Carlo simulations, and finally (4) the use of a two-stage least-squares Monte Carlo algorithm using a newly proposed Multi-Start Monte Carlo simulation.

All, but one, of the techniques drastically improve the accuracy of the results and significantly reduce the noise in the trigger boundary. Sobol's low-discrepancy sequences exhibit disappointing results however and are abandoned. With the successful

implementation of these refinements, the FLAVIA method is first verified extensively and then validated using a proof-of-concept application.

The proof-of-concept application is a Performance Improvement Package (PIP) development. The package is a set of technologies that can be retrofitted on a currently operating turbofan engine. The PIP development timeline is staggered with several phases and the value of the development program is driven to a large extent by the volatile jet-fuel price and to a lesser extent by the uncertain price of carbon emission allowances. In order to perform the economic evaluation of the PIP, the benefits to the operators (airlines) are quantified using the newly-developed *Integrated Cost And Revenue Estimation* method (i-CARE). Both the i-CARE and FLAVIA methods are implemented and linked together to analyze the PIP development program.

Results indicate that both managerial and timing flexibility have statistically significant values which are unaccounted for in typical discounted cash flow analyses. Trigger boundaries are generated and highlight the difference between a discounted cash flow-based investment policy and an optimal real option-based investment policy. The research concludes with several suggested improvements and extensions to the FLAVIA method.

CHAPTER 1: INTRODUCTION

Whether it be in a backyard workshop or inside a bicycle shop, neither Clément Ader nor Otto Lilienthal or the Wright Brothers probably understood the impact their flying machine inventions would have over mankind. Over the course of a century, private ventures led by self-driven individuals have evolved and given birth to a whole new industry at the forefront of technical and technological innovation while generating millions of jobs worldwide and drastically shifting the paradigms of travel and commerce. As the demand for air transportation grew, the requirements for faster and more reliable vehicles slowly evolved into requirements for more efficient and more environmentally-friendly vehicles. At the same time, the whole aerospace industry transformed itself into a complex mesh of stakeholders including manufacturers, suppliers, regulators, ground facility providers, maintenance providers, air traffic controllers, air carriers, and finally air travelers.

In its 2012 surveys [1], the *Aerospace Industry Association* (AIA) reports that, in the United States alone, the aerospace industry as a whole is employing 629,000 workers (including 430,000 directly related to the design and manufacture of aircraft, engines, and parts), has revenues exceeding 217 billion dollars (of which 118 billion dollars are aircraft-related with a backlog of over 2,600 aircraft), generates profits in excess of 20 billion dollars, and finally exhibits the highest positive trade balance of all major American industries with a surplus estimated over 63 billion dollars. Such glittering statistics is bound to attract some attention.

1.1 Aircraft manufacturing industry

The aerospace industry is very diverse. It encompasses a multitude of commercial, industrial, and military applications and spans a wide range of interests stretching from the design and manufacture, to the operation and maintenance of vehicles moving within the Earth atmosphere or further away in space. In the remainder of this document, the focus will be mostly on civil commercial application within the aircraft manufacturing industry.

Over the course of the first half of the XXth century, the aircraft manufacturing industry has sailed through its infancy. It was characterized by a multitude of manufacturers, each coming up with its own unique and sometimes extravagant design. During the second half of the XXth century, the young industry was characterized by a frantic effort to extend the overall flying envelope and to investigate the entire design space of aerial vehicles including specialized ones with the ability to hover or to fly at hypersonic speeds. By the turn of the new century, the industry has reached a mature status whereby most of the aircraft manufacturers have converged to a single aircraft configuration, the tube and wing conventional configuration. Nowadays, each new aircraft design iteration seems to consist mostly of further refinements and optimizations of this very configuration.

While the industry was going through these stages, the resources required to design, fly, and certify new aircraft kept growing. The investments required to fund the development of new aircraft forced the industry to transform itself. Many family-run private manufacturers were unable to cope with these radical changes and were forced to either merge and survive (*Breguet Aviation* and *Dassault Aviation*) or disappear (*de*

Havilland Aircraft Company). Later on, as the industry was maturing, much larger manufacturers were also forced to either merge (*McDonnell Douglas* and *Boeing*) or close doors (*Fokker*). This progressively led to a consolidation of the industry and eventually resulted in an industry dominated by two manufacturers of larger commercial aircraft, three manufacturers of smaller regional aircraft, and a handful of manufacturers of general aviation aircraft. At the lower end of the commercial aviation spectrum, *ATR*, *Bombardier*, and *Embraer* are registering most of the orders for aircraft with a capacity under one hundred seats, whereas at the other end of the commercial aviation spectrum, *Airbus* and *Boeing* are registering most of the orders for aircraft with a capacity greater than one hundred seats. On the general aviation side, *Cessna*, *Bombardier*, *Gulfstream*, and *Dassault* are taking most of the orders for larger private jets, while *Cessna*, *Cirrus*, and *Diamond* are getting most of the orders for smaller private aircraft.

1.2 Aircraft design process

Hazelrigg defines the design activity as the “*Use of available information to make intelligent decisions leading to optimal solutions.*” A dissection of this definition yields several important notions that are going to be analyzed next in the context of aircraft design.

The term ‘available information’ refers to the fact that designers are not scientists but rather engineers. Thus, their goal is not to push the state-of-the-art in terms of scientific knowledge. Instead, they strive to monitor, review, and finally use relevant pieces of knowledge to establish models for the purpose of sizing and designing an aircraft.

The term ‘intelligent’ refers to the fact that the models used are neither trivial nor black boxes. Making an intelligent decision requires some understanding of the underlying physics as well as of the data underpinning these models. In turn, this implies having an understanding of the limitations and suitability of these models for a specific application. This also requires a good understanding of their inputs and outputs, and of their respective level of uncertainty and accuracy.

The term ‘decision’ implies the existence of choices that designers must substantiate when confronted with sets of alternatives. Often, there is more than a single design that meets the requirements set forth by the end-customers, but some are better than others. However, designers do not always have a detailed understanding of the impacts of their design choices. Similarly, they may not have a definitive set of operating conditions for their design as requirements creep and design missions evolve. There is therefore a need to make and substantiate decisions with incomplete and uncertain knowledge.

Finally, the term ‘optimal solutions’ implies the presence of trade-offs and of non-trivial answers to the task to be completed. By nature, the aircraft design environment is multi-disciplinary, encompassing disciplines as diverse as aerodynamics, control theory, structural engineering, material sciences, chemical engineering, and many others. Since aircraft are more and more complex, they need to be viewed as systems-of-systems with conflicting requirements and constraints. Because a system-of-systems is usually greater than the sum of its systems, each system shall not be optimized independently of the others but rather concurrently with all the others. The essence of a designer’s task is

consequently to provide a synthesis through the means of multi-disciplinary optimizations.

1.2.1 Different phases of design

Having defined what the design activity involves, it should be clear that designing an aircraft is a complex multi-step and multi-disciplinary process. As such, the aircraft design process is usually decomposed into three phases: the conceptual design, the preliminary design, and the detailed design. Together with the flight testing and certification phase, these four phases form the aircraft development process.

During the conceptual design phase, the customer needs and requirements are first analyzed to answer some basic questions regarding the configuration of the design, the technologies to be used, the first weight estimations, and some primary analysis of the economic viability of the aircraft. Design requirements may include aircraft capacity, aircraft range, payload, take-off and landing distances, as well as emissions and noise thresholds. Using these design requirements, a sizing and synthesis of the aircraft is done next to ensure that requirements are met. An increasing number of trade-off analyses are performed to get more and more knowledge about the design configuration and its performance. Raymer [2] argues that the design at this stage is very fluid, evolving by the week as more and more sophisticated analyses are performed for every aspect of the design. Parametric studies and design space explorations are used to speed-up analyses in such a fluid design environment and to ensure that different alternatives are reviewed before a preferred design is eventually selected.

The next phase of the design process is the preliminary design phase which usually starts whenever major changes to the aircraft configuration are unlikely to occur.

It is characterized by a maturation of the selected design. During the preliminary design phase, experts from all disciplines are involved to design and analyze their part of the aircraft and to prepare for the detailed design stage. The preliminary design phase ends when the configuration of the aircraft is frozen and when experts have established confidence that the aircraft can be built on time, according to specifications, and for the projected budget.

The last phase of the aircraft design process is the detailed design phase during which the full scale development of the aircraft takes place. This is the most expensive part of the design process as detailed models are created for each and every part of the aircraft. Many sub-assemblies are created and tested. This is also the phase during which the production design takes place. During production design, engineers determine how the aircraft will be fabricated and what production tooling will be required. The detailed design phase ends with the production of the very first aircraft.

According to Schrage [3], the traditional development process can be described as shown in Figure 1. Looking more closely at the different phases of this design process, it appears that the commitment required from aircraft manufacturers increases with time. Indeed, Raymer [2] indicates that conceptual design may stretch from a couple of weeks up to half a year while preliminary design may stretch from a couple of months up to two years and detailed design may last for over two years. At the same time, as the design progresses through these different phases, the number of people involved keeps increasing, and the analyses carried out get deeper and deeper as more and more fidelity is required.

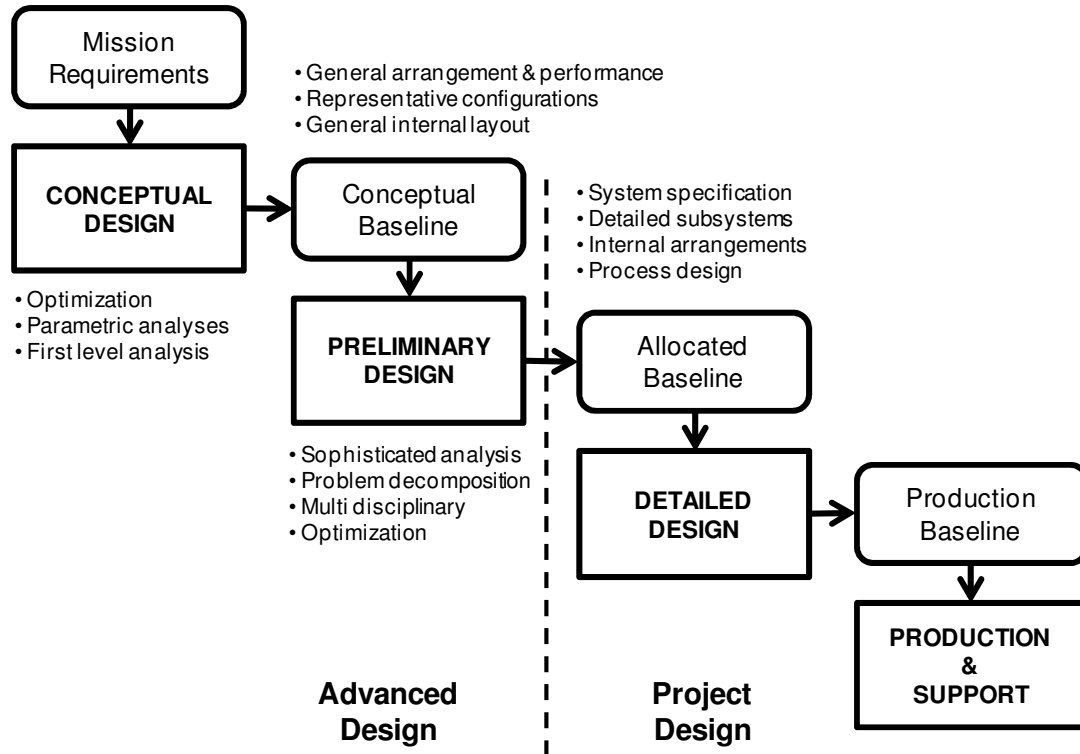


Figure 1: Traditional Development Process (adapted from [3])

1.2.2 Modern system design methods

By definition, aircraft and engine design is a multi-disciplinary endeavor. Over the years, Original Equipment Manufacturers (OEM) have developed or borrowed various systems engineering methods and processes to help them in their design tasks. A full review of all these methods and tasks is beyond the scope of this section but several of the most relevant methods are presented in the following paragraphs.

Integrated Product and Process Development

It is well known that the freedom to alter a design decreases substantially as it matures from a concept blueprint to a full scale production. At the same time, the costs induced by changes increase significantly as the design matures. In this context,

manufacturing decisions cannot wait for the design to be frozen. Tightly coupled design and manufacturing decisions must be made concurrently early in the design using the concurrent engineering approach defined as a “*systematic approach to the integrated, concurrent design of product and their related processes, including manufacture and support*” [4]. This approach has emerged as an effective way to drive Total Quality Management in each stage of a product life-cycle. Its underpinning philosophy is to bring together experts from the different phases of both the product (aerodynamics, propulsion, structure) and the manufacturing (producibility, supportability) processes early-on with the ultimate goal of minimizing the overall life-cycle cost of the product designed.

The Integrated Product and Process Development (IPPD) methodology was developed to implement this concurrent engineering approach for the purpose of aerospace design. This methodology, described by Marx. et al. [5], is illustrated in Figure 2 and depicts the interactions of four key elements to enable parallel product and process trades to be made: systems engineering methods, quality engineering methods, a top-down design decision support process, and a computer integrated environment. Beneath these, the interactions necessary to perform the parallel product and process trades are described. Using this methodology, knowledge is brought forward in the design process yielding greater flexibility to decision-makers. Higher fidelity trade-off analyses can then be performed early-on in the design to leverage the inexpensive design freedom still available. Indeed, the system synthesis is performed using a multi-disciplinary design optimization environment in which sets of alternatives are generated. For each of these feasible alternatives, an adequate risk assessment and uncertainty analysis are performed early-on with regards to performance, costs, and schedule.

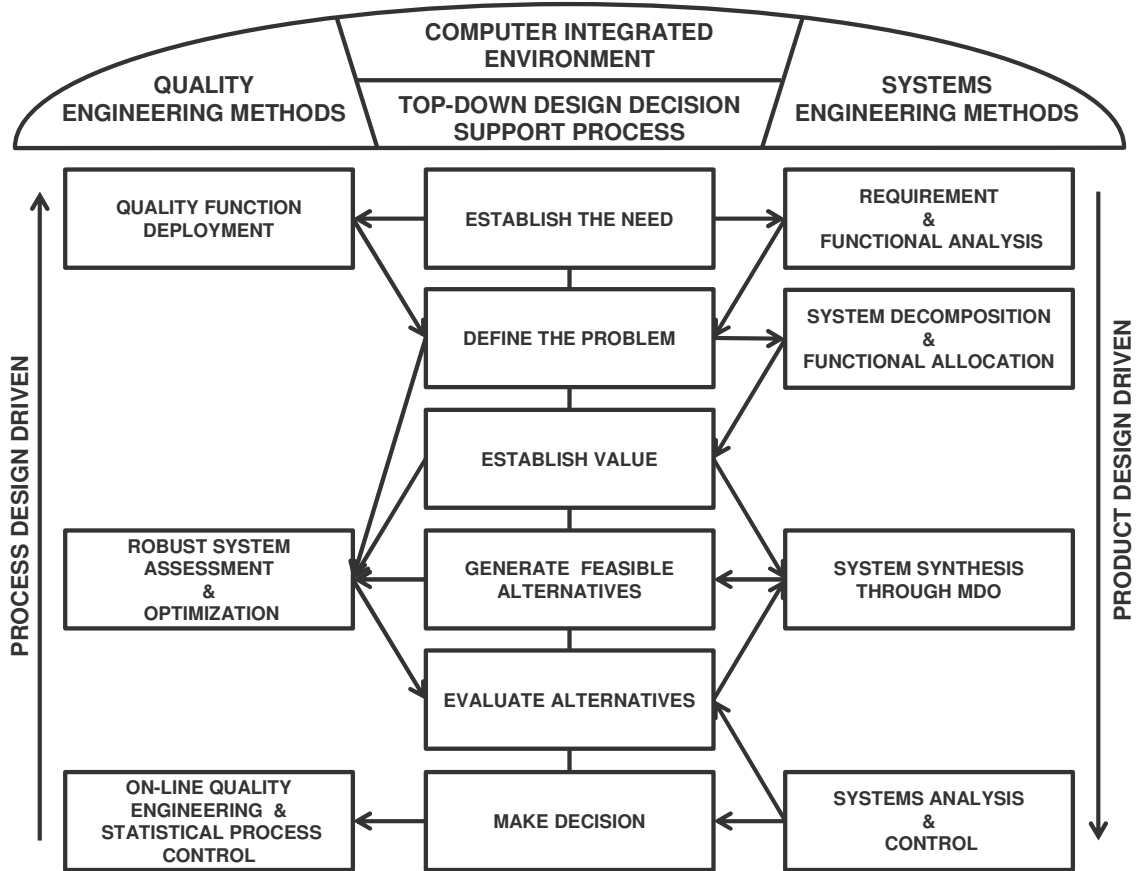


Figure 2: Georgia Tech IPPD methodology adapted from Marx et al. [5]

Design Space Exploration

The purpose of the design space exploration is to keep an open-minded approach to design and therefore to explore a wide range of variables and design choices. This process refers to the activity of exploring design alternatives prior to implementation. It involves the systematic evaluation of feasibility and viability of all available design choices to generate a down-selection of feasible alternatives. Some methods and models to evaluate the feasibility and viability of these design choices are therefore required. This exploration may be done at two different levels: a qualitative level which relates more to design choices and a quantitative level which relates more to the value assigned to variables associated with these design choices. To tackle the first level, the Interactive

Reconfigurable Matrix of Alternative (IRMA) [6] has been developed and proposes a systematic qualitative procedure to decompose the design, identify alternatives, and check for compatibilities. To tackle the second level, Monte Carlo simulations may be used to investigate design viability and overall vehicle performance for different levels of the design variables. This is performed by treating design variables as random variables and assigning distributions to them (usually uniform distributions at this stage) instead of single point estimates. This process is illustrated in Figure 3.

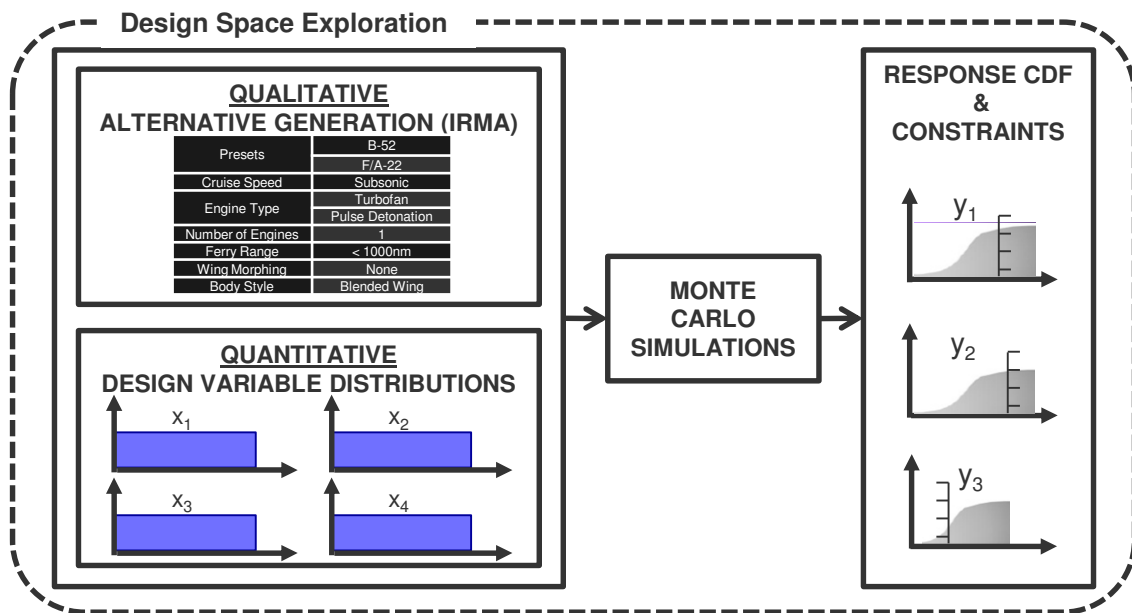


Figure 3: Design space exploration

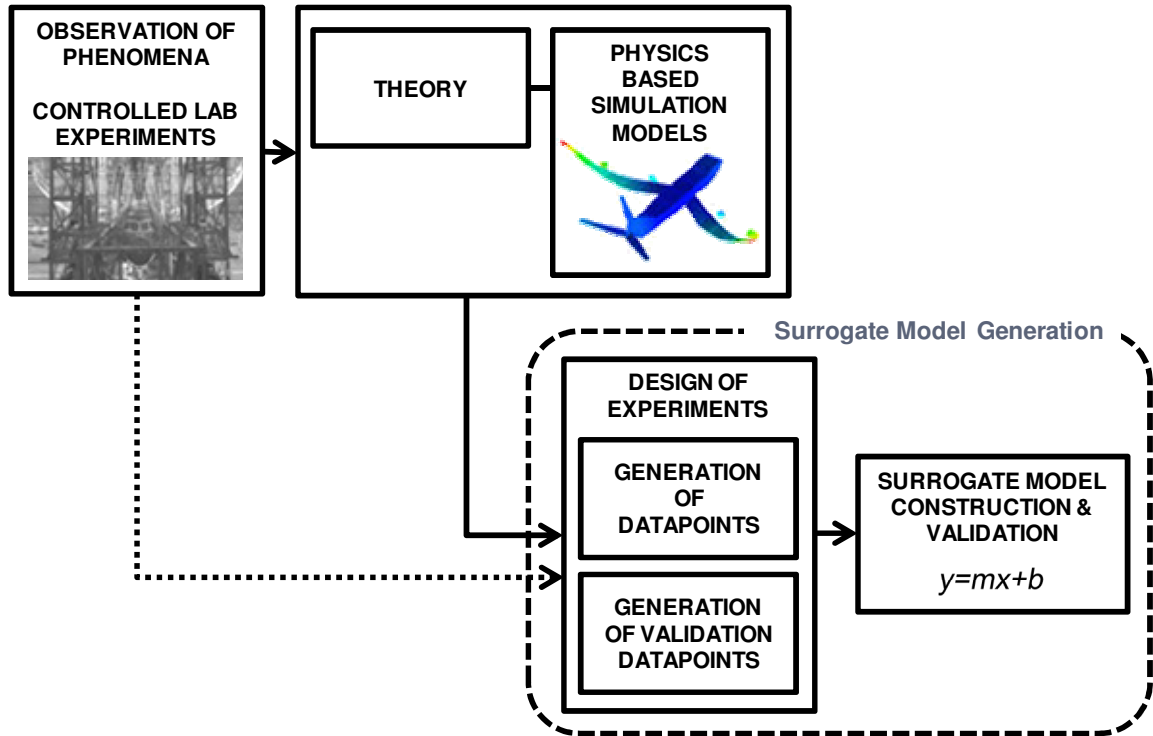
The main challenge of design space exploration arises from the sheer size of the design space that must be explored. Indeed, aircraft are large systems-of-systems defined by numerous design variables which may have wide ranges of values. Therefore, there may be millions of different possibilities to explore and enumerating each of them for further analysis may be prohibitive. If not for a reduction of the design space, there is at least a need to speed-up the analysis carried out for each design point. In addition, at this stage of the design, it is not obvious that higher fidelity physics-based models are

required. Instead, faster surrogate models or metamodels may be used to speed-up the Monte Carlo simulations and perform a better sampling, and consequently a better exploration, of the design space.

Surrogate Modeling

In the quest to reduce uncertainty and to mitigate risks, improvements to current physics-based models keep being made. These result in more accurate predictions of performance estimates of future aircraft and engine designs. These simulation models reflect the advancement of science and the greater knowledge and understanding engineers have of the environment they are working on. As these models achieve higher and higher fidelity, they usually grow in complexity and are therefore an impediment to the fast evaluation of alternatives as required for the design space exploration previously described. Indeed, routine tasks such as sensitivity analysis, design optimization, and what-if analysis may become impossible because they require millions of simulation evaluations and each of these evaluations may take hours or even days to run and to fully converge.

Surrogate modeling designates the activity of building models of models in order to speed-up analyses and to alleviate the burden of running many times the same high fidelity physics-based models. Surrogate models are constructed using a data-driven approach in which results from the higher fidelity models are used to yield simpler and more manageable models while still retaining most of the accuracy of the underlying physics-based models.



constructed. In this case, results from the surrogate model are directly compared to data provided by lab experiments or higher fidelity physics-based models.

Robust Design Simulation

The purpose of robust design simulation is to ensure that the design under review is robust with regards to external and internal perturbations. The emphasis on robust design comes hand in hand with the philosophy of quality planning and Taguchi's observations [7] that it is often cheaper to make a process insensitive to manufacturing variations than to control the causes of these variations. Robustness analysis helps by providing an estimation of the sensitivity of outputs to the variability of inputs described in terms of random variables and probabilistic distributions.

In the IPPD methodology presented earlier, Robust Design Simulation is performed using surrogate models of physics-based models to speed-up sensitivity analyses along with Monte Carlo simulations to model the variability of inputs. Indeed, probabilistic distributions are used to model uncertain parameters which abound during the design of aircraft and engines. Technology benefits for instance are better assessed by ranges rather than by single number estimates because they have yet to be corroborated and tested in the operating environment. Economic inputs, such as energy prices, may also be well served by ranges because they usually exhibit substantial volatility and change drastically over time. Probabilistic assessments are carried out as shown in Figure 5 to ensure that the design satisfies the customer needs and meets all constraints and requirements in an uncertain environment with uncertain technology benefits. According to Mavris and Bandte [8], the end-purpose of robust design simulation is to come up with

“a design that performs well in the environment for which it was designed but also in all environments”.

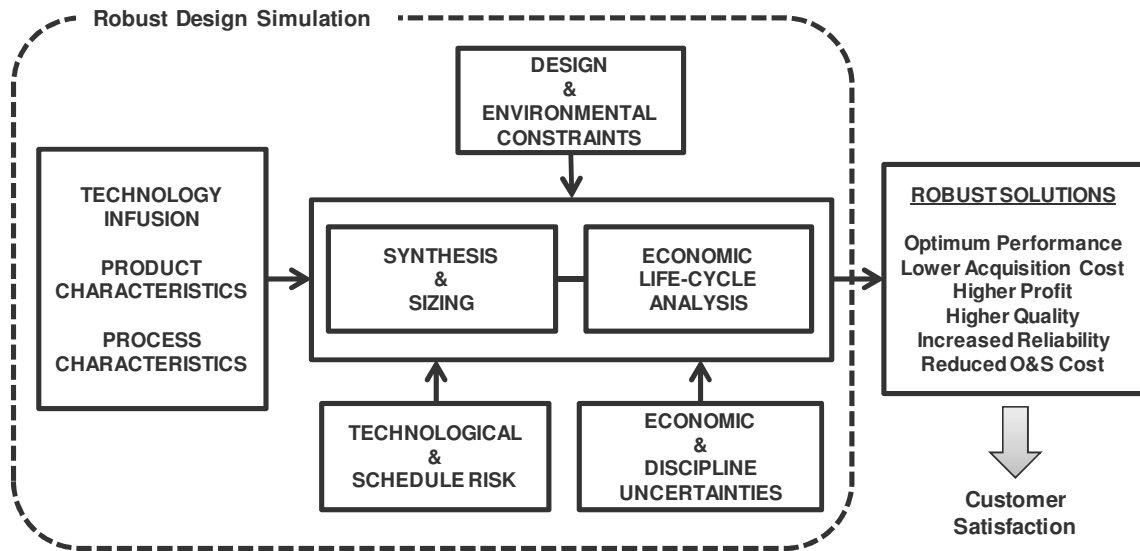


Figure 5: Robust Design Simulation (adapted from Mavris and Bandte [8])

Value-Driven Design

Value-driven design is a systems engineering framework where design choices are made to maximize the overall system value instead of to simply meet performance requirements. Value-driven design therefore creates an environment that enables design optimization by providing designers with an objective function while eliminating constraints expressed as performance requirements [9]. In fact, the objective function itself is already incorporating the important attributes and therefore implicitly the constraints, and outputs a score which is the overall value of the system. In that regards, value-driven design is quite close to multi-attribute decision making techniques because it aims at maximizing a single-objective value function. However, instead of using stakeholders to define marginal utility or sub-utility functions, it uses microeconomic analyses to yield a figure of merit which is usually a profit function expressed in a

monetary unit. This approach allows designers to compare different designs in terms of value scores and therefore to rationally make design decisions.

1.3 Traditional challenges

The design and production of aircraft is a complex and challenging process encompassing many disciplines ranging from market research to mechanical engineering and supply chain management. This process has traditionally been riddled with uncertainties and these uncertainties create many different types of risks spanning from market risks to technical risks, supply chain risks, and schedule risks. The overwhelming number of “things that may go wrong” during the design stage is probably one of the reasons so many aircraft designs have been struggling to become profitable [10] and the reason some of the newest aircraft developments have seen their development schedule slip over time [11].

1.3.1 Major investments and limited returns

One of the main challenges facing aircraft manufacturers is the unsustainable development costs required to fully research, develop, certify, and produce new aircraft. Although these investments are significant barriers to entry for new competitors and are therefore preserving the market for established players, development costs upward of ten billion dollars are in fact betting the future of the airframe and engine manufacturers as shown in Table 1. On the engine manufacturing side for instance, *Rolls-Royce* went into administrative receivership in 1971 after encountering technical problems while designing the new triple spool RB211 turbofan for the *Lockheed* L-1011 Tristar. This led to production delays which, in conjunction with limited sales, jeopardized the future of

the Tristar program. Although the technologies developed by *Rolls-Royce* for the RB211 paved the way for future commercial successes with the Trent family of triple spool high-bypass turbofan engines, the unprofitable Tristar development program ended in 1984 with the delivery of just two hundred and fifty aircraft.

Besides the high investments required, it is not uncommon to see development cost overruns as the development programs get into the detailed design, certification, and production phases during which technical issues may be encountered. In such cases, some parts need to be redesigned and the development schedules tend to slip. Similarly, suppliers of large aircraft subsystems may not have the ability to ramp-up production of complex parts as fast as airframe manufacturer initially projected [12] leading to additional delays [13]. As may be seen in Table 1, most of the new programs of this past decade have experienced some form of delays at entry into service which have resulted in drastic increases of already high development costs.

Table 1: Aircraft development costs: projected and final estimate at completion

| Aircraft Model (* denotes derivative) | Projected Development Costs [14] Billion 2012-US\$ | Final Development Costs [15] [16] [17] Billion 2012-US\$ |
|--|---|---|
| <i>Airbus A320 NEO*</i> | 2.0 | / |
| <i>Airbus A350-A350XWB</i> | 6.5 | ~15.0 |
| <i>Airbus A380</i> | 12.0 | ~ 15.0 |
| <i>Boeing 737 Max*</i> | 3.0 | / |
| <i>Boeing 787</i> | 8.0 | ~ 33.0 |
| <i>Boeing 777</i> | 6.0 | ~ 6.0 |
| <i>Bombardier C-Series</i> | 3.5 | ~5.2 |

Digging further into the high development costs and the uncertainty surrounding these programs, another issue seems to be the inability of aircraft manufacturers to

command prices that truly reflect the incremental value of the new aircraft relative to prior-generation products. Indeed, a substantial number of customers purchase these aircraft during the initial launch phase of a project once the manufacturer gets the authorization to offer the aircraft from its board, but before the development is completed. In addition, these pre-orders are often sold at a significant discount over the catalogue price in order to build momentum and gain market share. At this point, the airframe manufacturers are relying on costs projections. However, the business plan may underestimate development costs and program risks, leading to a significant number of airframes being sold at a loss or with unsustainable margins.

1.3.2 Technical and technological risks

Another challenge for aircraft and engine manufacturers is the management of technical and technological risks. Indeed, aircraft and engine manufacturers are at the forefront of technical and technological innovation, and pressure from both customers and the competition to design more and more efficient aircraft with shorter and shorter lead time force them to embrace new design methods, new manufacturing processes, and new technologies for each new aircraft. On the one hand, technical risks may be defined as an exposure to losses arising from the design and the manufacturing activities and is therefore related to the processes used by the aircraft manufacturer. The *Airbus A380* for instance suffered major setbacks during its development because different plants within the company were using incompatible versions of the same software preventing them to update in real time the common digital mock-up of the aircraft [18]. On the other hand, technological risks may be defined as the exposure to losses arising from the infusion of new technologies into the aircraft design. It is therefore related to the maturity level of a

technology and the level of experience the manufacturer has with using this technology. The *Federal Aviation Administration* (FAA) for instance announced a fleet-wide grounding and a comprehensive review of the *Boeing 787* critical systems after airlines reported issues and severe thermal runaways with the new lithium-ion batteries. At the end of 2014, the root-cause of these thermal runaways was still unknown and only palliative solutions had been provided by the manufacturer [19]. This is a prime example of technological risks taken during design by incorporating technologies not yet fully mature.

Closely related to the technical and technological risks are the performance risks which may be defined as the exposure of the manufacturer to losses stemming from unreached performance targets by the aircraft. Justin and Mavris [20] report that it is common for original equipment manufacturers to offer performance guarantees as well as “power-by-the-hour” type of maintenance contracts. These contracts are becoming very popular for engines: they allow the operators to pay a fixed price, set in advance, to cover the maintenance expenditures. The purpose of these contracts is to shift some of the operating risks from the airlines back to the manufacturers. As a consequence, original equipment manufacturers are now more than ever exposed to these performance risks: if the performance targets are not met, the manufacturers are liable in terms of compensations and penalties. In addition, if the operating costs of the aircraft are within specifications but higher than the design targets, they might still be exposed to some risks through the maintenance contracts signed by the manufacturers.

1.3.3 Timescale to bring new aircraft to the market

Another great challenge for manufacturers is related to the timescales involved. Aircraft developments are usually measured in years, typically ranging from two to three years for derivative aircraft and from five to eight years for brand new designs. Beyond the development phase, aircraft may remain in production for ten to twenty years leading to an overall program horizon of up to thirty years. However, the world does not remain static meanwhile and the needs, requirements, and regulations evolve. On the evolving requirements side, it is well accepted that the main reason for the commercial failure of *Concord* was the steep raise in energy prices after the 1973 oil crisis leading to unsustainable supersonic operations and reduced airlines' interest. On the evolving regulations side, the commercial success of the four-engine *Airbus A340* was substantially hampered by the evolving Extended Operations (ETOPS) regulations set forth by the *Federal Aviation Administration* that initially constrained twin-engine aircraft operations to airspaces close to diversion airports. As these regulations were extended, new certification levels were created allowing more and more routes to be flown by more efficient twin-engine aircraft and the need for four-engine aircraft progressively vanished.

At the same time, the airline industry is quite volatile with periods of acute crisis, low profitability, and scarce demand for new aircraft followed by periods of rapid expansion with greater profitability and increased demand for new capacity. Liehr et al. [21] argue that one of the root-cause of this cyclical behavior is endogenously generated by the long lead times between aircraft order and delivery, the rush for additional capacity during profitable periods leading to significant aircraft deliveries, and excessive

capacity industry-wide several years down the road. This cyclical behavior is further exacerbated by the elastic nature of the demand for air transportation which contracts significantly during periods of crisis and expands significantly during periods of growth. All in all, the uncertainties surrounding the airline industry coupled with the timescales involved for aircraft development force aircraft and engine manufacturers to speculate during the design stages regarding the size of the market, the future airline needs, the future certification requirements, and more generally, the future states of the world.

1.3.4 Supply chain management

Finally, another source of headaches for producers is related to the management of the supply chain. Over the course of the past decades, the network of suppliers has grown significantly and what used to be a limited number of partners designing some specific systems of the aircraft (such as the engine, the landing gear, and the auxiliary power unit) has evolved into a vast network featuring thousands of suppliers. These suppliers are classified in tiers according to how close they work with the airframer: tier-one suppliers work directly with the airframer, tier-two suppliers work for the tier-one suppliers, and tier-three suppliers work for the tier-two suppliers [11]. Such a complex multi-tiered supply chain spread all over the world calls for a significant increase in managerial oversight on the part of the airframer to ensure that partners deliver on time and within specifications and to ensure delays and costs overrun are minimized. Recent history with the developments of the *Airbus A380* and the *Boeing 787* has shown that this is not a trivial task. For instance, some of the delays of the 787 were attributable to poor oversight by *Boeing* of its network of suppliers [22].

1.4 Emergence of new challenges

1.4.1 Longer development cycles increase risk

In the previous section, the issue regarding the long development and certification times involved in aircraft design was raised. What is troubling is the evolution of these timelines: as aircraft get more and more sophisticated, the time to develop and certify grows substantially, going from an average of four years in the fifties to over eight years for recent designs. Some examples are provided in Table 2. This might sound counterintuitive given the advances made in the past two decades in Computer Aided Design (CAD) and in virtual manufacturing which were supposed to speed-up design tasks. However, at the same time, aircraft have become major interconnected systems-of-systems loaded with intelligent sensors and health monitoring devices. For instance, the *Airbus A380* is equipped with over 100,000 different wires totaling a length of 530 kilometers [23]. Designing such a piece of machinery is quite different from designing a plain metallic tube and wings featuring simple hydraulic systems as was done in the fifties.

At the same time, the world is uncertain, the global airline industry is uncertain, and the price of energy, one of the key aspects of airline profitability, is uncertain. Alan Joyce, CEO of Qantas [24], argues that the unprecedented volatility continues to shake and shape the aviation industry as a whole. However, it would be naïve to believe that this increase in volatility does only affect the aviation industry. Aircraft and engine manufacturers must also adapt themselves and their design processes to handle this volatile environment. For these manufacturers, the long development and certification

timelines in conjunction with the increased market volatility compound the uncertainties and therefore the risks.

Table 2: Aircraft development times

| Aircraft Model (* denotes projected) | Entry Into Service (Date) | Development Time (Years) |
|---|--------------------------------------|-------------------------------------|
| <i>Boeing 707</i> | 1959 | 4 |
| <i>Boeing 747</i> | 1970 | 4 |
| <i>Airbus A320</i> | 1988 | 4 |
| <i>Bombardier CRJ100</i> | 1992 | 4 |
| <i>Boeing 777</i> | 1995 | 5 |
| <i>Airbus A380</i> | 2007 | 7 |
| <i>Embraer E-170</i> | 2004 | 5 |
| <i>Boeing 787</i> | 2011 | 7 |
| <i>Airbus A350</i> | 2015 | 8 |
| <i>Bombardier CS100*</i> | 2016* | 9* |

In this context, it becomes paramount for this industry to be able to react to unforeseen changes in the business environment, to update business plans as uncertainty unfolds, and more generally, to become more flexible.

Observation:

Aircraft and engine developments are characterized by longer and longer development cycles and are therefore subject to significant risk due to the uncertain and volatile business environment. Design methods and design processes must evolve accordingly to provide enough flexibility to managers to steer programs into profitable directions as the uncertainty unfolds.

1.4.2 New competition emerging

Curiously enough, the aircraft manufacturing industry does not exhibit one characteristic trait of mature industries which are often described as having stagnant or declining profits leading to a reduced attractiveness for new entrants in the competition. Indeed, in its 2012 surveys [25], the AIA reports that the net profits of this industry keep increasing despite some cyclical variations and the profit margins increase as well to a lesser extent. Using the example of commercial transport and more particularly the short-to medium-haul single aisle market, Justin et al. [26] claim that after years of declining competition (with the exit of *McDonnell Douglas*, *Tupolev* and *Yakholev*), the competition is drastically increasing with new offerings from established aircraft manufacturers (*Bombardier C-Series*, *Sukhoi SSJ-100*) as well as new designs from new entrants from Russia (*Irkut MS-21*), Japan (*Mitsubishi MRJ*) and China (*Comac C919*). The aerospace industry has been traditionally regarded as an entrenched industry not vulnerable to the threat of new competitors and several reasons are underpinning this paradigm. Historically, manufacturers in Europe, the United States, and the former Soviet Union have been at the forefront of aerospace developments and have therefore accumulated years of experience, skills, and know-how to develop new aircraft. Next, the high barriers of entry into this industry are essentially preventing other competitors from entering the market. Indeed, airlines are expecting a high level of maturity for their newly acquired aircraft and engines at entry into service so as not to add risk to their own operations. These customers are therefore reluctant to place orders to new manufacturers with limited track record of their ability to deliver on time and on specifications. In addition, airlines require the availability of global networks of spare parts and

maintenance facilities in order to limit downtime and schedule disruption when unexpected maintenance arises. This hinders the ability of new competitors to enter the market as these networks are costly to setup and operate.

However, the state of the business is quickly evolving and the competitive landscape is being remodeled for the coming decades. The biggest growths in demand for air transportation are coming from China and India, and experts do not foresee any changes to this in the coming years [27]. As these markets continue to grow, the demand for additional capacity and new aircraft are bound to increase which will probably nurture the aspiration of the homegrown industry to play a larger role in the aircraft and engine developments. In addition, Friedman [28] argues that the educational expertise of Chinese and Indian schools has been comparable to that of Western schools for several years, meaning that an educated and skilled workforce is available to fulfill the aspirations of these homegrown industries.

Besides, it is no secret that there has been a political push at the highest level in many developing countries, and particularly in China, to become more independent of Western and Russian aircraft manufacturers for their booming air transportation needs. In these countries, policies have been set up to help nurture, grow, and develop the local aerospace industry in the hope of letting them compete in the worldwide arena later on. This process may take many shapes but usually involves local manufacturers contracting work from the established aircraft manufacturers (for instance, wings of the *Bombardier C-Series* aircraft are produced by *Shenyang Aircraft* in China). Then, technology transfers are sought whenever large acquisitions are made by state-sponsored entities or technology acquisitions are made through the purchase of established manufacturers

overseas (such as the acquisition of *Cirrus Aircraft* by the *Aviation Industry Corporation of China* holding). Following these acquisitions, local manufacturers aim at subcontracting or licensing the final assembly of aircraft designs of established aircraft manufacturers (e.g., manufacture of *McDonnell Douglas MD-90* by *Shanghai Aircraft* in China and production of *Piper* aircraft in Brazil under the *Embraer* brand). Progressively, this leads to the ability of the local manufacturers to both develop and manufacture self-sufficiently new designs once the technologies are mastered and the skills and processes are well established (e.g. C919 aircraft by the *Commercial Aircraft Corporation of China*).

These recent developments have led to an increase in competitive pressure while it is not obvious that the organic growth of the market will be able to sustain so many manufacturers. In this context, it becomes paramount for the aircraft manufacturers to account for this increase in competition while making business plans and assessing market penetration and profitability, as well as to offer to the market a portfolio of products that both meet the requirements and are differentiable from competing products.

Observation:

Consistent profitability and politics stir up the interest for a homegrown aircraft development industry which leads to a substantial increase in the competitive pressure. Without a corresponding growth in the aircraft demand, manufacturers will need to account for the competitive environment early-on in the design to ensure the business plan is sound and the product portfolio is both competitive and well positioned in the market.

1.5 Challenges to the aerospace industry: a summary

In the previous sections, a brief introduction to the aircraft and engine manufacturing business was provided. Some design processes and methods were presented and have led to the identification of several challenges that are affecting development programs. Table 3 summarizes some of these challenges that were classified as either traditional, for which existing design processes and methods are well adapted, or newly emerging challenges, for which design processes and methods may need to be adapted in order to mitigate potential adverse impacts.

Table 3: Identified challenges to the design process

| Source | Traditional Challenges | Emerging Challenges |
|-----------------------------------|--|--|
| Design | Technical risk Technological risk | |
| Manufacture and Production | Schedule risk Supply chain risks Regulatory risk | Increased reliance on tier-one suppliers |
| Market | Business cycle risk | Competitive threats Increased market volatility |
| Economics | Development cost overruns | Longer development cycles |

CHAPTER 2: MOTIVATION

The purpose of this chapter is to introduce the reader to three different cases of aircraft development programs. The first example analyzes some aspects of the *Airbus* A340-500 and A340-600 development, describes the lackluster sales of these aircraft, and highlights the need for extensive scenario investigations to ensure the design is economically viable in a wide variety of scenarios. The second example describes the way aircraft developments are currently assessed from an economic standpoint, points out the uncertain profitability of the *Airbus* A380 program, and highlights the need for economic evaluations that account for the flexibility offered to decision-makers. The third example describes the successive iterations of the *Airbus* A350 development and highlights the need for closer cooperation between aircraft manufacturers and suppliers to ensure a close match between technical capability and market expectations. This series of examples stems from observations of the aircraft industry during the past two decades and will contain back of the envelope calculations to highlight possible discrepancies and deficiencies in the current design approaches. These examples further emphasize the challenges and observations identified in the previous chapter in order to define where new methods might be warranted and what their purposes should be.

2.1 First motivating example

Launched in 1997, the *Airbus* A340-500 and A340-600 aircraft are derivatives of the *Airbus* A330 and A340 product lines and were aimed at replacing the aging family of classic *Boeing* 747 (747-100, 747-200, 747-300, and 747-SP). While both the *Airbus*

A340-500 and A340-600 aircraft have a lot of similar systems, the former was primarily developed as a niche aircraft for ultra-long thin routes while the latter was developed as a stretch with a substantially longer fuselage, larger wings, and up-rated engines. Back in 1997, the *Airbus* announcement was promising as *Boeing* had delivered a total of 675 classic *Boeing 747* (passenger and freight versions) and these were due for retirement in the subsequent years with no sign of a competing design being offered by *Boeing*. In addition, the A340-600 was launched with the goal of carrying a similar amount of passengers while carrying twenty-five percent more cargo at lower trip and seat costs.

In the following paragraph, a simple economic analysis of the development program is performed in order to get a net present value estimate and an internal rate of return estimate. It is not the intention of the author to perform a deeper analysis since inputs for an accurate economic analysis are proprietary. Therefore, educated assumptions are made and summarized in Table 4. The program development costs are derived from the launch aids that *Airbus* received during the development. The United Kingdom provided US\$200 million in repayable loan to *British Aerospace* [29] (which represents 20% of *Airbus*) and the launch aids were capped at 33% of the entire development cost. The weighted average cost of capital (WACC) of *Airbus* is estimated to be 12.5% as per documents from EADS [30]. Catalogue price of aircraft and engine are traditionally inflated and do not reflect what customers actually pay. Therefore, the aircraft are estimated to be sold at a 50% discount from the catalogue price [31]. A down-payment of 10% of the aircraft price is done when manufacturing starts one year prior to delivery, and the remaining 90% are due at delivery. The manufacturing costs are

estimated to be 50% of the sales price. The schedule of deliveries [32] is described in Figure 6.

| | |
|---|---------|
| A340-500/A340-600 Development Costs (Billion) | US\$3.0 |
| A340-500 List Price (Million) | US\$233 |
| A340-600 List Price (Million) | US\$240 |
| Customer Discount (Catalogue Price %) | 50% |
| Production Cost (As Price %) | 50% |
| EADS WACC (2003) | 12.5% |
| Risk Free Rate (10Y T-Bond 1996) | 6% |

Table 4: Main assumptions for economic estimations

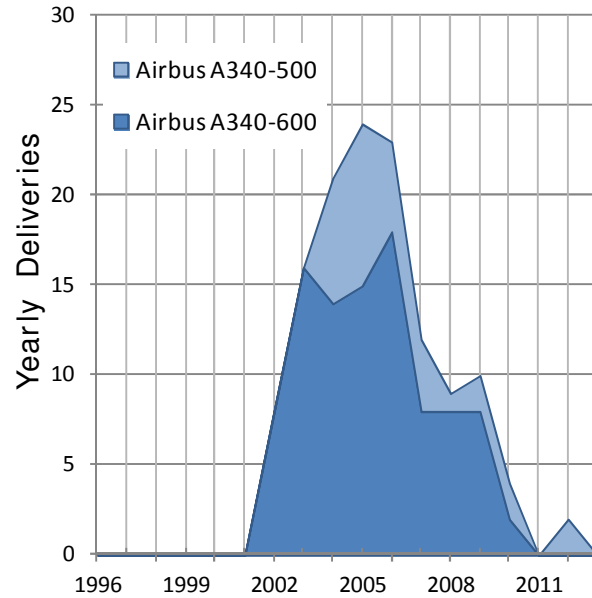


Figure 6: A340-500 and A340-600 deliveries since program launch

Using these assumptions, the value generated by the program is estimated using a discounted cash flow analysis (some more details about this type of analysis is provided in Chapter 3, section 3.1.2). For the sake of simplicity, this analysis is performed without accounting for interest or taxation effects. While most of the inputs used for this analysis can be substantiated with references, the production cost is highly proprietary and there is no reference to back this number, therefore, a very conservative estimate of 50% of the overall price is used. To reflect the uncertainty about this number, value estimates are computed with production costs ranging from 50% to 70% of the overall price and the results are provided in Table 5. First, let's preface any type of analysis by saying that these results are rough estimates following a first order back of the envelope estimation.

Yet, the results are astonishing: in most cases, the analysis indicates that the program would incur significant losses. In fact, as long as the production cost exceeds 50% of the sale price, the program is at a loss. In other words, the manufacturer needs a 52% profit margin on sales to breakeven. This is not accounting for additional discounts given to launch customers and additional expenditures to fund weight reduction programs (lighter wings) and other refinements (increased gross weight versions).

Table 5: Economic estimation for the value of the A340-500/600 programs

| | | | | |
|---|-------|--------|--------|--------|
| Production Cost (As % of Price) | 40% | 50% | 60% | 70% |
| Project Value (Estimation in Million US\$) | 572 | -36 | -644 | -1252 |
| Project Internal Rate of Return (IRR in %) | 3.32% | -0.22% | -4.05% | -8.19% |

What happened? In fact, the answer lies in the sales forecasts. Over the course of the program, *Airbus* delivered only 130 aircraft in a market which was calling for at least 675 classic *Boeing 747* replacements (assuming no organic growth). Additionally, *Boeing* launched the development of two competitor aircraft in 2000, namely the *Boeing 777-200LR* and *Boeing 777-300ER* which have reached over 655 deliveries [33] as of June 2013. Even though both aircraft types have similar capabilities (*A340-600* and *777-300ER* for the higher capacity derivatives, *A340-500* and *777-200LR* for the longer range derivatives), the four-engined *Airbus* aircraft has better field performance but is heavier and its operating costs are impacted by high energy prices. Between the commercial launch in June 1997 and the first delivery in May 2002, the price of jet-fuel had already increased by 28% from US\$0.52/gal to US\$0.67/gal according to data from the US Energy Information Administration [34]. Between the commercial launch and the last delivery in November 2012, the price of jet-fuel had increased over 450% from

US\$0.52/gal to US\$2.96/gal as shown in Figure 7. With fuel cost making up about 40% of direct operating costs for long range aircraft, the difference in fuel-burn between the two aircraft became too large [35] for the *Airbus* quad to overcome.




Figure 7: US Gulf Coast kerosene-type jet-fuel spot price

Why did it happen? Investigating the jet-fuel price time series from 1994 till 1997, the yearly volatility stands at 33% while the yearly return stands at 9.9%. In fact, with such a high volatility it is not surprising that the prices of oil could move substantially over the course of the five years of design. Let's now use these properties to get more insight about possible jet-fuel price evolution scenarios.

One popular stochastic process to model the evolution of volatile commodities, such as the price of jet-fuel, is the geometric Brownian motion [36]. In simple terms, a geometric Brownian motion is a stochastic process featuring a drift (representing the long-term trend) and some noise around this drift (modeling the volatility or variability around the long-term trend). APPENDIX A: offers a quick introduction as well as a

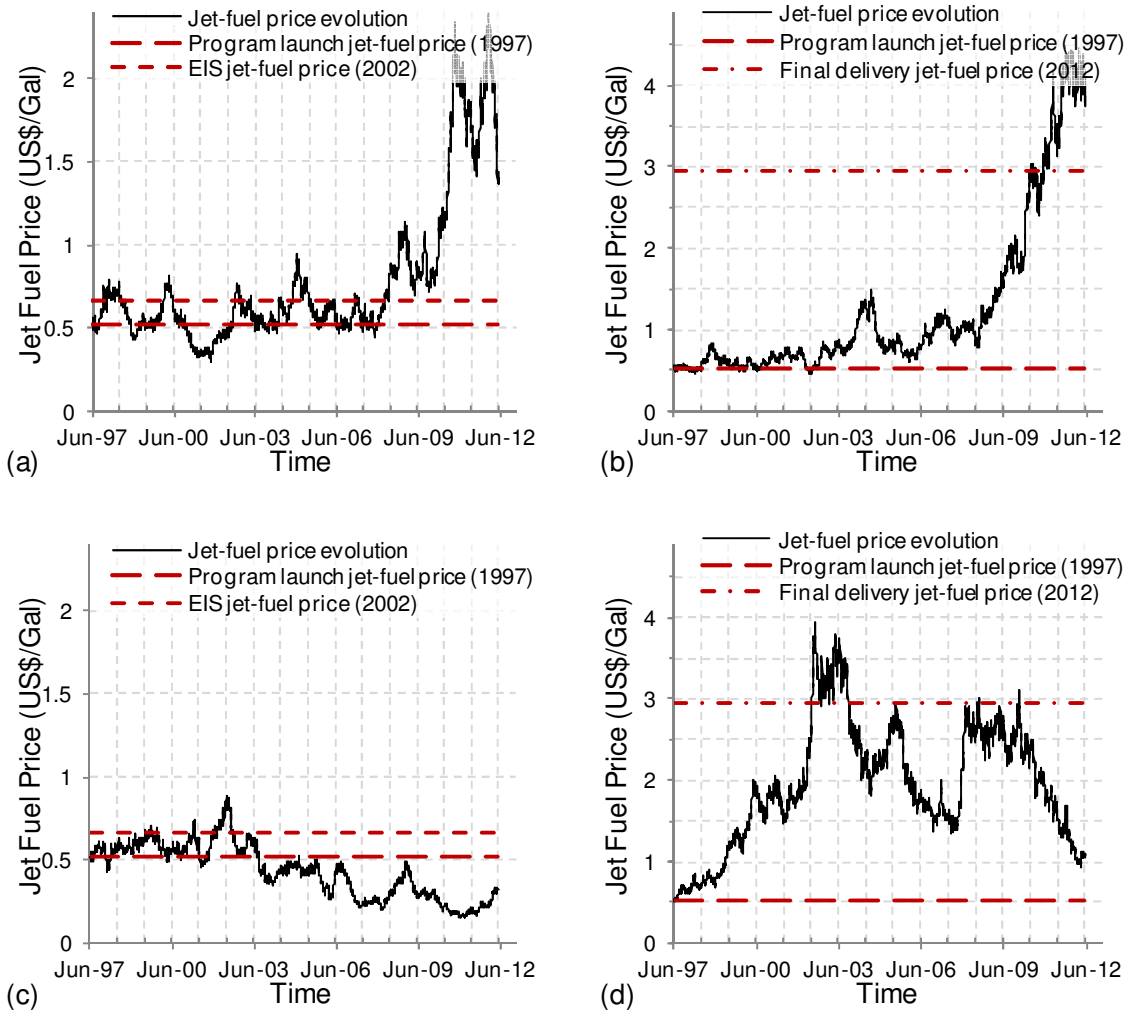
description of some of the most useful properties of the geometric Brownian motion. In particular, I derive the closed-form formula to compute the probability that a geometric Brownian motion exceeds a given threshold within an allocated timeframe (first hit time probability for a barrier). Some other formula such as the expected time for the first hit of a barrier and the probability that a geometric Brownian motion is above a threshold at a given point in time are also presented. These derivations allow revisiting some aspects of the A340-500 and A340-600 business plan using data that was available at the commercial launch of these aircraft in June 1997.

Indeed, Justin et al. [37] show that a geometric Brownian motion is a good model to represent the stochastic process followed by the price of jet-fuel between 1994 and 1997. This enables the investigation of various jet-fuel price scenarios and an estimation of their likelihood. Using the information available to decision-makers in 1997, the following estimations are performed:

- The probability that jet-fuel price ended up at least where it was at the aircraft entry into service or EIS (US\$0.67/gal in 2002). One such case is notionally represented in exhibit (a) of Table 6.
- The probability that jet-fuel price ended up at least where it was at the time of the last aircraft delivery (US\$2.96/gal in 2012). One such case is notionally represented in exhibit (b) of Table 6.
- The probability that jet-fuel price hit at one point the price it actually reached at entry into service (US\$0.67/gal) prior to entry into service (2002). One such case is notionally represented in exhibit (c) of Table 6.

- The probability that jet-fuel price hit at one point the price it actually reached at the time of last aircraft delivery (US\$2.96/gal) prior to the time of last aircraft delivery (2012). One such case is notionally represented in exhibit (d) of Table 6.
- The expected time for the jet-fuel price to first hit the price it actually reached at entry into service (US\$0.67/gal).
- The expected time for the jet-fuel price to first hit the price it actually reached at the time of last aircraft delivery (US\$2.96/gal).

Table 6: Four different types of jet-fuel price scenarios investigated



Derivations enabling the computation of these probabilities and expected time to hit are explained in detail in APPENDIX A: and the results are summarized in Table 7 for convenience. Looking at these results, it is possible to conclude that the surge in the jet-fuel price experienced during the 1997 to 2012 timeframe was not unlikely to occur within the lifetimes of the A340-500/600 programs. These results also highlight one fundamental issue that arises with the use of expected values when dealing with a leptokurtic distribution¹ with heavy asymmetric tails: even though the surge of the price of jet-fuel was not unlikely to happen and this scenario was not far-fetched by any means, the expected time for such a scenario to happen was projected to be much further into the future.

Table 7: Some jet-fuel time-series results related to the A340-500/600 programs

| Probability of event | | | | Expected time to event | |
|--|---|---|--|-------------------------------|-------------------------------|
| Jet-fuel reaching US\$0.67/gal at entry into service | Jet-fuel reaching US\$2.96/gal at last delivery | Jet-fuel hitting US\$0.67/gal before entry into service | Jet-fuel hitting US\$2.96/gal before last delivery | Jet-fuel hitting US\$0.67/gal | Jet-fuel hitting US\$2.96/gal |
| 48% | 21% | 72% | 34% | 5.8 years | 40 years |

Besides these issues, *Airbus* main competitor launched two competing designs in 2000, almost three years later, giving *Boeing* the opportunity to first gauge the market and then to benefit from the observation that oil prices were indeed surging and that significant emphasis ought to be given to the fuel-burn metrics during design. The main conclusion of this analysis is that the combination of two non-robust designs coupled

¹ A distribution is leptokurtic if its kurtosis is larger than the kurtosis of a normal distribution which means that it features positive excess kurtosis (kurtosis minus three). Leptokurtic distributions are characterized by a more acute peak around the mean as well as fatter tails than a normal distribution. Fat tails indicate that extreme observations are more likely to occur and that the risks associated with these outlier events are increased.

with the later entry into service of competing designs accelerated the demise of the A340-500 and A340-600.

First Issue:

In a competitive industry with long development cycles, there are few opportunities in the later part of the development process for manufacturers to change course as the uncertain environment unfolds. In this context, robust design simulation must be coupled with extensive competitive scenario investigations to ensure that the realization of uncertainty does not undermine a design that otherwise meets all customer requirements.

2.2 Second motivating example

Towards the very end of 2000, *Airbus* commercially launched the *Airbus A380* aiming to break the monopoly *Boeing* had for decades on the very large aircraft market segment with its *Boeing 747*. At the turn of the century, *Airbus* was embracing the hub and spoke philosophy which stipulates that airlines operate most of their flights out of a couple of mega hubs and funnel air traffic between major hubs with large-size aircraft [38]. Air traffic congestion and more generally insufficient airport infrastructure would force airlines to consolidate some of their capacity with larger aircraft and would therefore drive the demand for very large aircraft. In this context, *Airbus* was extremely bullish with the A380 development and estimated in its *2002 Global Market Forecast* [39] that over 1,100 new very large aircraft would be delivered over the 2001 to 2020 time period. It is worth mentioning that at the same time, *Boeing* made the opposite bet and projected that congestion would be alleviated by a fragmentation of the market [38]. Instead of funneling passengers through congested airports, the launch of more efficient

long range and smaller capacity aircraft such as the 787 Dreamliner would enable sustainable non-stop point to point routes bypassing traditional hubs. Few aircraft development programs have been as much scrutinized as the *Airbus* A380 program due to the size of the program and the media exposure given to this *Airbus* flagship aircraft. Data retrieved from the public domain and published around the commercial launch of the A380 program is used to perform a rough business case calculation using traditional economic valuation methodologies and the results are presented in the following paragraphs.

The assumptions and inputs for the business case analysis are summarized in Figure 8 and Table 10. The program development costs were initially projected to be US\$10.7 billion [40] in addition to the US\$700 million spent before the actual aircraft commercial launch. The aircraft catalogue price was set at US\$250 million and *Airbus* was hoping to get 50% of the overall market consisting of 1,100 new passenger airframes and 300 new cargo airframes with deliveries reaching 45 per year starting in 2009 [41]. Like in the previous example, the weighted average cost of capital (WACC) of *Airbus* is estimated to be 12.5% as per documents from EADS [30]. The aircraft are estimated to be sold at a typical 50% discount from list price with a down-payment of 10% when manufacturing starts one year prior to delivery, and with the remaining 90% due at delivery. The manufacturing costs are estimated to be close to 50% of the final price which is consistent with Noel Forgeard's initial estimate of a program break-even at 250 aircraft sold [42]. Using this set of inputs, rough estimates of the project value and project internal rate of return are provided in Table 9 with values for the unknown production cost varying from 40% to 70% of the final sale price of the aircraft.




Figure 8: Initial A380 yearly production forecasts

| | |
|--------------------------------|-------|
| R&D Costs (Billion US\$) | 10.7 |
| A380 List Price (Million US\$) | 250 |
| Cost Escalation (Yearly) | 2% |
| Cust. Discount (Price) | 50% |
| Risk Free Rate (Yearly) | 6% |
| EADS WACC (2003) | 12.5% |

Table 8: Assumptions for the A380 business case estimations

| | | | | |
|-------------------------------------|--------|--------|-----|-------|
| Gross Profit (% of price) | 30% | 40% | 50% | 60% |
| Project NPV Estimate (Million US\$) | -4,119 | -1,627 | 866 | 3,359 |
| Project IRR Estimate (%) | 9% | 12% | 15% | 17% |
| Program breakeven (Aircraft) | / | 1,255 | 580 | 399 |

Table 9: A380 project value and internal rate of return estimations for different production costs

Like in the previous example, these results are not meant to replicate a full-blown business case analysis performed by financial analysts within the company. Indeed, most of the information required to run a proper business case analysis is proprietary to the company and cannot be found in the literature. Besides, there are many aspects that are not taken into account in this simplistic analysis such as the expected revenues from selling spare parts, the future follow-up orders that may occur to replace early deliveries with improved and further optimized versions of the aircraft, and the expected costs of providing customers with performance guarantees regarding the design.

Still, this analysis should provide some indication regarding the profitability likelihood of the program. Table 9 reveals that under the assumptions made, the program does not look very profitable. In fact, the results show that the program is profitable only if the gross profit margin reaches 47%. In other terms, the program remains profitable only if the production cost of a new A380 remains below 53% of its sale price. This looks like a

risky proposition. Besides, even with a gross profit margin of 60%, the internal rate of return does not exceed 17%. With a weighted average cost of capital at 12.5%, it seems unlikely that the management would accept a hurdle rate of only 17% for such a risky flagship program.

These results are flabbergasting. The A380 received so much media exposure at launch-time that it seems almost paradoxical that the business case could be so weak. In fact, the author posits that most of the large-scale aircraft and engine development programs valued using traditional neoclassical valuation methodologies would yield similar shaky business cases. Yet, new aircraft and engine developments are announced every year and both *Airbus* and *Boeing* are profitable [43].

What is happening? Part of the issue may be related to the use of discounted cash flow analysis to perform economic valuations. As will be discussed in Chapter 3, section 3.1.2, the discounted cash flow analysis is a sound method to value cash flows, but may not be the most adapted valuation method to perform economic valuation when uncertainty abounds. Indeed, this type of analysis is perfect for valuation with little or no uncertainty as most investments are treated as a “now or never” decision, and once committed, the investments cannot be revised or rescinded. This static approach to capital budgeting is to put in contrast with the very dynamic nature of capital budgeting required in a highly volatile and uncertain environment. As Justin et al. [44] mention, the static approach fails to recognize the flexibility offered to management when uncertainty unfolds and the value added by management when steering development programs into profitable directions. This results in a systematic undervaluation of long-term

development programs as the discount factor used to perform the valuation may not reflect the risk-mitigating impact that (successful) managers have.

During the development of the *Airbus A380*, managerial flexibility has been used at least once to significantly alter the development program and improve its profitability. It happened in early 2007 when *Airbus* decided to stop and delay the development of the A380F freighter version [45] following limited sales and production issues that required diverting engineering capabilities to work on the passenger version of the aircraft. In this case, the research and development effort had been committed at the program commercial launch, yet only a fraction had been spent to prioritize the more promising passenger version. As a sideline benefit, the sunk costs of the A380F were not completely lost as the development work for the freighter version had been reused to develop increased gross weight of the passenger version.

Second Issue:

In a volatile industry sensitive to business cycles, uncertain energy prices and evolving customer requirements, managerial flexibility defined as the ability of management to actively steer research and development programs into profitable direction is valuable and must be accounted for when business plans are laid-out. Traditional capital budgeting methods do not usually account for this flexibility and consequently undervalue significantly long-term aircraft and engine developments.

2.3 Third motivating example

Airbus launched in late 2004 the A350 as a new design based on the fuselage of the highly successful *Airbus A330*. The new *Airbus* would be a new long-range wide-body twin-engine aircraft featuring a new wing with new materials as well as new

engines and updated systems. The main objective for *Airbus* was to rejuvenate the product line to better compete with the *Boeing 787 Dreamliner* launched almost two years before and which was getting significant traction on the market. Marketed as a highly efficient and environmentally-friendly aircraft, the competitor from *Boeing* was indeed threatening the lucrative position of the A330 and risked making the *Airbus* aircraft obsolete.

However, giving birth to its new aircraft was not an easy exercise for *Airbus*. In fact, the history of the A350 starts in early 2004 when *Airbus* first presented an updated A330 with better aerodynamics and engines dubbed A330 Lite as a response to the Dreamliner [46]. Later that year, *Airbus* decided to change strategy, significantly revamp the program by updating the wing and empennage, and call it A350 [47]. Following lukewarm reception by the airlines who wanted a clean-sheet design, *Airbus* announced a year later at the 2006 Farnborough air show that it was redesigning again the aircraft with a wider fuselage and that it was now calling it the A350XWB [48].

What happened? It is well accepted that *Airbus* was taken off-guard with this latest offering from *Boeing* and initially scrambled to provide a definite strategic answer to the new threat [49]. After all, even John Leahy of *Airbus* [49] jokingly mentioned these successive redesign attempts “*Everyone was writing that we redesigned the aircraft six or seven times. We didn't. We redesigned it three times, and that was enough.*” However, besides the initial surprise and the fact that it took *Airbus* several design iterations to fully address the needs of airlines and settle with the A350XWB, significant design changes were far from over: the A350XWB design from July 2006 is indeed quite different from the one that first flew in June 2013. Following criticism from potential

customers, the design of the A350XWB fuselage was updated in 2007 to reflect the latest advancements in carbon fuselage design and switch from a hybrid metallic fuselage frame and carbon panels to an all composite fuselage frames and carbon panels [50]. Some analysts, such as Gary Chapman [51] from Emirates Group Services and Data, believe that *Airbus* was probably not technologically ready to build an all-carbon fuselage in 2006 and that “*Airbus has probably [caught-up and] learned a lot from what Boeing has done with the 787.*” As *Airbus* was fine-tuning its manufacturing capabilities at the same time the aircraft design was progressing, the design space opened up incrementally leading to many subsequent revisions to yield a more competitive aircraft to customers.

This example illustrates a fundamental problem facing many companies at the forefront of technological innovation: the timing adequacy between when technologies become available and mature enough to be used for a commercial application and when the company has the ability to develop and add a new element in its product line. In this last sentence, the term “ability” is to remain generic and may have different meaning depending on the situation: for the aerospace industry, it could be engineering manpower (constrained by limited skilled workforce), market acceptance (constrained by market demand), or financial resources (constrained by limited capital expenditures).

Third Issue:

In an industry where manufacturers can neither afford to have a gap in their development pipeline (to retain skilled workforce) nor develop two clean-sheet designs concurrently (due to limited engineering capabilities), the time at which technologies become mature for commercial application becomes crucial. These timings issues need to be anticipated with both the company and the competition product development pipelines in mind.

2.4 Research motivation and thesis organization

The first chapter presented a list of challenges affecting the aircraft and engine manufacturing business. These challenges are very diverse and affect many different aspects of the business. Some of these are directly related to the design activity itself and include the management of technical and technological risks as well as the management of long-term development programs with evolving customer and regulatory requirements. Some are related more closely to the production and manufacturing, and include the supply chain management and the ability to meet production ramp-up targets. Finally, some challenges are related more closely to the economics of these developments and the ability to establish a proper business case in an uncertain and competitive environment.

The second chapter presented three real-life aircraft development programs in order to illustrate how some of the aforementioned challenges may have impacted or hindered decision-makers during some phases of the design. In particular, three major problems were identified: the first is the drastic impact uncertainty can have on non-robust aircraft designs in competitive environments; the second is the establishment of business cases with methods and techniques that may not adequately capture uncertainty and its impact on development program management; the third is the need for techniques and methodologies that enable detection of precursors to ensure optimal synchronization between (enabling) technology portfolio maturation and product development schedules. These observations lead to the overarching motivation for this research which is stated as follows:

How can current state-of-the-art design methodologies be updated and improved so that they help:

- *Identify precursors of technical, technological, and market opportunities leading to successful aircraft and engine developments*
- *Account for the widespread uncertainty surrounding aircraft and engine developments in competitive environments*
- *Establish a business case which reflects the entire spectrum of means available to management to steer development programs into profitable directions*

The main motivation for this research is formulated in a more compact setting below.

With a more thorough and profound understanding of the implications of this problem statement, this motivation will be declined into several subsequent research questions.

Overarching Research Question – Improvement of value-based design methods

Within the context of aerospace research and development optimization, how can value-based design methods be improved to identify precursors of technological and market opportunities while reflecting the specific challenges associated with long-term and uncertainty-plagued developments, and while accounting for the competitive nature of the business?

In the third chapter, a literature review is performed to review and identify the most appropriate methods for the present research. These methods are sorted into four different categories: methods for the construction of a business case; methods for the valuation of development programs with uncertainties and managerial flexibility; methods for the identification of precursors and trigger events; and finally methods to perform competitive assessments.

In the fourth chapter, an in-depth review of the real options literature is performed to analyze why current real options methodologies fail to be widely used despite their many theoretical advantages. The weaknesses and limitations of current methods are highlighted while investigations take place to relax some of the most constrictive hypotheses. This leads to a series of hypotheses regarding potential improvements to current real options methods.

In the fifth chapter, the original problem statement is revisited and issues identified in previous chapters are synthesized. The various research questions and hypotheses formulated in the previous chapters are revisited and a mapping between research questions and hypotheses is performed to facilitate the layout of verification and validation processes. Three types of hypotheses are formulated: *Method hypotheses* which concern a set of ordered procedures to investigate and resolve real-life problems faced by practitioners in the industry, *Modeling hypotheses* which address generic mathematical representations of some aspects of real-life, and *Technical hypotheses* which deal with specific mathematical techniques to solve specific mathematical problems

In the sixth chapter, a novel methodology is constructed via cross-fertilization of techniques and methods used in the actuarial sciences, statistics and quantitative finance industry to improve the current state-of-the-art in evaluation methodologies. It builds upon traditional methods but makes use of advanced evaluation techniques presented in the previous chapters which aims at assessing staggered investments featuring flexibility.

In the seventh chapter, an experimental plan is proposed to determine a set of experiments necessary to prove or disprove each of the technical hypotheses. In this case,

a pure mathematical verification is usually sufficient to ensure they properly address and solve the identified mathematical problems. The experimental plan is carried out using a series of canonical examples starting with a traditional stochastic process and then moving to a complex process featuring jumps.

In the eighth chapter, a proof-of-concept is introduced to demonstrate the applicability of the proposed methodology in a typical aerospace industry setting. This proof-of-concept concerns a Performance Improvement Package (PIP) which is being offered to aircraft operators as a retrofitting option to improve the operating economics of a currently out of production aircraft. The uncertainties affecting the value of the PIP development program are identified and calibrated using market data. The section concludes with the development of a market model to estimate the adoption of the PIP by engine operators worldwide.

In the ninth chapter, the aforementioned case study is used to validate the method and modeling hypotheses formulated in previous chapters. The method and modeling hypotheses propose a mathematical abstraction to represent some tangible aspects of the real-life and the validation ensures that the mathematical abstraction is adequate, is suitable for the envisioned application, and finally, represents all pertinent aspects of the problem.

Finally, in the tenth chapter, general conclusions are drawn and the main contributions of this research are summarized. Several improvements to the proposed methodology are also suggested for future research.

CHAPTER 3: PROBLEM DEFINITION

3.1 Valuation methods

Valuation methods are sets of procedures and techniques used in order to assess the economic value of a business or a prospect. The purpose of a valuation technique can be summarized by the deceptively simple question: “how much is the business or prospect worth?” It is an essential step in any capital budgeting decision as it provides the rationale for selecting business ventures that add value to a firm. Over the years, many techniques have been proposed with higher and higher levels of sophistication. This increase in complexity results from the observation that more and more *hidden value* needs to be accounted for in the evaluation of prospects (value of flexibility) and the realization that a venture being merely profitable is not sufficient to warrant a significant investment (investment efficiency and capital constraints). This section reviews several ubiquitous as well as cutting edge valuation methods.

3.1.1 Payback period

A time-honored capital valuation method that was once prevalent, the payback period is the selection criteria that most business firms used in order to select capital investments [52]. Its main idea is that the sooner the break-even point of a venture, the better. For a capital investment project, the payback period is the time, usually expressed in years, needed to payback the initial investment from the future expected cash flows. It

is therefore based on assumptions regarding the future cash outflows and cash inflows. Mathematically, this may be computed with the formula shown in Eq. 1.

$$T = \frac{\text{Initial Investment}}{\text{Expected Annual Net Cash Flows}} \quad \text{Eq. 1}$$

All else being equal, an investor comparing different capital projects would prefer the one with the shortest payback period. Byrne et al. [53] argue that investments with short payback periods are relatively liquid investments that minimize the lost opportunity risks. Payback period as a valuation method is still used, especially in smaller firms, due to its simplicity. It has nonetheless serious limitations in that it does not account for the time-value of money, the risks associated with the capital investment, and the opportunity costs. Last but not least, the main shortcoming of the payback period is its failure to provide any information about the expected profitability of the investment.

3.1.2 Discounted cash flow analysis and the net present value

Recognizing the limitation of the payback period method, Irving Fisher [54] and John Bur Williams [55] formalized the discounted cash flow analysis using the concept of time-value of money. The adoption of the discounted cash flow analysis has been flabbergasting and it is believed to be the most widespread method to assess the economic performance of large investments made by corporations. In 1972, Klamer [56] reported that 19% of firms in a survey were using discounted cash flow techniques. This number increased to 75% in 2001, as reported by Graham and Harvey [57].

A discounted cash flow analysis starts with a forecast regarding both the investment outlays (the cash outflows) and the revenue prospects (the cash inflows) stemming from the investment. These forecasts may be assessed using historical analysis

(the past is a good starting point to predict the future), statistical analysis (comparable investments provide information regarding the outlays and revenue characteristics), educated guesses from subject matter experts, or simply best-guess estimates from managers. To account for the time-value of money, the cash inflows and cash outflows are each discounted according to their respective risks. Finally, the sum of all these discounted cash flows is called the net present value. An investment featuring a negative net present value is not economically viable whereas an investment with a positive net present value is viable. When funding constrains the number of investments to be undertaken, an investor comparing different capital projects would, all else being equal, select the investment having the highest net present value. As a side note, the discounting may be done according to two different conventions [58]: if the model is set up in continuous-time, the continuous compounding discount rate is used, while if the model is discrete, then a periodic discount rate is used. These two ways of discounting are shown in Eq. 2, to the left for continuous compounding, and to the right for discrete compounding.

$$NPV = \sum_{i=0}^n CF_i \cdot e^{-rt_i} \qquad NPV = \sum_{i=0}^n \frac{CF_i}{(1+r)^i} \qquad \text{Eq. 2}$$

Where CF_i is the cash flow at time t_i and r is the discount rate.

Theoretical underpinning of the discounted cash flow analysis

The theoretical underpinning of the discounted cash flow analysis may be explained using consumption choice arguments. To simplify the argumentation, let's follow Ross et al. [59] and assume a two-period economy where investors may borrow or lend at an equilibrium interest rate r . The investor has an initial wealth today of X and the

consumption choices offered to him may be represented by the graph shown in the exhibit (a) of Figure 9.

In this graph, the investor is currently in position *A* which represents his ability to consume in Period 1 all of his income in Period 1 and his ability to consume in Period 2 all of his income in Period 2. This is however not the only option available to an investor having access to the capital market. One extreme option for the investor is to consume all of the income of the first period and borrow money against the income from the second period which is represented by point *B* with a first period consumption of $X+Y/(1+r)$. Another extreme option is to not consume anything during the first period and instead lend at the interest rate r in which case the wealth available for consumption during the second period is represented by point *C* and is equal to $Y+(1+r)X$. Finally, any point *in-between* situations *B* and *C* is possible: moving from *A* to *B* indicates that the investor consumes more today and borrows, while moving from *A* to *C* indicates that the investor defers consumption and therefore lends. The line from *B* to *C* is straight because no individual has any effect on the market-driven interest rate. Each point along this line represents a specific consumption choice given the wealth for the two periods. It is called the intertemporal budget line.




Figure 9: Consumption choice in a two-period model

Let's now introduce another point, A' , as shown in exhibit (b) of Figure 9. This new point represents a new investment opportunity available to the investor. What is interesting about this new point is that it is placed above the original intertemporal budget line. Does undertaking this investment increase the wealth of the investor? To answer, let's look at how the investor could replicate this new opportunity. Replicating the investment means hopping from the original intertemporal budget line to the new one where A' lies. Moving from A to D' is one way of achieving this goal and since the investor is moving towards the right of the graph, it means that the consumption in the first period is increased without any change to the consumption in the second period. Similarly, moving upwards from A to E' means that the consumption in the second period is increased without changing the consumption in the first period. In both cases, higher overall consumption is achieved, which indicates that replicating this investment opportunity increases the wealth of the investor. The distance between A and D' or between A and E' is what characterizes the wealth created by the investment and is called the net present value (NPV) of the investment.

One take-out from this simple two-period example is that the net present value is a simple criterion to decide whether to undertake an investment: as long as the investment

exhibits a positive net present value, then it can be replicated by giving cash to the investor which is equivalent to creating wealth and therefore should be undertaken. On the contrary, an investment exhibiting a negative net present value can be replicated by giving up cash which destroys wealth and should therefore not be undertaken.

Another take-out from this example is that no investor preference has been used to assess whether the investment should be undertaken. This means that the investment decision is made only by estimating the net present value regardless of the individual preference for consumption now or in the future. This is the basis of Fisher's separation theorem which states that the objective of a corporation is the maximization of its present value and all of the owners (shareholders and debt holders) will agree on which investment to undertake regardless of their individual tastes for consumption and savings.

Theorem: Fisher Separation Theorem

Each person, after or while first choosing the option of greatest present worth, will then modify it by exchange so as to convert it into that particular form most wanted by him. This implies, as we have seen, that each person's degree of impatience, or rate of time preference, will at the margin, be brought to equality with the market rate of interest and, therefore, with the marginal preference rates of all the other persons.

Selecting a discount rate

In the previous section, the analysis is performed using a *market interest rate* for lending and borrowing purposes. The nature of the interest rate is, however, not explained and the term remains pretty vague; it could indeed be a risk-free rate or a risk-adjusted

discount rate. In fact, the discount rate needs to reflect the cost of borrowing money or better said the ease of access to capital.

For private companies, raising capital is usually done in two ways: either by issuing bonds or by issuing stock. Issuing bonds is equivalent to raising money by getting debt whereas issuing stock is equivalent to raising money by diluting ownership of the company. Both bondholders that have a creditor stake in the company and stockholders that have an equity stake of the company needs to be compensated through either coupon and principal payments or dividend payments which are funded using the cash flows generated by the company. Consequently, investments made by the company must have a return that is sufficient to pay for both the cost of debt and the cost of equity.

The Weighted Average Cost of Capital (WACC) is the rate that a company is expected to pay on average to all its bondholders and stockholders to finance its assets. It is therefore the minimum return that a company must seek on its asset to satisfy its creditors and owners. In the model with two sources of capital, the WACC may be estimated as shown in Eq. 3. In this formulation, the cost of debt and the cost of equity are weighted¹ according to the total debt and total equity of the company. Total debt and total equity can be estimated quite simply by looking at the outstanding debt and the number of outstanding shares and their market price. The cost of debt and the cost of equity are, however, more complex to estimate.

¹ The cost of debt is adjusted using the corporate tax rate since debt interest is tax deductible.

$$WACC = \frac{D}{D + E} K_d(1 - T_c) + \frac{E}{D + E} K_e$$

Eq. 3

D: total debt, E: total equity, K_d: cost of debt,
K_e: cost of equity, T_c: corporate tax rate

For the cost of debt, a historical approach or a market approach may be used. In the historical approach, financial statements of the company are used to estimate an average interest rate by weighting the interest of each individual loan according to its principal. A more accurate approach reflects current information from the marketplace (as opposed to past information) and uses the yield-to-maturity¹ of the outstanding debt. This requires that up-to-date market prices of the company's outstanding bonds exist which is not often the case. Finally, a practical approach that still incorporates market information consists in first looking up the company's debt rating (*Moody's, S&P, Fitch*) and then looking up the market yields of bonds with similar maturity and rating.

The cost of equity reflects the rate of return that a well-diversified shareholder would require given the exposure of the company to non-diversifiable risks. One way to estimate the cost of equity is to use the Capital Asset Pricing Model (CAPM) formalized by Treynor [60], Sharpe [61], Lintner [62] and Mossin [63] following the earlier work of Markowitz [64] on diversification and the modern portfolio theory. There are several assumptions used in the capital asset pricing model and the intent here is just to give the reader the essence of the model. For a more complete derivation, the user is referred to APPENDIX E: as well as [65] for some information about modern portfolio theory and

¹ Yield to Maturity (YTM) of a bond is the internal rate of return of a bond held until maturity and for which all coupon and principal payment are paid on schedule

the efficient portfolio frontier, and to APPENDIX F: as well as [65] for some more information about the derivation of the CAPM and about why the market portfolio is an optimal portfolio.

The capital asset pricing model uses the concept of market portfolio which is a basket of investments containing every asset available in the financial market, with each asset weighted according to its total presence in the market. Since the market portfolio is completely diversified, it is subject only to systematic risk. In this setting, the risk premium of any asset over the risk-free rate of return is directly related to both its contribution to the market portfolio rate of return (expressed in $E(r_e) - r_f$) and its contribution to the market portfolio risk (expressed as a covariance $cov(r_e, r_M)$). The capital asset pricing model yields the mathematical expression in Eq. 4 for the theoretical rate of return of an asset.

$$E(r_e) = r_f + \frac{cov(r_e, r_M)}{\sigma_M^2} [E(r_e) - r_f] \quad \text{Eq. 4}$$

With r_e the return of an asset, r_f the risk-free rate of return, r_M the return of the market portfolio, and σ_M the volatility of the market portfolio.

Using the expressions for the cost of debt and the cost of equity, it is now possible to compute the weighted average cost of capital to be used for discounting purposes as shown in Eq. 5.

$$WACC = \frac{D}{D + E} K_d (1 - T_c) + \frac{E}{D + E} \left[r_f + \frac{cov(r_e, r_M)}{\sigma_M^2} [E(r_e) - r_f] \right] \quad \text{Eq. 5}$$

Shortcomings of the discounted cash flow approach

The discounted cash flow approach is the current “standard” for evaluating business prospects and assessing their profitability. Its widespread acceptance may be traced to several reasons including a rigorous foundation and an early and widespread exposure to future practitioners in colleges and business schools. Besides, the method is both transparent and straightforward requiring only a single “black box” parameter, the discount factor, which may be hard to come-by on an ad-hoc basis but which is often supplied company-wide by the upper management.

However, the method is not devoid of shortcomings. One of the recurrent pitfalls is related to the use of the same discount factor to discount both the expenditures and the revenues. In many business ventures, the costs may have much less uncertainty (risk) than the revenues. Consequently, some expert argues that the cash outflows and the cash inflows should not be discounted using the same discount rate.

Another issue pertaining to the discounted cash flow analysis is related to the evaluation of long-term business ventures. In these cases, one (a most likely) or several (bad, good, and most likely) scenarios are identified and investments as well as revenue streams are generated for each case. However, this type of approach does not recognize the fact that long-term investments often present many opportunities for the management to react to the realization of uncertainty and to alter the course of the project. In turn, this managerial flexibility brings additional value to the firm. The value of *flexibility* is not accounted for in the rather “static” discounted cash flow analysis which assumes a passive management once capital investments are committed.

3.1.3 Internal rate of return

Whereas the discounted cash flow analysis attempts to estimate the value of an investment, the internal rate of return attempts to quantify the yield or quality of an investment. Both approaches are in fact very similar and mostly differ in what is considered to be an input and what is considered an output of the analysis. While the discounted cash flow analysis starts with estimates of cash flows that are then discounted to find the investment value, the internal rate of return analysis starts with estimates of the cash flows and then estimates the discount rate that yields a net present value of zero. In other words, the internal rate of return is the rate of return that makes the net present value of the investment equal to zero. This discount rate is compared next to the cost of capital to investigate whether the investment is worth pursuing.

In mathematical terms, if the cash flows resulting from the investments are estimated on an annual basis, the internal rate of return r_{irr} may be computed as the solution of Eq. 6. Most of the time, there are no analytical solutions to this equation. Yet, if there are alternating positive and negative cash flows, then there may be several solutions to this equation. Different root-finding algorithms may be used to estimate the solution to this equation: if a single solution is suspected, a bisection approach could be appropriate, if faster convergence is sought, a secant method may be more appropriate.

$$0 = \sum_{i=0}^n \frac{CF_i}{(1 + r_{irr})^i} \quad \text{Eq. 6}$$

The secant method is a time-honored root-finding algorithm that performs successive linear approximations to the function to be solved. Depending on the function and on the initial guess for the internal rate of return, convergence is not guaranteed.

However, an initial guess sufficiently close to the solution is usually sufficient. The successive iterations for the estimate of the internal rate of return are given in equation Eq. 7.

$$r_{irr}^{n+1} = r_{irr}^n - NPV^n \left(\frac{r_{irr}^n - r_{irr}^{n-1}}{NPV^n - NPV^{n-1}} \right) \quad \text{Eq. 7}$$

Despite the apparent hurdle of solving for the internal rate of return iteratively, the method has been embraced by practitioners and is one of the most popular methods to assess the economic viability of investments. Practitioners find it easier to use than the discounted cash flow analysis when comparing investments of different sizes because it yields a single figure of merit that is non-dimensional. In addition, it measures investment efficiency and may therefore give better insights in capital constrained situations.

Shortcomings of the internal rate of return

To illustrate some of the shortcomings of the internal rate of return methodology, let's introduce a notional example. Let's assume that a major aerospace company may elect to pursue one of two different investment opportunities: one is to invest in a low cost derivative aircraft, the other is to invest in a higher cost brand new aircraft as shown in Table 10.

Table 10: Aircraft development notional example

| | Low Cost Derivative Widebody Aircraft | High Cost New Widebody Aircraft |
|--------------------------------|--|--|
| Initial investment | 4.0 Billion | 8.0 Billion |
| Yearly net cash flow | 1 Billion | 1.9 Billion |
| Production run | 10 Years | 16 Years |
| Development time | 3 Years | 5 Years |
| Risk-free discount rate | 5% | 5% |

Using this simple example, let's plot the graph showing the net present value of these two aircraft development projects for different discount rates. First, what is striking in Figure 10 is that the two curves representing the two projects cross each other. This means that depending on the discount factor chosen, the supposedly better project changes from the derivative to the new aircraft. Looking at which project yields the highest internal rate of return, it appears that the derivative is the most capital efficient one as highlighted in Table 11. However, looking at the net present value using a discount factor equal to the weighted average cost of capital, it seems that the new aircraft is creating more value to the firm. Given the contradictory nature of these two results, how to substantiate decision making?




Figure 10: Development program value for different discount rates

| | Low Cost Derivative Widebody Aircraft | High Cost New Widebody Aircraft |
|--|--|--|
| NPV (Billion US\$) (WACC=12.5%) | 0.25 | 0.37 |
| Project IRR (%) | 16.7 | 15.8 |

Table 11: Conflicting NPV and IRR values for the new and derivative projects

Let's first investigate why the two results differ. The first reason is that the internal rate of return approach estimates the yield or capital efficiency of an investment and not the value created and added to the company by the investment. Therefore, investments with large initial expenditures may be erroneously turned down for shorter-term investments with smaller initial expenditure. Another reason is that the internal rate of return approach assumes that the interim cash flows are reinvested in assets having the

same internal rate of return as the investment under review. This poses a problem particularly relevant for investments with large internal rate of returns as there may not be any alternative investment available offering a similar rate of return. In this case, the computation may overestimate the real internal rate of return of the investment.

There are several reasons to this conundrum but all of the explanations point to one direction: the internal rate of return is a good metric to assess the efficiency of an investment but should not be used to compare mutually exclusive projects. This precludes the use of the internal rate of return for the research problem investigated.

3.1.4 Modified internal rate of return

As its name implies, the modified internal rate of return (r_{mirr}) is a modification of the internal rate of return in order to avoid the over-estimation induced by the assumption that intermediate cash flows are reinvested at the same rate of return as the investment itself.

In order to achieve this goal, the modified rate of return uses both the present value of negative cash flows and the future value of positive cash flows as seen in Eq. 8. The negative cash flows are discounted to the present time at the cost of capital (r_{wacc}). The positive cash flows are compounded to the final time (last period of interest) at a proper re-investment rate (r_{inv}) typical of the business. Another improvement over the classic internal rate of return is that the modified internal rate of return formula yields a single solution. There is therefore no guessing required to establish which solution is appropriate when several are plausible.

$$r_{mirr} = \sqrt[n]{\frac{\sum_{i=0}^n [\max(0, CF_i) \cdot (1 + r_{inv})^{n-i}]}{\sum_{i=0}^n \left[\frac{\min(0, CF_i)}{(1 + r_{wacc})^i} \right]}} \quad \text{Eq. 8}$$

Despite these improvements, the modified internal rate of return suffers from the same issue as the internal rate of return for the selection of mutually exclusive projects. Indeed, it estimates a metric of investment efficiency and not the aggregate added value of an investment.

3.1.5 Real options analysis

Another evolution of the discounted cash flow analysis that uses both cash flows over the entire life of the investment and a market-derived opportunity cost of capital is the real options analysis. Real options analysis is an emerging field in corporate finance where it is used to substantiate capital budgeting decisions when uncertainty abounds. Its emergence at the turn of the 21st century stems mainly from two facts: (1) the realization that a pure discounted cash flow approach does not reflect the flexibility offered to decision-makers and (2) the recent adaptation of option valuation techniques originally developed for financial trading to capital budgeting analysis.

Real options analysis goes beyond discounted cash flow analysis because it recognizes that managers do not stand still while uncertainty is unfolding, but rather actively steer projects into profitable directions. Decision-makers react to changes in the business environment, abandon projects that are not economically viable, and add resources to those that are promising given the latest realization of uncertainty.

Since the analysis accounts for the abandonment of unprofitable ventures, their values may be understood to be similar to the values of financial call options that are

exercised only if the values of the underlying assets are larger than the exercise prices. Like a financial option, a real option is the right but not the obligation to undertake a business decision. As such, the value of a research and development project may be viewed as the value of the option to fund research and the value of the option to fund the development program. In this sense, real options analysis is an extension of the seminal work pioneered by Black, Scholes, and Merton [66] [67] regarding financial options. In the case of real options, however, the underlying assets are not stocks, futures or forward contracts but usually real assets such as research and development programs. Typical examples of real options used in early works were the options to expand, shrink, or abandon investments in the mining industry.

What is an option?

Before getting into the technicalities of the options analysis, it is necessary to introduce the jargon and establish the parallelism between financial options analysis and real options analysis. Option contracts belong to a larger family of financial instruments called derivative securities or *derivatives*. As the name suggests, these are financial instruments whose prices are determined by the prices of other securities and are also called *contingent claims* because their payoffs are contingent on the prices of other securities. Because their values depend on the values of other securities, these instruments can be used for both hedging and speculation purposes. Options may be traded on exchange as well as in the over-the-counter market which means that they may come in many different flavors and forms. The simplest options, also known as vanilla options, are the *European call* options and the *European put* options. More complex

options are often referred to as exotic options. In its simplest form, a European call option is a contract with the following properties:

- It has a prescribed time in the future known as the *expiration date* or *maturity*
- It specifies a prescribed asset known as the *underlying asset* or *underlying*
- It gives the holder of the contract the possibility to buy at maturity a prescribed amount of the underlying at a prescribed price known as the *exercise price* or *strike price*

The fact that this contract gives the holder the *possibility* to purchase the underlying is important because this means that this contract provides a *right* and not an *obligation*. This provides flexibility to the holder to wait and see the evolution of the price of the underlying before committing to the purchase. Because this contract confers to its holder a right with no obligation, it has a value which must be paid at the time of opening the contract: this is the *price* of the option contract. When analysts mention option valuation techniques, they refer to those techniques that permit the pricing of these options. Similarly, a European Put Option is a contract with the following properties:

- It has a prescribed time in the future known as the *expiration date* or *maturity*
- It specifies a prescribed asset known as the *underlying asset* or *underlying*
- It gives the holder of the contract the possibility to sell at maturity a prescribed amount of the underlying at a prescribed price known as the *exercise price* or *strike price*

There are many other types of options and a review of these is not necessary at this point of the dissertation. An interested reader might find more information in Espen Gaarder Haug [68].

Drawing a parallelism with the theory of financial options, a real option is the right but not the obligation to undertake a business decision sometime in the future. These business decisions may have many different shapes and goals. The ones of interest in this research are development programs and more generally investments in the aerospace industry. These investments usually present many opportunities over the course of their existence, such as the possibility to abandon a non-profitable investment (abandonment of the *Boeing* Sonic Cruiser), the opportunity to expand a profitable investment (increase the rate of production of the *Boeing* 787), the opportunity to develop derivatives (develop the *Airbus* A321 if the *Airbus* A320 is successful), or the possibility to defer a risky investment (delay the development of the *Boeing* 777-300ER). In these cases, the relationship between financial options and real options may be established by mapping financial parameters used in the financial option literature to parameters used for investment valuation and capital budgeting. A framework to establish this duality is presented in Table 12.

Table 12: Mapping financial options jargon and real options jargon

| Parameter | Financial options | Real options |
|----------------------------|---|--|
| Option | Legal Contract | Business Decision Regarding an Economic Endeavor |
| S | Stock | Present Value of Program Cumulative Cash Flows |
| K | Strike Price | Present Value of Delayed Capital Investment |
| T | Maturity | Option Life |
| r | Risk-Free Rate | Risk-Free Rate |
| σ | Standard Deviation of Return (Volatility) | Standard Deviation of Returns of Project Value |

Theoretical underpinning of the real options analysis

The theoretical underpinning of real options analysis relies on option valuation techniques developed in the finance and trading industry and later applied to capital budgeting decisions. Therefore, a review of some of the work developed within the context of the finance industry is warranted to fully appreciate both the similarities and differences between financial options and real options.

Using the previously defined nomenclature, the payoff at expiration T of a financial option on an underlying stock S with a strike price K and with a purchase premium of V may be represented by Figure 11 with exhibit (a) for a call option (option to buy stock) and exhibit (b) for a put option (option to sell stock). Exhibit (c) and (d) represent the final payoff without accounting for the option purchase premium which is more customary in the option literature. As shown in exhibit (c), when the price of the underlying stock exceeds the strike price at expiration, the payoff of a European call option is positive because the contract holder will exercise the option and sell the stock immediately after, making a profit of $S-K$. Similarly, exhibit (d) shows that when the price of the underlying is below the strike price of the European put option, the contract holder will immediately exercise the option (sell the stock) and buy it back at the prevailing market price, therefore making a profit of $K-S$.

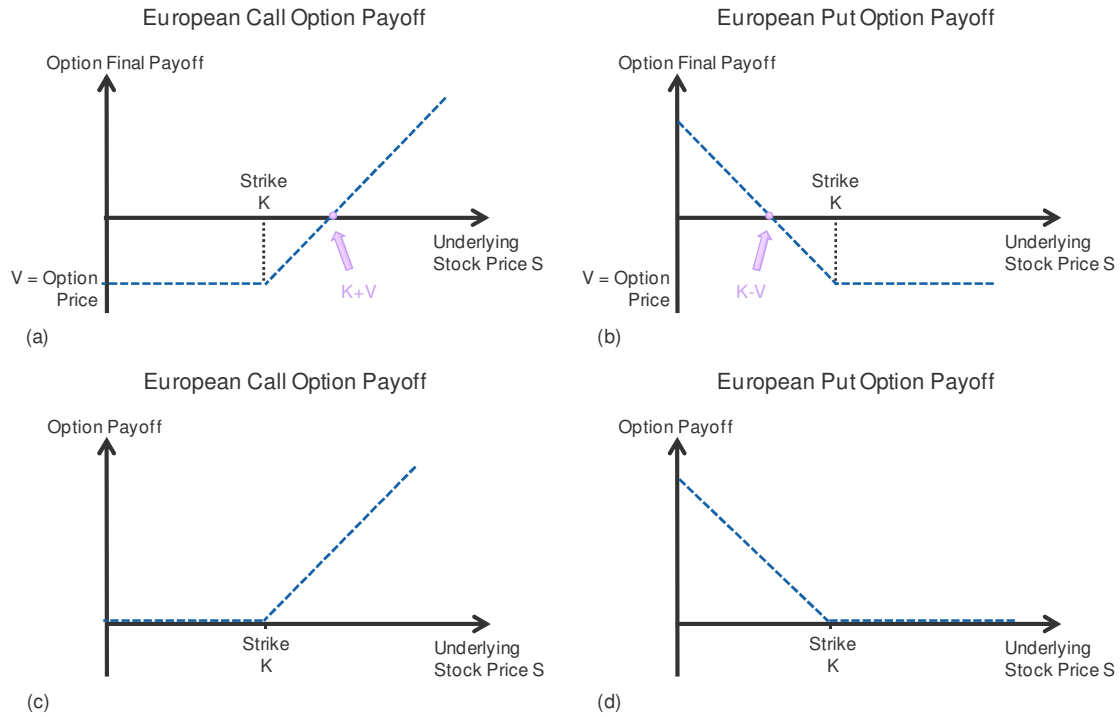


Figure 11: Vanilla options payoffs at expiration (accounting for the option premium)

Having defined the payoff of the options at expiration is the first step for options valuation purposes. How can the analyst use this information to estimate the price of an option? This is in fact much more complicated than it seems. A naïve answer suggests using the expected distribution of the asset price at expiration to estimate an expected payoff and then discount this payoff to the present time to find the arbitrage-free price of the option. Even though this approach is theoretically correct, it falls short because there is no rigorous way to estimate the discount factor to be used¹ for present value calculation.

Fisher Black, Myron Scholes [66], and Robert Merton [67] circumvented this problem in 1972 by setting up a hedged portfolio made of the underlying stock and the

¹ The option is not risk-free and therefore using the risk-free rate of return is not appropriate. The risk associated with the option is also different from the risk associated with holding the underlying, so a discount factor associated with the underlying would not work either.

option itself such that its return is exactly the risk-free rate of return. Knowing the value of the portfolio and using its riskless property ensures that its expected value can be discounted to the present time using the risk-free rate of return (alternatively, a *replicating portfolio* can also be constructed using both the underlying stock and the risk-free asset such that it exactly replicates the payoff of the option). APPENDIX G: actually provides a description of the Black-Scholes-Merton model as well as a derivation of the acclaimed Black-Scholes formula for pricing European call options for which Robert Merton and Myron Scholes earned the Nobel Memorial Prize in Economic Sciences (*Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel*). The beauty of the Black-Scholes-Merton setup is that it enables the derivation of a deceptively simple closed-form expression for the price of a European option.

For real options, similar techniques are used to value staggered corporate investments subject to numerous uncertainties. In this case, an initial investment outlay is required to acquire the right to exercise another investment later down the road. Using the development of a commercial airliner presented previously, the initial investment is the cost associated with funding a market research or even the conceptual design phase of the development. With this investment done and if the market environment favors a further development of the aircraft, this initial expenditure provides the company the possibility to fund the next stages of the development, namely the preliminary and detailed design phases. In this case, the initial investment is the value of the option to further develop the aircraft, the investment required to fund the preliminary and detailed design phases of development is the strike price, and the uncertain present value of the aircraft program cash flows is the uncertain underlying price. Similarly to financial options, these real

options come in many different flavors but always have something in common: there is flexibility at one point to alter the forecasted development, whether it be to scale-up the development, to delay the development, or to abandon the development as illustrated in Table 13.

Table 13: Different types of real options

| Option Type | Description | Example |
|---------------------------|---|---|
| Delay option | Possibility to fund only the first phase of a development in order to wait for the realization of uncertainty before committing to further developments | Funding market research before committing to the development of an aircraft |
| Growth option | Funding a development program enabling the option to be further extended | Development of a family of aircraft including several derivatives |
| Abandonment option | Possibility to resell the assets if the uncertainty is realized in unfavorable directions | Possibility to sell technology research patents |
| Compound option | Possibility to fund only the first phase of a development in order to wait for the realization of uncertainty before committing to a second phase and later to a third phase of development | Multi-phase aircraft development program with decision toll-gates |

Shortcomings of current real options approaches

Real options analysis has many advantages stemming from its ability to capture the value of flexibility offered to managers over the course of a development program. On the other hand, most of its shortcomings stem from the fact that practitioners have tried to transpose a methodology developed by and for the quantitative finance industry to capital budgeting problems.

The first type of issues is related to the assumptions underlying the Black-Scholes-Merton models. These assumptions state that the underlying stock follows a geometric Brownian motion with constant volatility and constant drift. The geometric Brownian motion is generally accepted as a “good enough” approximation for the return

of stocks traded in the financial markets but might not be adequate to model the value of research and development programs. There is therefore a need to depart from the traditional Black-Scholes model to be able to accommodate more diverse stochastic processes, and especially those that are more relevant for technologically-driven ventures.

The second type of issues is more fundamental and is related to the nature of the underlying. In the case of real options, the research and development program is not a traded asset and therefore there is no consensus regarding its fair market value. Instead, the projected cash flows stemming from the program are the best-guess estimates of analysts and this might result in a skewed analysis. In addition, the market might be incomplete in that it might not be possible to fully replicate the evolution of the value of the underlying research and development program behavior in all possible states of the world. Therefore, the Black-Scholes-Merton model which is articulated around a hedged position using a replicating portfolio might not be adequate, and the valuation of real options may fall under the more generic umbrella of option valuation in incomplete markets.

3.1.6 Valuation methods: a summary

Table 14 summarizes the main strengths and weaknesses of the methods previously described. From this table, it appears that none of the method is perfect to perform valuation of long-term and highly uncertain research and development programs. However, it stands out that real options techniques are more adapted for the valuation of research and development programs in the aerospace industry given their highly uncertain and long-term nature. In addition, real options valuation provides a good platform to build upon and improve the current valuation state-of-the-art. This

comparison of various valuation techniques allows stating one of the research hypotheses underlying this work:

Hypothesis 1—Real options for valuation with flexibility and uncertainty

In the context of aerospace research and development programs, real options methodologies enable the valuation of business cases and the development of value-driven design frameworks accounting for the value created by managerial flexibility in an uncertain environment.

Table 14: Strengths and Weaknesses of a Selection of Valuation Methods

| | Payback Period | Discounted Cash Flow | Internal Rate of Return | Modified Internal Rate of Return | Real options Analysis |
|--|----------------|----------------------|-------------------------|----------------------------------|-----------------------|
| Ease of implementation by analysts | + | + | + | + | - |
| Applicability to R&D program valuation | - | + | + | + | + |
| Applicability to long-term program valuation | - | - | + | + | + |
| Valuation handles well uncertainty | - | - | - | - | + |
| Valuation includes flexibility | - | - | - | - | + |

3.2 Marketing and competitive analysis methods

A competitive analysis is a marketing and strategic management analysis aimed at assessing the strength and weaknesses of both current and potential competitors. This

analysis is aimed at formulating a product development strategy that is robust with regards to uncertain moves by the competition and optimal for the economic viability of the project under review. In the context of aerospace engineering, a competitive analysis is a critical step during a product development cycle for several reasons. First, aircraft developments are long-term endeavors and it is likely that the competitive landscape will evolve in the meantime. Therefore, properly assessing beforehand the impacts of potential competitor moves is critical to the economic viability of the program. Next, aircraft developments are multi-billion dollar ventures which tie the funding abilities of manufacturers for a long time and preclude them from substantially altering the product development cycle once it reaches the detailed design phase. It is therefore paramount to ensure that the product meets the customer requirements and that there is enough space for multiple competitors within a market niche before actually committing to it. Finally, the competitive analysis provides decision-makers with preliminary information regarding the market reaction in various competitive scenarios. In turn, this enables the formulation of sales volume estimates helping the construction of scenario-based business plans and the estimation of profitability.

3.2.1 SWOT analysis

The SWOT analysis has disputed origins (often attributed to Albert Humphrey as part of his *Team Action Model* research) and stands for the evaluation of Strengths, Weaknesses, Opportunities and Threats involved in a business venture. It involves specifying the objectives of the business venture and listing internal and external factors that are both favorable and unfavorable to achieving these objectives. This listing is done to investigate the adequacy between the internal and external environments, to bring the

company in balance with the external environment, and to ensure it remains so over time. The first step is therefore the collection of relevant data to assess the capability of the company. This data is next sorted into the four categories mentioned previously with Strengths and Weaknesses generally stemming from the company itself, and Opportunities and Threats usually stemming from the outside. For each business venture under review, the information is presented in a *SWOT* matrix such as the one in Table 15 to check the adequacy between the objectives and the strengths. Finally, the last step consists in incorporating this analysis in the strategic decision process to ensure that the corporate strategy is in balance with these four attributes.

Table 15: SWOT Matrix

| | Beneficial for Company Objective | Detrimental to Company Objective |
|-----------------------------|---|---|
| Internal Origins | Strength | Weaknesses |
| External Origins | Opportunities | Threats |

Strength and shortcomings of the SWOT analysis

The SWOT analysis is a precursor to decision making in the strategic planning process to ensure the current and future endeavors are in line with the company's core competencies and strengths. This analysis is typically performed by a panel of experts leading to factual and informed conclusions. It remains however qualitative and is insufficient to judge the economic viability of a strategy.

3.2.2 Five forces analysis

The five forces analysis is a competitive assessment proposed by Michael Porter [69] in 1979 to help a company better position itself within an industry. As its name

implies, it is articulated around five forces, the collective strength of which ultimately determines the profit potential of an industry. These are the threat of new entrants, the threat of substitute products, the bargaining power of suppliers, the bargaining power of customers, and the intensity of rivalry amongst competitors. The role of management is therefore to find a position and then steer the company towards this position where the company is most likely to defend itself successfully against these threats.

The threat of entry is related to new competitors wishing to enter the market and gain market share. The seriousness of this threat depends both on the natural barrier to the entry and the reaction of established players. In the aerospace industry, the threat of entry is quite low due to the capital investment required as well as the steep learning curve to reach a point where products and services become competitive. Political impetus may lower this entry barrier by subsidizing research and development but cannot remove it completely. In the context of civil aircraft development, the availability of maintenance networks around the world to quickly service and ensure continuous and smooth operations of airliners is an issue that new entrants keep facing.

The threat of substitute products is related to their ability to offer price-performance trade-offs sufficient to entice customers to switch. Unless the legacy product can differentiate itself sufficiently from the substitute, it is likely that the substitute will get some market share and therefore limit profits. In the context of the aerospace industry, few substitutes exist to commercial aircraft with the exception of high speed rail.

The bargaining power of customers is related to their ability to put the firm under pressure when they have the ability to switch manufacturers. This is prevalent when the

buyer concentration to firm concentration ratio is quite low, when buyers purchase in large volumes, and when the products are undifferentiated and can easily be substituted by other products from the competition.

The bargaining power of suppliers is related to their ability to raise price without the possibility for the manufacturer to pass this price increase to the final consumer in elastic markets. This characteristic is mostly prevalent in industries where suppliers are few, where their products are differentiated, and where the switching costs from one supplier to another are high.

The last force identified by Porter is the intensity of rivalry which may take several forms such as price competition, race for first entry into service, and advertising. The strength of this last force is related to the number of competitors present in the market as well as to their relative sizes. The absence of differentiation between the products offered and the absence of growth potential within the market further exacerbates this rivalry. For civil aircraft manufacturers, the competition is intense despite the limited number of competitors. Indeed, each of them is fighting “tooth and nail” to gain market share which ensures a stream of revenues down the road with maintenance services and the sale of spare parts.

Strength and shortcomings of the five forces analysis

Porter’s five forces analysis provides a framework to perform a competitive assessment to check the positioning of a company within the industry sector and to ensure that the company avoids as much as possible intense competitive pressure. It is therefore a good starting point to evaluate a company’s strategy. Unfortunately, this is a qualitative analysis which does not help in estimating or forecasting what the outcome would be in

terms of sale volumes, market share, and overall profitability of a particular competitive scenario.

3.2.3 Game theoretic analysis

Game theory is a means of approaching, analyzing, and optimizing decision making problems featuring several parties, each with a rational behavior but possibly conflicting interests. When the analysis includes competitors and alternative product developments, then it is suitable and pertinent for use in strategic planning and to substantiate decision making. Even though there is evidence of prior use of game theoretic rationale in economics [70] and in evolutionary sciences [71], game theory as it is known today was formalized only in 1944 by mathematicians John von Neumann and Oskar Morgenstern in the *Theory of Games and Economic Behavior* [72].

A game may be a model of a competitive situation, and game theory is a set of mathematical methods for analyzing these models and selecting optimal strategies. Even without complete knowledge of other stakeholders' decisions or resources, game theory is useful for enumerating the decisions available and for evaluating these options or "moves". When a company's investment decisions are contingent upon the competitors' moves, it becomes a helpful tool in evaluating strategic decisions because it includes a means of predicting how competitors will behave. Of interest are competitive games which may be described in a four dimensional space with the players, the actions available to them, the timing of these actions, and the payoff structure of each possible outcome. Enumerating these elements may be done differently depending on the situation. Two popular choices used for simple games are the normal form, or matrix

format, and the extensive format. These are shown for the famous prisoners' dilemma respectively in exhibit (a) and (b) of Figure 12

Once a competitive game is defined, the purpose of a game theoretic analysis is to solve for an equilibrium concept, of which the Nash equilibrium is the most famous. This concept was first formulated by John Nash [73] in 1950. "The Nash equilibrium is a profile of strategies such that each player's strategy is an optimal response to the other players' strategies" [74]. In other words, the quest for a Nash equilibrium is an optimization process performed in the action space which searches for a set of actions and reactions from which none of the competitor has any incentive to deviate [75]. To illustrate the concept of the Nash equilibrium, the prisoners dilemma presented in Figure 12 may be used. In this case, the couple (Strategy 2, Strategy 2) is the only Nash equilibrium from which two non-cooperating players have no incentive to deviate. It is dominated by the couple (Strategy 1, Strategy1) but this couple is not a stable equilibrium as the two players both have an incentive to deviate and to switch strategy.

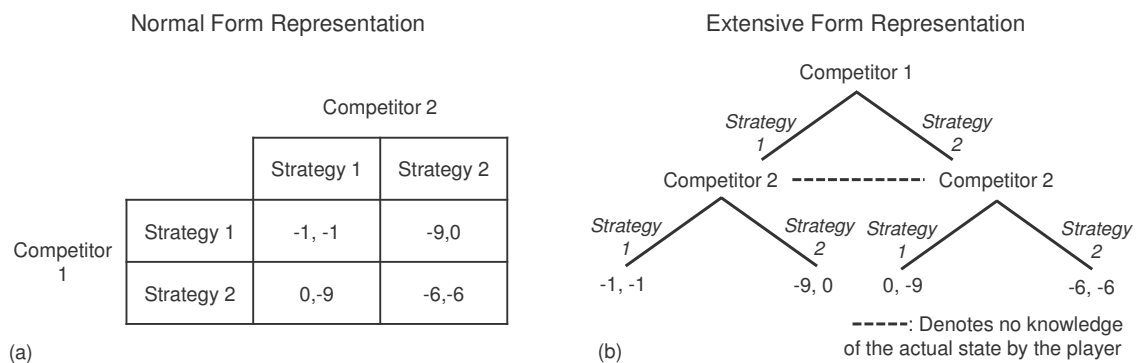


Figure 12: Normal form and extensive form representation of a strategic game

More complex scenarios may be analyzed using a game theoretic approach with different equilibrium concepts such as subgame-perfect Nash equilibrium, Bayesian Nash equilibrium, and perfect Bayesian equilibrium. These equilibriums will be introduced as

necessary over the course of this dissertation and an interested reader is referred to Gibbons [76] for a more in-depth treatment of these equilibriums.

Strength and shortcomings of the game theoretic analysis

Game theory has been widely embraced by both academics and practitioners. Game theoretic related work has been linked to eight different Nobel Memorial Prizes in Economic Sciences. The success of this analysis can be attributed to several factors. The first one is its ability to perform both qualitative and quantitative competitive analyses which are in great demand in an increasingly competitive world. The second is the versatility of the method which can handle competitive problems ranging from simple simultaneous games, to more involved dynamic games with incomplete information, and to signaling games. Another reason for this success is the transparent implementation combined with a theoretically sound foundation leading to a wide acceptance amongst academics and therefore a wide exposure to future practitioners.

However, game theory relies heavily on the concept of rationality which, in traditional economic models, is the maximization of utility. As a result, players will always act to gain as much as possible regardless of how these actions affect other players. But is this really the case?

It is indeed not obvious that economic agents always take rational decisions and French economist Maurice Allais [77] showed that sometimes they do not make decisions that appear rational according to the prevalent expected utility model. In particular, behavioral patterns have shown both a *certainty effect* whereby decision-makers overweigh outcomes that are considered certain relative to outcomes that are merely probable, and a *reflection effect* whereby decision-makers exhibit a shift from risk-aversion for

positive prospects to risk-seeking for negative prospects. These effects imply a departure from the utility-maximization principle of neoclassical economics that defines rational decision making and that underpins game theoretic approaches.

Furthermore, Nash equilibriums do not always yield *plausible* solutions that are likely to be observed in the real world.¹ Observing this conflict between predictions from rational models and observed behaviors, economists have developed less stringent models of rationality helping to connect the rational and the psychological. Models of bounded rationality [78] [79] argue for instance that the rationality of economic agents (decision-makers in particular) is limited by the information they have, the cognitive limitations of their minds, and the finite amount of time they have to make a decision. Other economic models assume that only a sufficiently large number of economic agents can be approximated to act rationally.

3.2.4 Prospect theory

Further improvements to the game theoretic analysis have been made using behavioral economic theories. Daniel Kahneman and Amos Tversky are experts in this field and have coined the term “prospect theory” in their seminal paper [80] to describe a new behavioral theory that may be used for a behavioral game theoretic approach.

Prospect theory aims at describing how economic agents chose between probabilistic alternatives that involve risk and uncertainty when the probabilities of outcomes are known. It recognizes that decisions are based on judgments which are biased assessments about the external state of the world and that decisions are

¹ Repeated prisoner dilemma

challenging because of the difficulty of assessing their consequences and because of the conflicting internal trade-offs they usually require. The theory addresses how these choices are *framed* and *evaluated* in the decision making process. The authors argue that there are two phases in the choice process, an early phase of *editing* and a subsequent phase of *evaluation*.

The *editing* phase is a preliminary analysis of the proposed prospects often leading to a reformulation and a simplification of their representations. Several operations transform the outcomes and probabilities associated with these prospects in a fashion that mimic how decision-makers actually process the information. The first type of operation stems from the observations that economic agents usually do not formulate the outcome of a choice in terms of absolute values or magnitude but rather as relative values with regards to something that they are familiar with. During this operation, the prospects are therefore reformulated in terms of gains and losses with respect to a *reference point* which is usually the baseline strategy or “business as usual”. The second and third types of operations are combination and segregation. They aim at simplifying the situation by either lumping together prospects with similar outcomes or, on the opposite, segregating prospects between their risky component and their riskless component. Next is the cancellation operation which consists in discarding components that are shared by all the prospects. Finally, another round of operations is the simplification which consists in discarding extremely unlikely outcomes, and the detection of dominance which consists in scanning alternatives to detect dominated ones that are rejected right away.

The *evaluation* part is the second part of the analysis during which the edited processes are evaluated in order to select the one with the highest value. This evaluation

is done using two scales: one associates a decision weight to the probabilities of each prospect, while another one assigns a subjective value to each outcome. Going further into the details of the generation of these scales is beyond the scope of this thesis. One aspect, the shape of the value function is quite interesting however. First, the value curve is “centered” on the reference point obtained earlier which determines whether the prospect outcome is a relative loss or a relative gain. Next, using the *reflection* principle which states that economic agents have a different attitude with regards to losses and gains, the curve exhibits risk-aversion for the profit section (concave shape) while it exhibits risk-seeking for the loss section (convex shape). This is in stark contrast with usual utility curves. Finally, another salient characteristic is that the losses are perceived as more detrimental than gains. This leads to a steeper slope for the value curve in the losses quadrant than in the gains quadrant as shown in Figure 13.




Figure 13: Notional value function representation

Strength and shortcomings of the prospect theory

The prospect theory proposed by Kahneman and Tversky is a descriptive model.

It tries to model real-life behaviors as observed in many controlled clinical experiments

and therefore exhibits ad hoc patterns consistent with some of the critics of purely rational utility-based models. In particular, it accounts for the fact that economic agents have different risk attitudes with regards to losses and gains, exhibiting risk-aversion in choices involving sure gains and risk-seeking in choices involving sure losses. It also models the fact that losses are considered more penalizing than gains.

However, being a descriptive model, it lacks a strong theoretical foundation as well as practical ways of estimating both the value function and the weighting function required for the evaluation step.

3.2.5 Competitive methods: a summary

Table 16 summarizes the main strengths and weaknesses of the methods previously described. From this table, it appears that none of the methods is perfect to perform a competitive analysis. Both the simpler game theoretic approach and the behavioral game theoretic approach are able to perform a quantitative analysis which will help substantiate a business plan. On one hand, a game theoretic approach presents a simpler and more transparent way to perform the analysis at the cost of assuming perfect rationality of decision-makers. On the other hand, the behavioral game theoretic approach proposes an alternative to better model human behavior at the cost of defining complex and potentially inaccurate value and weight functions. It is therefore believed that a simpler game theoretic analysis will better model the highly analyzed and highly substantiated choices made by decision-makers in the aerospace industry.

Table 16: Strengths and Weaknesses of Competitive Analysis Methods

| | SWOT Analysis | Five Forces Analysis | Game Theoretic Analysis | Behavioral Game Theoretic Analysis |
|--|---------------|----------------------|-------------------------|------------------------------------|
| Ease of implementation by analysts | + | + | + | - |
| Ability to handle qualitative analysis | + | + | + | + |
| Ability to handle quantitative analysis | - | - | + | + |
| Handle observed economic agent behaviors | - | - | - | + |

Comparisons between these competitive scenario investigative techniques allow the statement of the second research hypothesis underlying this work:

Hypothesis 2—Game theory for investigation of economic robustness with competition

In the context of aerospace research and development programs, game theory methods enable transparent and traceable analyses that allow decision-makers to investigate the economic robustness of selected technology and product development streams in a competitive environment characterized by uncertain moves by the competition.

CHAPTER 4: REAL OPTIONS THINKING

4.1 Genesis and challenges

In the previous chapter, some of the methodologies applicable to decision making problems for capital investments in aerospace research and development programs were presented. Some candidate methods were down-selected on the assumption that they *conceptually* fit the requirements of this research endeavor. Still, these methods come in many different flavors depending on the application and more generally the context of their use. In this chapter, a thorough review of the real options literature is presented to identify which methods are most suitable for this research. Being a rather new paradigm, the real options thinking field is quite fluid with an ever changing state-of-the-art.

4.1.1 *Borrowing a paradigm from the finance industry*

There is no doubt that real options inspired methodologies present an attractive concept for (scarce) capital allocation budgeting problems due to their abilities to better mimic the decision processes that take place as uncertainty unfolds. These methodologies inspired by the concept of contingent claims in finance enable valuations that account for the flexibility offered to management to react, update business plans, and change tactical plans to steer projects into profitable directions. Indeed, similarly to a financial option which is the right but not the obligation to exercise a predefined action within an allocated timeframe, a real option is the right but not the obligation to *take action*. ‘Take action’ is purposefully a vague term as it encompasses many different notions such as

abandoning a research and development investment, continuing the funding of a staggered research and development investment, expanding a promising research and development investment, or finally, deferring a research and development investment until conditions improve. This ability to better relate to what is actually happening daily within companies has been the driver for most of the research in the field of real options. Krychowski [81] reports that the literature on real options has increased exponentially since Myers [82] first coined the term in 1977. Moreover, real options inspired methodologies have been used in the aerospace industry for many different applications: valuation of aircraft purchase option at *Airbus* [83], valuation of adaptability in aerospace systems [84], investment under uncertainty in the air transportation infrastructure [85], and aircraft development investments at *Boeing* [86] [87] and *Embraer* [88].

4.1.2 An interesting concept harder to implement in practice

Usual assumptions in the framework of Black, Scholes and Merton

Many of the early applications of real options theory revolved around the transposition and subsequent use of Black-Scholes inspired formulae to value corporate investments featuring flexibility. In 1998, Luehrman [89] described a step-by-step methodology in the Harvard Business Review to value phased-investment opportunities using the Black-Scholes formula for call options. The application case was the evaluation of a growth option opportunity by a chemical company wishing to expand its production facilities. Later, Shank et al. [90] use the Black-Scholes-Merton model and the resulting call option valuation formula to estimate the value of investing in internet infrastructures to support the potentially growing *e*-business. More recently, Pinon [91] uses the Black-

Scholes formula to value flexible technology investments in underutilized regional airports to relieve capacity-constrained airports in metroplex areas.

Despite the breadth of applications using the Black-Scholes setting and the Black-Scholes formula, there is generally little effort to justify whether the model and its assumptions are actually suitable. In Chapter 3.1.5, the real options analogy was first introduced while APPENDIX G: presented the setting of the Black-Scholes model and the derivation of the Black-Scholes option valuation formula. In Table 17, the main assumptions underpinning the model are reminded, while Table 18 attempts to translate these assumptions for use in real options valuation.

| | |
|-------|---|
| (i) | The market has no arbitrage |
| (ii) | The market has no fees or trading costs |
| (iii) | The asset does not pay any dividend |
| (iv) | The asset follows a Geometric Brownian Motion |
| (v) | Both volatility of asset and risk-free interest rate are constant |
| (vi) | Asset and bond may be bought in any quantity, including negative amount and fractions |
| (vii) | Claim can only be exercised at maturity |

Table 17: Main assumptions underpinning the Black-Scholes model

| | |
|--------|---|
| (i') | <i>Not applicable</i> |
| (ii') | <i>Not applicable¹</i> |
| (iii') | The underlying project has no value leakage |
| (iv') | The underlying project value follows a Geometric Brownian Motion |
| (v') | Volatility of underlying project value and risk-free interest rate are constant |
| (vi') | <i>Not applicable</i> |
| (vii') | Taking action to continue or change course can only be made at maturity |

Table 18: Translating main assumptions for real options valuation using the Black-Scholes model

A closer look at the set of assumptions leads to their classification in two categories. One set of assumptions is made to ease the derivation of the Black-Scholes

¹ It is sometimes argued that buying a real option involves some form of trading cost (cost to invest in human resources, cost to invest in infrastructure). Nevertheless, in many cases these costs can be incorporated into the required investment to fund one phase of the business venture.

formula. It leads to an academically pleasing yet somewhat unrealistic setting. The other set of assumptions relates to the foundation of the methodology. Practice has shown that most of the assumptions belonging to the first category may be somewhat relaxed and the Black-Scholes formula can either be adjusted easily, or can be left as is and yet provide reasonably accurate solutions.

For instance, assumptions (ii) and (vi) are unrealistic at best as there are few, if any, market devoid of any trading costs that would materialize by a zero spread between the bid and ask prices. Besides, even if it were possible to short sell and take a negative position in a security, it is not usually possible to take a fractional position in a security as assumed by the model. Assumption (iii) relates to the modeling of the underlying security and whether or not this security is issuing dividends over the course of the option life. This assumption may be relaxed by using the Black-Scholes formula with dividends. Assumptions (iv) and (v) relate to the modeling of the underlying asset behavior and an ad-hoc time-series analysis needs to be performed to ensure the model is not misspecified. Most asset prices do not follow a true Geometric Brownian Motion as evidenced by the volatility smile for deep in-the-money and deep out-of-the-money options and by the heteroscedasticity displayed by most financial time series. The deviation is nevertheless not large enough to invalidate the results and the Black-Scholes formula still provides a good enough approximation.

In a real options environment, the assumptions related to the dynamics of the underlying asset are directly translated into assumptions related to the dynamics of the value of the underlying project featuring flexibility. Consequently, as long as the project value follows a Geometric Brownian Motion as prescribed in assumption (iv) and as long

as the volatility and risk-free rate are constant over time as prescribed in assumption (v), these assumptions are still valid. Similarly, if the flexibility offered to management in the underlying project can be modeled as a European-type real option, then assumption (vii) still holds. Finally, if the project does not lose some of its value over time (no value leakage due for instance to the cost to defer a decision), then assumption (iii) regarding the dividend payments also holds true. If not, a modified Black-Scholes with dividends framework may be used as was proposed earlier for financial options.

Assumptions (i), (ii) and (vi) are more difficult to translate as they relate to the ability to replicate any claim with a self-financing replicating portfolio. Indeed, the Black-Scholes model relies on the assumption that in a complete market, it is possible to replicate every claim with an arbitrary payoff using a self-financing portfolio consisting of a dynamically adjusted linear combination of the basis assets present in the market. Therefore, the no-arbitrage price in a complete market can be calculated using this self-financing replicating portfolio. Assumption (i) ensures that, whatever the state of the world, the self-financing portfolio having the same payoff as the claim must have the same price. Assumption (ii) ensures that no loss occurs whenever the replicating portfolio is constructed and continuously adjusted to replicate the claim. Finally, assumption (vi) ensures that the claim is attainable, which means that it is always possible to replicate the claim using a linear combination of assets present in the market. This includes the ability to short some assets (ability to borrow and sell these assets which mathematically means the ability to have assets within the portfolio that have a negative weight) and the ability to have fractional quantity of some assets (weight of some assets in the portfolio need not be integers which allows to better track the claim payoff).

These assumptions are problematic for a real options application because the underlying project value is not a traded asset in any market. Therefore, there is no arbitrage-free price for the underlying project and therefore, there is no guarantee of a single price for the replicating portfolio made up of the underlying project and some other securities. In addition, it is not obvious that the market can be complete. In fact, the market is likely to be incomplete and the claim is most probably not attainable. This means that its payoff cannot be replicated with a self-financing portfolio made up of a combination of the basis assets in the market. Finally, even if these two assumptions still hold true, it is not conceptually possible to construct a replicating portfolio with no restriction on the ability to short sell and no restriction on the ability to take fractional positions on the underlying project: after all, how to borrow half of a project and sell the other half? With a major assumption underpinning the Black-Scholes formula derivation violated, any use of the Black-Scholes formula for real options applications now looks suspicious. Fortunately, beyond the original hedge and partial differential equation pioneered by Black and Scholes, some other techniques may be used to value financial options.

Relaxing assumptions with the martingale approach

Another popular approach is the martingale¹ approach initially proposed by Cox and Ross [92]. It is more mathematically involved as it borrows the concepts of probability measure, equivalent probability measure, and change of probability measure

¹ A martingale is a stochastic process whose current value is its expectation. In finance, the discounted price of an asset is assumed to be a martingale since the asset current price is its future expected discounted price. Mathematically, a stochastic process S_t is a martingale with respect to a probability measure \mathbb{Q} if and only if the following two conditions are satisfied: $E_{\mathbb{Q}}(|S_t|) < \infty$ and $E_{\mathbb{Q}}(S_t | \mathfrak{F}_s) = S_s$ for $s \leq t$

from probability theory. In essence, the martingale approach is based on five steps to value claims. In the first step, the Girsanov theorem is invoked to change the probability measure from the physical probability measure to an equivalent probability measure \mathbb{Q} such that the discounted stock process is a martingale under this new measure. In the second step, a stochastic martingale E_t under the new probability measure \mathbb{Q} hitting the discounted claim value X at expiration is constructed using a conditional expectation $E_t = E_{\mathbb{Q}}(B_t^{-1}X|\mathfrak{F}_t)$. In the third step, the martingale representation theorem is invoked to find a previsible process ϕ_t that relates the variation of the martingale E_t hitting the discounted claim at expiration to the discounted asset martingale $B_t^{-1}S_t$, which is expressed as follows: $d(E_t) = \phi_t d(B_t^{-1}S_t)$. This step is powerful because it allows to mathematically link the evolution of the discounted claim process to that of the discounted asset process through a stochastic process that is previsible¹. The fourth step is to construct a self-financing portfolio made up of the discounted asset in quantity ϕ_t and the discounted bond in quantity ψ_t such that its value exactly replicates the value of the discounted claim process E_t . This portfolio therefore hits the value of the discounted claim $B_T X$ at expiration and its value is expressed as $V_t = \phi_t B_t^{-1}S_t + \psi_t B_t^{-1}B_t$. The non-discounted counterpart portfolio made of real asset S_t and real bond B_t is given by $X_t = \phi_t S_t + \psi_t B_t$. This portfolio hits the value of the claim at expiration and it is self-financing since the change in its value is given by $dX_t = \phi_t dS_t + \psi_t dB_t$, which is exactly the change in the value of the assets contained within the portfolio. The final step

¹ A previsible process is a process which only depends on the information available up to the current time, but not on any future information. The concept is interesting to construct portfolios since the existence of previsible processes, each of which models the weight of one asset within the portfolio, ensures that the portfolio can be constructed in real time with contemporary information.

invokes the no-arbitrage condition to establish the current price of the claim X_t which must be the current value of this replicating self-financing portfolio. This is also exactly the expectation under probability measure \mathbb{Q} of the discounted claim $X_t = B_t E_{\mathbb{Q}}(B_T^{-1} X_T | \mathcal{F}_t)$. For a more thorough analysis of the martingale approach, the reader is referred to the textbook by Baxter and Rennie [93].

The martingale approach presents several advantages for derivative pricing. The first advantage is that the claim price is formulated as an expectation instead of a partial differential equation. This allows the use of Monte Carlo simulations to numerically compute the expectation and therefore estimate the claim value [94]. Another advantage of the martingale approach is that no restrictions are made regarding the claim payoff except that it is attainable. This means that the claim payoff can be arbitrary and even path-dependent. This more general setting also allows the relaxation of assumption (vii) that restricted the Black-Scholes partial differential equation approach to European-type claims. With the martingale approach, American, Bermudan, and Asian types of claims may be priced, provided the claim payoff is attainable. For instance, an American call option would be valued by maximizing the expectation over all possible exercise times τ up to the option maturity and would simply result in the following expectation computation $X_t = \sup_{t \leq \tau \leq T} [B_t E_{\mathbb{Q}}(B_{\tau}^{-1} X_{\tau} | \mathcal{F}_t)]$. A third advantage of the martingale approach is that it does not require that the underlying asset volatility and the risk-free interest rate remain constant over time, thus relaxing assumption (v). Indeed, the martingale approach requires only the computation of an expectation and therefore, the distribution of the claim value at any time up to the expiration is sufficient. In contrast, using the original approach, time-varying parameters would lead to a partial differential

equation featuring time-dependent coefficients and there is no clear way to proceed further and solve that equation.

Relaxing assumptions with the change of numéraire approach

Another approach is the change of numéraire¹ approach initially proposed by Jamshidian [95] for bond option pricing and generalized later by Geman et al. [96] for derivative pricing. It is very similar to the martingale approach in that it uses some change of probability measure and requires the computation of an expectation. It may thus be seen as a generalization of the martingale approach. In the martingale approach, all processes are discounted using a risk-free bond. This discounting may be interpreted as a normalization step since each and every process is now defined relative to the risk-free bond process. In this relative pricing environment, the risk-free bond plays the role of the valuation standard, also known in mathematical finance as a numéraire or a deflator. However, nothing prevents the use of a different numéraire or in other words, nothing prevents the use of a different standard to perform valuation: what if an asset price for instance is used as the new normalizing standard or the new numéraire? Although this approach seems complicated, this is in fact very much akin to a valuation made in two different currencies for which the unitary bill (or coin) of each currency represents different numéraires.

Like the martingale approach, the change of numéraire approach can be decomposed into five main steps. In the first step, the Girsanov theorem is invoked to change the probability measure from the physical probability measure to an equivalent

¹ A numéraire is a basic standard by which value is computed. Acting as the numéraire is one of the functions of money which serves as the accounting unit to measure the worth of different goods and services relative to one another.

probability measure \mathbb{Q}' so that the risk-free bond process deflated by the asset price process B_t/S_t is a martingale under this new measure. In the second step, a stochastic martingale E_t under the new probability measure \mathbb{Q}' hitting the deflated claim value at expiration is constructed using the numéraire S_t and the following conditional expectation process is constructed $E_t = E_{\mathbb{Q}'}(S_t^{-1}X|\mathcal{F}_t)$. In the third step, the martingale representation theorem is invoked to find a previsible process ϕ_t that relates the variation of the martingale E_t hitting the deflated claim at expiration to the relative risk-free bond martingale $S_t^{-1}B_t$. This relationship is expressed as follows: $d(E_t) = \phi_t d(S_t^{-1}B_t)$. The fourth step is to construct a self-financing portfolio made up of the deflated risk-free bond process in quantity ϕ_t and the deflated stock in quantity ψ_t such that its value exactly replicates the value of the deflated claim process E_t . This portfolio therefore hits the value of the deflated claim X/S_t at expiration and its value is expressed as $V_t = \phi_t S_t^{-1}B_t + \psi_t S_t^{-1}S_t$. The non-deflated counterpart portfolio made of real asset S_t and real bond B_t is given by $X_t = \phi_t B_t + \psi_t S_t$. This portfolio hits the value of the claim at expiration and it is self-financing since the change in its value is given by $dX_t = \phi_t dB_t + \psi_t dS_t$, which is exactly the change in the value of the assets contained within the portfolio. The final step invokes the no-arbitrage condition to establish the current price of the claim X_t which must be the current value of this self-financing replicating portfolio and which is also exactly the expectation under probability measure \mathbb{Q}' of the deflated claim: $X_t = S_t E_{\mathbb{Q}'}(S_T^{-1}X_T|\mathcal{F}_t)$. In summary, the change of numéraire technique is very similar to the martingale approach as it follows the same steps and makes use of the same theorems. The only difference is that the equivalent martingale measure is not made using the risk-free bond deflator but rather another asset, which in this case, is the

underlying asset S_t since there are only two assets in the economy. A more thorough analysis of change of numéraire is given in Duffie [97].

Being similar to the martingale approach, the change of numéraire approach retains most of the advantages described earlier. It is however more general as the numéraires can be chosen so as to simplify expectation computations. This is helpful to neutralize one source of risk when the option payoff contains several different sources of risk. For instance, it is customary to use this approach when pricing under a time-varying random interest rate in case the price of a bond maturing at the option expiration is used as numéraire. A more relevant example for a real options application is the case of complex options on n different assets. In this context, the change of numéraire approach simplifies the problem by neutralizing one source of risk and therefore reducing the problem to $n-1$ different sources of risk. For instance, let's hypothesize an environment where a company nurtures a portfolio of two competing, uncertain, but potentially promising projects. The company does not have the financial resources to fully fund these two projects concurrently and consequently must choose at one point in time which one to pursue. This setting is very similar to a rainbow "call on max" option which gives the option holder the ability to select one of two assets to purchase. In this case, the change of numéraire is useful if the new numéraire is taken as one of these two assets. Indeed, the problem is reduced to a single source of risk: the asset used as numéraire has a relative or deflated process which is trivially constant and equal to one, while the other asset deflated process is a ratio of two stochastic processes which is treated as a single stochastic process [98].

In the preceding paragraphs, three approaches to derivative pricing were presented and Table 19 summarizes the steps involved in each of them. This is obviously not an exhaustive list as many other methods have been used including binomial trees [99], utility-based techniques [100] as well as empirical methods such as implied binomial trees [101]. Still, the methods presented have progressively relaxed some of the original assumptions of Black and Scholes which may be useful for real options analysis: the interest rate does not need to be constant, the volatility parameter does not need to be constant, the option does not need to be of European type, and the asset may have dividends. As a result, these methods grow in terms of applicability with the original hedge and partial differential equation approach being the most restrictive, while the martingale approach and the change of numéraire approach provide a richer and wider domain of applicability. Still, one fundamental assumption of the original Black and Scholes setting remains: the requirement that claims be attainable and markets complete, which underpins the ability to use no-arbitrage arguments and find replicating portfolios. This observation and quest for a rigorous option-thinking framework leads to the first sub-research question as follows.

Research Question 1.1—Creation of an option-thinking framework

In the context of uncertain corporate investment analysis, how can state-of-the-art option-based valuation methodologies be altered and improved upon to ensure their domain of application is consistent with their underpinning assumptions? More precisely, how can practitioners benefit from an option-thinking perspective while acknowledging the issues surrounding the no-arbitrage argument arising from the market incompleteness and the resulting inability to find replicating portfolios of real assets?

Table 19: Required steps for Black-Scholes formula derivation using three valuation techniques

| Hedge and PDE Approach | | Martingale Approach | | Change of Numéraire Approach | |
|--|---|---|---|--|--|
| <i>Black, Scholes – 1973 – [66]</i> <i>Merton – 1973 – [67]</i> | | <i>Cox, Ross – 1976 – [92]</i> <i>Harrison, Kreps – 1979 – [102]</i> | | <i>Jamshidian – 1989 – [95]</i> <i>Geman, El Karoui, Rochet – 1995 – [96]</i> | |
| Step 1 | Using Ito’s lemma, derive PDE describing evolution of claim value dX_t function of (t, S_t) | Step 1 | Find a probability measure \mathbb{Q} so that the discounted stock is a \mathbb{Q} martingale $E_{\mathbb{Q}}(B_t^{-1}S_t \mathfrak{F}_s) = B_s^{-1}S_s = Z_s$ | Step 1 | Find a probability measure \mathbb{Q}' with numeraire S_t so that B_t/S_t is a \mathbb{Q}' martingale $E_{\mathbb{Q}'}(S_t^{-1}B_t \mathfrak{F}_s) = S_s^{-1}B_s = Z_s$ |
| Step 2 | Construct self-financing portfolio of both underlying asset S_t and bond B_t having same dynamics | Step 2 | Construct a \mathbb{Q} martingale process for the discounted claim with value X at maturity $E_t = E_{\mathbb{Q}}(B_T^{-1}X \mathfrak{F}_t)$ | Step 2 | Construct a \mathbb{Q}' martingale process for the normalized claim with value X/S_T at maturity $E_{\mathbb{Q}'}(S_T^{-1}X \mathfrak{F}_s) = X_s/S_s$ |
| Step 3 | Identify drift and diffusion terms by Unique Decomposition Theorem and add boundary conditions | Step 3 | Find a previsible process ϕ_t such that: $dE_t = \phi_t dZ_t$ | Step 3 | Find a previsible process ϕ_t such that: $dE_t = \phi_t dZ_t$ |
| Step 4 | Solve PDE to yield option price using Fourier transform or change variables to transform PDE into heat equation | Step 4 | Construct self-financing portfolio V_t holding ϕ_t of discounted stock $Z_t = B_t^{-1}S_t$ and ψ_t of discounted bond such that it replicates E_t $V_t = \phi_t Z_t + \psi_t B_T^{-1}B_t = E_t$ | Step 4 | Construct self-financing portfolio V_t holding ϕ_t of normalized bond $Z_t = S_t^{-1}B_t$ and ψ_t of normalized stock such that it replicates E_t $V_t = \phi_t Z_t + \psi_t S_T^{-1}S_t = E_t$ |
| | | Step 5 | Establish no-arbitrage price X_t of claim X as present value of replicating portfolio $B_t V_t$ $B_t V_t = \phi_t S_t + \psi_t B_t = B_t E_t$ which is also $X_t = B_t E_{\mathbb{Q}}(B_T^{-1}X \mathfrak{F}_t)$ | Step 5 | Establish no-arbitrage price X_t of claim X as present value of replicating portfolio $S_t V_t$ $S_t V_t = \phi_t B_t + \psi_t S_t = S_t E_t$ which is also $X_t = S_t E_{\mathbb{Q}'}(S_T^{-1}X \mathfrak{F}_t)$ |

4.1.3 Substantiating real options thinking: the Marketed Asset Disclaimer

Substantiating the availability assumption of a “twin security” in the financial markets that can be used to perfectly replicate the value of the business prospect is

difficult. There is indeed little reason to believe that the value of a corporate investment, which is subject to both private and market risks, would exhibit over its entire life a perfect correlation with one particular stock in each and every possible state of the world. This is a weakness facing many real options methods since the lack of a twin security to construct a replicating portfolio precludes a priori the use of no-arbitrage arguments for pricing purposes.

In 2000, Copeland et al. [103] argue that in the absence of an explicit market-traded twin security, the value of the business prospect without flexibility and therefore computed as a net present value is the best known proxy for a traded security having perfect correlation with the corporate investment value. They state that *“the option pricing approach gives the correct value because it captures the value of flexibility correctly by using an arbitrage-free replicating portfolio approach. But where does one find the twin security? We can use the project itself (without flexibility) as the twin security, and use its NPV (without flexibility) as an estimate of the price it would have if it were a security traded in the open market. After all, what has better correlation with the project than the project itself? And we know that the DCF value of equities is highly correlated with their market value when optionality is not an issue. We shall use the net present value of the project’s expected cash flows (without flexibility) as an estimate of the market value of the twin security. We shall call this the marketed asset disclaimer.”*

In 2001, Copeland and Antikarov [104] restate this assumption as follows *“we are willing to make the assumption that the present value of the cash flows of the project without flexibility (i.e., the traditional NPV) is the best unbiased estimate of the market value of the project were it a traded asset”*.

The *Marketed Asset Disclaimer* or *MAD* assumption is extremely powerful: by acknowledging that a twin security probably does not exist in the financial market and by supposing that the best unbiased surrogate for this twin security is the *subjective* estimation of the business prospect value without flexibility, practitioners can now use this fictitious twin security to build a replicating portfolio and therefore use the no-arbitrage argument for the economic valuation. The assumption also implies that the net present value of the prospect is the best known unbiased estimate of the project's market value if it were a traded asset and that no-one can "arbitrage" this project valuation. Still, it is important to take a step back and not be carried over by this assumption. The assumption allows practitioners to bridge a gap in the real options analysis and to transpose a method applied for financial option valuation to corporate investments valuation. It states that, when no twin-security can properly be found and used to build a replicating portfolio, then the best *subjective* surrogate is the value of the investment itself. The word *subjective* carries a lot of weight as the net present value of a corporate investment relies on assessments, many of which are subjective. For an aircraft development application, these subjective inputs may be the expected market penetration stemming from the sale of a new more efficient aircraft, the extra revenues generated by these sales as well as the costs to develop, certify, and produce the new aircraft. Borison [105] indicates that the assumption "*ensures that the 'Law of One Price' is maintained internally between the investment and the options*" but that due to the subjective nature of the valuation "*arbitrage opportunities may be available between the corporate investment and traded investments if any traded investments are available.*" In other words, the MAD assumption only ensures that the valuation is internally consistent but

arbitrage opportunities may still exist if the investment valuation is biased and if some traded assets that can act as the twin-security are available.

Consequently, the MAD assumption is best used as a last resort if there are no traded assets that could be used in the replicating portfolio. Therefore, Copeland and Antikarov [106] advise analysts to rely primarily on capital markets to substantiate inputs in the prospect valuation since they believe that *“the analysis would be incomplete if it ignored information contained in available market prices.”* Borison [107] echoes this statement and argues that *“if investments are evaluated using subjective, non-market assessments of these risks, the possibility of arbitrage is introduced”* and that avoiding arbitrage possibilities requires that practitioners analyze *“relevant spot, future, and option prices to determine the prices that capital markets have already established for an investment’s public risks.”*

So, how can this piece of advice be implemented in practice? Let’s assume for instance a performance improvement package (PIP) that improves the fuel-burn of a turbofan engine. Much of the value of the package for an airline is derived from the lower fuel consumption and therefore the lower operating costs which are directly related to the uncertain market price of jet fuel. Much of the package value to the airline remains uncertain: if the jet-fuel spot price goes up, so does the value of the PIP; on the other hand, if the jet-fuel spot price goes down, so does the value of the PIP. To preclude the possibility of arbitrage, the analyst should closely examine jet-fuel futures contract that have already established a market price for the jet fuel at different horizons. By using several jet-fuel prices, each corresponding to a different time horizon and each derived

from the jet-fuel futures, the analyst has included as much market information as possible in the construction of the performance improvement package business case.

4.1.4 What about the dynamics of the underlying real assets value?

Most of the literature on financial options and real options assumes that the underlying stock or the underlying real assets are following a Geometric Brownian Motion (GBM). A mathematical description of the Geometric Brownian Motion and its main characteristics is provided in APPENDIX A: . For financial stocks, the Geometric Brownian Motion assumption relies on the proof provided by Nobel Memorial Prize in Economic Sciences laureate Paul Samuelson [108] who argues that “*properly anticipated prices fluctuate randomly*”, an argumentation echoed by Fama [109]. Later, Samuelson [36] suggests using Geometric Brownian Motion to model the price evolution of risky assets.

The model is interesting for several reasons. The first is its mathematical simplicity since it is parameterized by only two variables: a drift to account for the long-term evolution and a volatility to characterize the diffusion. The second is that prices remain positive, which is in agreement with limited liability of stakeholders. The third is that returns can be either positive or negative and are uncorrelated, which is in agreement with the efficient market hypothesis and with the fact that no-one should be able to predict future returns based on past performance. Despite its mathematical elegance and its widespread use in financial models, the Geometric Brownian Motion is a mathematical model. Like many models used to capture and simulate complex phenomena, it has several shortcomings such as the inability to explain fatter tails in

observed return distributions [110], observed autocorrelations [111], and observed heteroscedasticity [112].

For real options applications, the use of Geometric Brownian Motion is widespread and applied to many different problems such as oil field exploitation value, intellectual property value, and technology portfolio value. Kemna [113] uses a Geometric Brownian Motion to simulate the value of exploiting an off-shore oil field subject to uncertain commodity prices. Weeds [114] assumes that the value of a technological patent evolves according to a Geometric Brownian Motion. Pinon et al. [115] assume that the value of airport technologies follow a Geometric Brownian Motion driven by the uncertainty in transportation demand.

Despite this widespread use, the case for using Geometric Brownian Motion in real options applications is not clear-cut. Indeed, implicit in many applications is the fact that if the uncertainty follows a Geometric Brownian Motion, so does the business prospect value. This supposition is often made when dealing with prospects deriving their value from the price of an uncertain commodity (coal price, oil price, jet-fuel price...) or from an uncertain aggregated indicator (air transportation demand, market size...). A closer inspection reveals that this assumption is debatable for two reasons. First, it obviously requires that the uncertainty driving the value of the business prospect follows a geometric random walk. Better said, it requires that the Geometric Brownian Motion be a good enough approximation of the dynamics of these commodity prices. Unfortunately, this verification is seldom done by practitioners. Secondly, if the implication were to be true, it would require that the cash flows of the project conserve two things: the independence of the increments and the Gaussian nature of the distribution of increments.

There is no reason to believe that this is true, especially for complex cash flows that are not simple additions, subtractions, or multiplications of uncertain random quantities. Borison [105] argues that “*While there may be good arguments for GBM with respect to equilibrium prices in highly liquid, widely accessible markets, there is no reason to believe that subjective assessments [...] of the value of the underlying investment should follow GBM*”. This is because “*the assessed value of the underlying investments may be driven by specific events in specific time periods in a manner that looks nothing like random drift.*” Following this observation, there is a need to extend current real options methodologies to ensure they can handle non geometric random walks. Many popular real options methodologies, such as the one advocated by Copeland and Antikarov [104], are based upon the marketed asset disclaimer hypothesis reviewed previously and make use of binomial trees for the valuation of the real options. Thus, they assume a geometric random walk process for the underlying corporate investment value. It is believed that a more generic approach able to relax this assumption would be beneficial.

Research Question 1.1.1 — Enlarging the domain of applicability of real options

How can the domain of application of current state-of-the-art real options methodologies be extended to include corporate investments with value processes that do not follow classic geometric random walks?

4.2 Numerical recipes for real options

In the previous paragraph, the fundamental assumptions underpinning real options analysis have been reviewed. It is now time to investigate how a versatile analysis framework for real options thinking can be constructed. The framework shall be generic so as to be able to handle the wide spectrum of applications that real options practitioners may face while retaining most of the mathematical rigor required by the models and the assumptions underpinning these models.

4.2.1 *Three venues for real options evaluation*

In this section, three of the most common valuation techniques are presented. These include the partial differential equation approach, the lattice approach, and the Monte Carlo approach. Their respective strengths and weaknesses are highlighted to establish the most appropriate set of methods for corporate investment analysis in the aerospace industry.

Partial differential equation approach

The partial differential equation approach consists in solving the partial differential equation that represents the evolution of the real option value over time. The nature of the option and its payoff at expiration usually define the boundary conditions for the partial differential equation. Closed-form solutions to the partial differential equation may exist and are usually found by performing some changes of variables so as to transform the original equation into a simpler equation for which solutions are well known. This type of approach was chosen for instance by McDonald and Siegel [116] to estimate the value of delaying a corporate investment, and by Grenadier [117] to value

real options in leasing contracts. However, closed-form solutions do not generally exist for partial differential equations and therefore numerical approaches are used. Amongst them, finite-difference methods are popular and consist in both discretizing the time and price space while writing the differential terms of the equation in terms of central, forward, or backward differences. Numerical approximations to solve partial differential equations for real options valuation have been used for different types of applications. For instance, Majd and Pindyck [118] value the options to delay, slow, or speed-up sequential corporate investments, Bernardo and Chowdry [119] study how firms choose between different types of corporate investments, and Dias and Rocha [120] analyze the valuation of oil exploration concessions featuring some horizon extension possibilities.

For the partial differential equation approach to work, an explicit formulation of the underlying corporate investment process must be known. In the studies presented above, the underlying processes were either diffusion processes (McDonald and Siegel [116]), diffusion with leakage processes (Majd and Pindyck [118]), or diffusion with jump processes (Dias and Rocha [120]). The main challenge is that it is not always possible to find a well-known stochastic process that properly models the underlying corporate investment value. Another challenge is that calibrating these models is not trivial if little or no historical data is available to estimate the corresponding parameters: how to calibrate the volatility of a business prospect value following an assumed diffusion process if similar prospects have never been attempted before? In the aforementioned literature, there is little effort made to substantiate why a specific stochastic process is retained in the valuation models besides a generic “*we assume that X follows Y*”.

Although partial differential equations were initially embraced by academics, they have somehow fallen out of favor because they end-up being impractical. Indeed, as much as a discounted cash flow analysis is simple enough to be implemented in spreadsheets for use within a company, the valuation of a real option by solving a partial differential equation is harder to implement: it requires substantial mathematical skills to derive the differential equation itself and specific solvers to numerically resolve the equation. In the end, practitioners may reject the valuation that they consider a “black box” valuation.

Lattices and trees

Lattice methods have been undoubtedly some of the most popular methods for real- options pricing thanks to their visualization appeal and simplicity. The idea is to discretize continuous processes in many different time steps and then to restrict the evolution of the underlying asset by imposing a maximum number of outcomes in which the system may end-up at the following time step. Usually, this restriction is set to either two for binomial lattices, three for trinomial trees, and four for quadrinomial trees. There are many different flavors of lattices but the construction and use is very similar. Therefore, only the simplest binomial lattices will be covered in this dissertation. A more thorough description of lattice methods may be found in Cox, Stephen and Rubinstein [99].

Similarly to the fact that a normal distribution can be approached by a large number of repeated binomial experiments, the normal distribution of the asset return is approximated by many repeated binomial experiments whereby the underlying asset return is either going up or down. These methods are divided into two steps: the first is

the construction of a lattice representing the evolution of the value of the underlying business prospect, and the second is the evaluation of the real option value at each and every node of the tree. To do so, a lattice starting at the current present value of the underlying asset and extending until the maturity of the real option is constructed, as shown in Figure 14. At each time step, the value of the asset may either go up with a certain probability or down with the complementary probability. The choice of the up and down probabilities as well as the choice of the up and down tick sizes are made in order to match both the volatility of the underlying asset and to ensure that the option value asset price dynamics is risk-neutral. In other words, the lattice parameters are determined such that the up and down tick sizes (one variable since the down-tick size is the inverse of the up-tick size) match the asset price process volatility, and such that the up and down probabilities (one variable again since both probabilities sum to one) ensure that the discounted asset price process is a martingale.

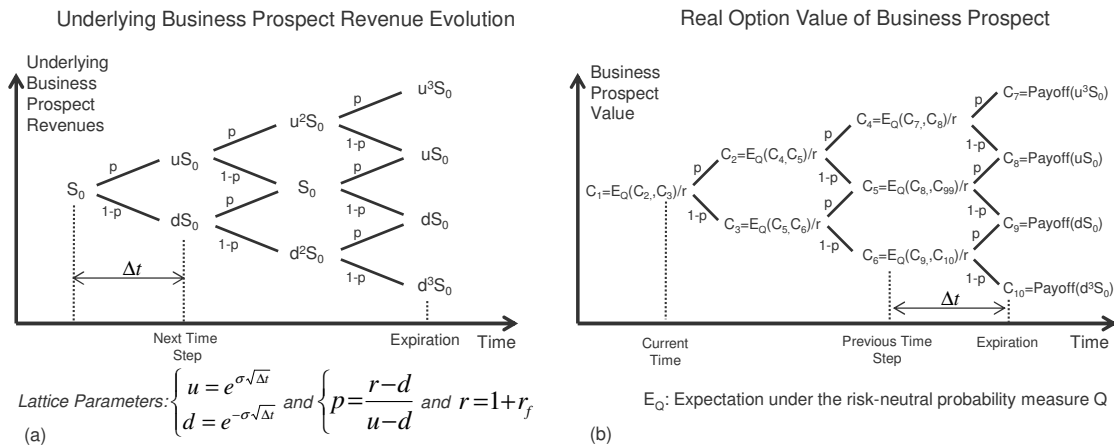


Figure 14: Real options pricing with binomial lattices

The next step is to superimpose the real option value to this graph by means of back-propagation. Back-propagation means that the real option value is first computed at

expiration using the payoff formula for the option under review and that the calculated values are then propagated backward in the lattice up to the present time. This step can vary somewhat depending of the type of option, whether it is a European, American, Asian, Bermudan, or any other type of option. For a European-type option, the value at expiration is estimated for each expiration-time node in the lattice using the value of the underlying business prospect previously computed. This value at expiration is then back-propagated to the prior nodes of the lattice by computing the expected discounted value of the real option. This step is repeated until the very first node in the lattice is reached and the current value of the real option is found.

Lattice methods have been widely used in the real options literature with many diverse applications. For example, Trigeorgis [121] uses binomial trees to value embedded options in leasing contracts such as the option to cancel, extend, or buy the leased asset. Stonier from *Airbus* [83] [122] uses binomial lattices to value real options embedded in aircraft purchase contracts for additional purchase rights, as well as for switching aircraft size. Lewis et al. [123] use binomial trees to value deferral options in research and development projects that are used to wait for more information to become available.

The popularity of lattice methods is mostly due to their transparency and ease of application. Indeed, once a stochastic process is accepted to model the underlying asset value, the following steps consisting in the construction of the lattice and the valuation of the real option are both easy to implement and straightforward. This allows the implementation of lattice methods in spreadsheet which are commonly used in companies. Besides, the lattice methods provide practitioners with the ability to visualize

the whole underlying asset value evolution and the resulting decisions regarding the exercise or not of the real option. In addition, lattice methods allow an easier valuation of complex real options, such as American options, which are more realistically modeling decision making within a company since the tree allows the comparison between exercising the option prior to maturity and holding-on the real option.

Like the partial differential equation approach, lattice methods require an explicit formulation of the dynamics of the underlying corporate investment process. The main challenge remains the unavailability of well-known stochastic processes to properly model the value of any underlying corporate investment. Even when it is possible, calibrating these models is not trivial due to the absence of historical data to estimate the volatility of the underlying business prospect. Unfortunately, this constrains the variety of applications for real options as many practitioners resort to the use of established models such as the ubiquitous Geometric Brownian Motion without substantiating the assumptions. Nevertheless, some efforts have been observed in recent years to depart from these limitations. For instance, Hahn and Dyer [124] use a modified lattice to value oil and gas switching options with dual-correlated single factor mean reversing stochastic processes. Bastian et al. [125] use a modified lattice to value flexibility and regime switching options that arise during the production of alternative fuels by choosing to favor either sugar or ethanol production. Still, this is a somewhat limited landscape: what about processes with jumps for instance?

Monte Carlo Simulations

Monte Carlo simulations are also popular to price real options and have been widely used to price financial options. Monte Carlo simulations originated in the 1940's

with Ulam and Metropolis [126]. The approach consists in randomly generating many numbers following a given probability distribution to then perform some deterministic computations and to finally aggregate the results. The original argument for using Monte Carlo simulations to price real options is attributed to Boyle [94]. It is based on an observation made earlier in this dissertation stating that under the martingale or change of numéraire approach, an option value can be expressed as an expectation under the new equivalent martingale probability measure. If the option value can be reduced to an expectation, then it lends itself pretty well for Monte Carlo simulations because it only requires the random generation of many prices for the underlying asset using its probability distribution.

Using the strong law of large numbers, it is known that the average of a sample converges almost surely to its expected value. For real options pricing purposes, it means that by generating a sufficiently large number of underlying asset price trajectories and therefore a sufficiently large number of option payoffs, it is possible to recover the expected value of the option payoff at maturity. Recalling the martingale approach presented in section 4.1.2, pricing options using Monte Carlo methods can be decomposed into four main steps. In the first step, the dynamics of the uncertain business prospect revenues is modeled with a stochastic process using both market and historical information. The market uncertainties that have the largest impact on the revenues of the business prospect are first identified and listed. They are then modeled using stochastic processes so that they can be used in the valuation of the business prospect. If these uncertainties are correlated, the correlation must be accounted for so that a proper behavior of the uncertain quantities can be used for the valuation. Some of the most

useful stochastic models used for real options analysis are shown in Figure 15. They include pure diffusion processes and mean reverting processes, both with or without jumps.

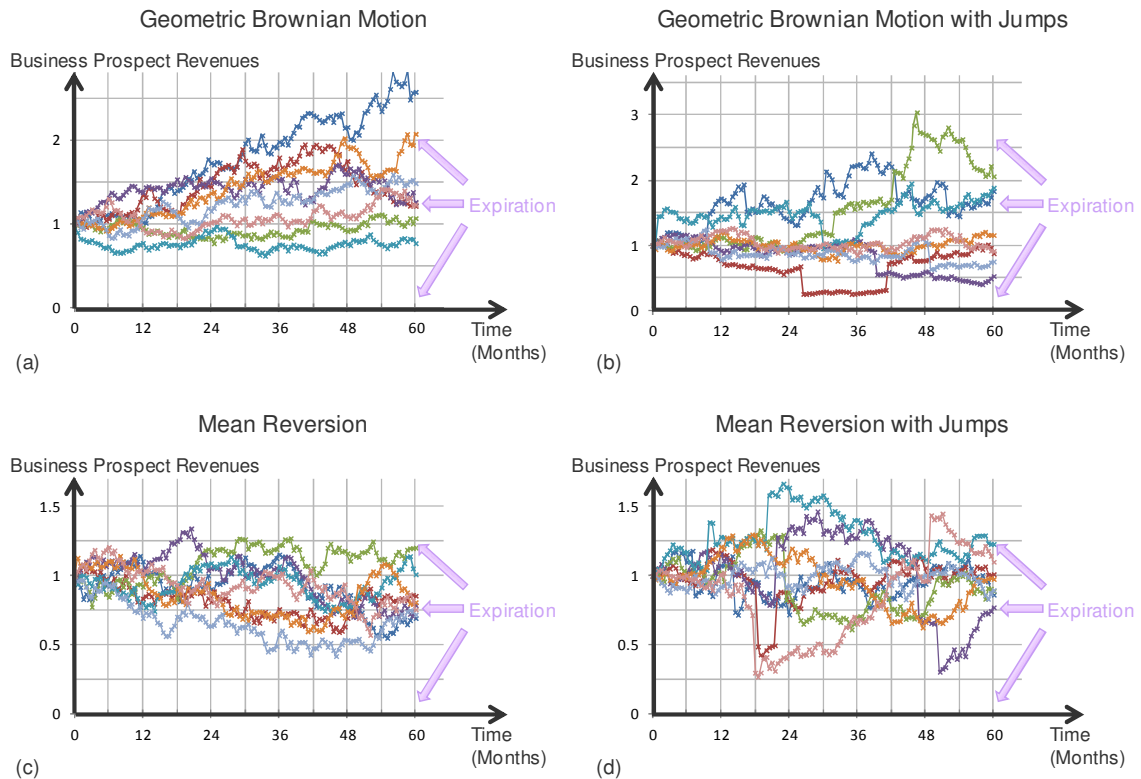


Figure 15: Monte Carlo simulations using some popular stochastic processes

So far, the stochastic processes which model the observations of analysts are defined under the physical or historical or observable probability measure. However, it was previously shown that for options pricing purposes, the underlying asset value process must be defined under the equivalent martingale measure also known as risk-neutral measure: this is made to ensure that the terminal option payoff can be discounted at the risk-free rate. Therefore, the purpose of the second step of the analysis is to define this equivalent martingale measure and to express the dynamics of the business prospect under this synthetic probability measure. For some of the most popular stochastic

processes, the mathematical expression under the risk-neutral probability measure is known and a closed-form expression can be used. Generally speaking, it requires the removal of the risk premium from the drift of the underlying stochastic process.

The numerical implementation is the third step of the analysis. Many simulations are run to generate different trajectories for the value of the business prospect. This step can be implemented in a Monte Carlo simulator as shown in Figure 16 to yield a sampling of the terminal value distribution. In the fourth and final step, the real option payoff is estimated for each and every trajectory generated during the simulations. This enables the estimation of the average payoff which is then discounted to the present time using the risk-free discount rate.

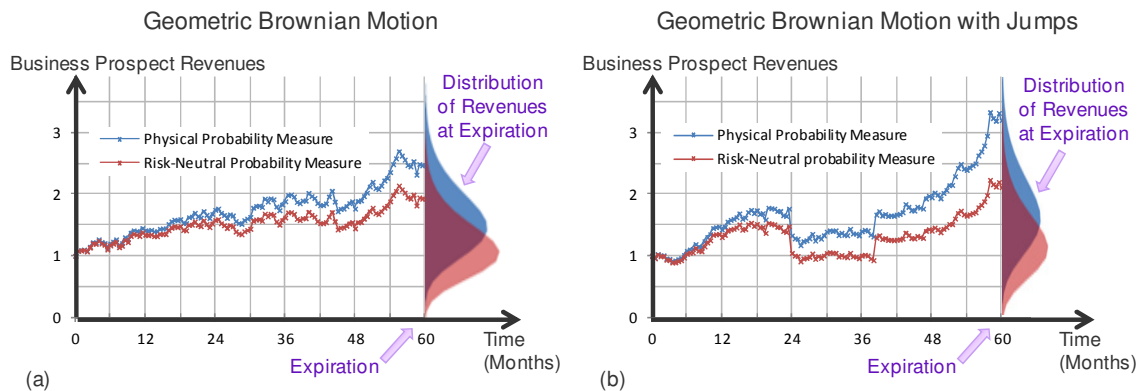


Figure 16: Simulations and resulting business prospect value distributions at expiration under physical and risk-neutral probability measures

Despite the computational flexibility offered by Monte Carlo valuation methods, few academic papers highlight their use and application for real options valuation. This is both surprising and in stark contrast to the financial industry where Monte Carlo methods have been embraced for valuing financial options [127]. Still, Rose [128] uses Monte Carlo simulations to value the option offered to a Government to take back ownership of a toll road concession. The author uses a Geometric Brownian Motion to represent

uncertain traffic volumes on toll roads and a mean reverting process to represent uncertain interest rate evolutions. Tseng and Barz [129] use Monte Carlo simulations of a mean-reverting process to value electricity generating power plants with operating constraints over short-term periods (such as start-up times, minimum up-times, and minimum down-times). More recently, Justin and Mavris [20] use simulations to capture the uncertain deterioration and possible failures of turbofan engine components in order to price maintenance guarantees embedded in some aircraft and engine purchase contracts.

There are many advantages to the use of Monte Carlo simulations for real options valuation. The first advantage is that they allow the simulation of complex processes which would prove almost intractable with more conventional partial differential equations and lattice methods. This is particularly obvious for multi-dimensional real options when the underlying real assets are subject to several sources of uncertainties or when the real options derive their values from several underlying real assets. In these cases, it starts to get impractical to code, draw, and visualize lattices whenever the dimension exceeds two or three. For instance, Rodriguez [130] uses Monte Carlo simulations to value the flexibility offered by the ability to supply different markets with liquefied natural gas. Each of these markets has an uncertain prevailing price for the natural gas, and each is represented by a specific stochastic process.

In addition, these dimensions may not be independent and some correlations may exist between them. A typical example would be price and sales volume: when demand is high, both price and volume increase while when demand is down, both price and volume decrease. Monte Carlo methods present a simple framework to capture these correlations

by generating correlated paths by way of Cholesky decompositions [131]. An example is Yang [132] who captures the correlation between gas and electricity prices while performing simulations to estimate the impact of uncertain future climate-change policies on power investments made by utility companies. Justin and Mavris [37] also use Monte Carlo simulations to model the uncertainties related to the two correlated quantities, jet-fuel price and carbon emission permit prices. In turn, they use simulations of their trajectories to estimate the value of staggered research and development investments in more fuel-efficient technologies for commercial aircraft.

Another advantage of Monte Carlo simulations is the ability to use more complex stochastic models and still implement them with relative ease. More complex models such as those featuring a mean-reverting behavior or those featuring jumps, have proven popular in recent years to improve over some of the deficiencies of pure diffusive processes. Mean-reverting processes have been proposed to model the price of some commodities [133] because the forced return towards a long-term mean is better suited to account for the demand and supply forces that act when prices get away from an equilibrium level. Besides, while analyzing stock returns, Fama [109] realized that many of them were exhibiting leptokurtic distributions with heavier tails than those predicted by pure diffusive processes. He introduced the idea that jumps may be responsible for those heavy tails representing large and sudden shocks. Later, Merton [134] proposed a revised framework for pricing options when jumps are present. Subsequently, Ait-Sahalia [135] proposed a methodology to disentangle diffusion and jumps. This enables the calibration of the diffusive and jump parts of stochastic processes. All in all, there is little doubt that a methodology that can handle these complex processes is superior, for it can

be used in more general settings. As a matter of fact, Monte Carlo inspired methodologies can easily simulate trajectories featuring mean-reverting behaviors and jumps, and can therefore be useful for real options valuation.

Despite these advantages, simulation-based real options valuation has some issues related to the computational complexity. Indeed, even though the implementation is simple, a Monte Carlo approach requires the simulation of a huge number of trajectories (millions) to converge to the expectation value. In turn, this means that the computational cost is quite significant and that runtimes are not short. In addition, despite some recent breakthroughs thanks to Longstaff and Schwartz [136], pricing path-dependent options such as American-type real options or Bermudan-type real options remains challenging. Also, like the partial differential equation approach and the lattice approach, Monte Carlo methods require an explicit formulation of the dynamics of the underlying corporate investment process for two reasons. The first reason is to be able to express the stochastic process under the risk-neutral probability measure since the expectation is made using this equivalent martingale measure. If the process being used is well-known, then the adjustments needed to model it under the new risk-neutral probability measure are usually well-known (this is the case for the Geometric Brownian Motion, the mean-reverting Brownian Motion, as well as the Geometric Brownian Motion with Poisson jumps). The second reason is that the explicit formulation of the process must be known to generate trajectories for the simulations. Consequently, the issues mentioned earlier regarding the estimation of parameters for the stochastic process modeling the underlying prospect value remain. These issues are particularly relevant when little historical data is available to perform a proper calibration.

4.2.2 Summary of methodologies for real options evaluation

Table 20 summarizes the strengths and weaknesses of the different approaches to price options that have been introduced so far. This table includes the (rare) case when closed-form solutions exist (Black-Scholes formula for some European options, Geske formula for some compound options...), the partial differential equation approach, the lattices and trees approach, and finally the Monte Carlo simulations approach. The criteria retained to compare these methods account for their versatility (ability to handle many different problems), their mathematical and economic rigor (whether the model is mathematically and economically sound), their ability to capture the reality of the problem (whether the underlying assumptions are validated), and finally their ability to be implemented easily by practitioners within a company.

Some of these criteria are stemming from the list of requirements set up during a real options symposium held at Georgetown University in 2003 (the *Georgetown Challenge* [106]) where academics and practitioners reached a consensus on what was necessary for real options methodologies to get traction within companies. Reviewing the information contained in Table 14 and being thoughtful of the application of the method, which is the valuation of flexibility for unique investments within the aerospace industry in an environment riddled with uncertainties leads to the following hypothesis:

Hypothesis 1.1.1 – Monte Carlo methods for real options analyses

Monte Carlo methods and lattice-based methods present the most promising approaches to solve for the arbitrage-free value of corporate investments featuring flexibility. Within the context of the aerospace industry, Monte Carlo methods offer the ability to integrate well with other probabilistic methods.

Table 20: Real options valuation methods, strengths and weaknesses

| | Closed-Form Solution | Partial Differential Equation Approach | Lattices and Trees Approach | Monte Carlo Simulations Approach |
|---|----------------------|--|-----------------------------|----------------------------------|
| Intuitively dominates other decision-making methods | + | + | + | + |
| Soundness of the method (mathematically and economically) | + | + | + | + |
| Ease of implementation of the method by practitioners | + | - | + | ~ |
| Ability to visualize uncertainties and the decision process | - | - | + | ~ |
| Ability to capture a complex reality with intertwined uncertainties | - | - | - | + |
| Ability to handle path-dependent real options | - | + | + | ~ |
| Ability to handle corporate investments featuring exotic options | - | ~ | + | + |

This immediately leads to a research question regarding the usability of Monte Carlo methods for their intended use by practitioners within companies.

Research Question 1.1.2 – Improving Monte Carlo methods for real options analyses

Monte Carlo methods seem appropriate to value corporate investments featuring managerial flexibility and programmatic optionality. With usability by practitioners in mind, how can these methods be modified to alleviate the complexity of finding the proper equivalent probability measure required for the expectation computation while maintaining their rigor?

4.3 Probability measure and change of probability measure

In the previous section, numerical methods using Monte Carlo simulations to value real options were introduced. These methods always rely on the computation of a discounted expectation to price real options. The expectation is not computed using the observable historical probability measure but rather the equivalent martingale measure also known as risk-neutral probability measure. What is a probability measure? A probability measure¹ is a function defined on a set of events and returning real numbers in the unit interval, assigning zero for the empty set and one for the entire space. It also satisfies the property of countable additivity which means that for disjoint events, the probability measure of the union of these events is the sum of the probability measures of all the events. Informally, if A is an event within the sample space S and if the function N denotes a measure on that sample space S (for instance the number of occurrences of an event), then the commonly used probability measure is given by $P(A) = N_A/N_S$. Still, nothing in the definition of a probability measure relates to the *observed likelihood* of an event happening as we usually understand it.

In fact, a probability measure is simply defined as a real function with some specific properties. Thus, a probability measure need not be unique and there may exist some other measures besides the usual probability measure (also known as natural, historical, or observable probability measure). If several probability measures exist, one may want to find out if these measures are related, and if so, how they are related. This

¹ A probability measure \mathbb{P} on a sample space S is a real-valued function defined on the collection of events of a random experiment that satisfies the three properties: $P(A)$ is non-negative for any event A in S , $P(S)=1$, and if A_i is a countable collection of pair-wise disjoint events, then $P(\cup_{i \in I} A_i) = \sum_{i \in I} P(A_i)$

leads to the introduction of equivalent probability measures¹. Probability measures are said to be equivalent if they agree on what is possible and what is impossible. Indeed, according to Schreve [137], the equivalent probability measures “*must agree on what is possible and what is impossible; they may disagree on how probable the possibilities are.*” For real options applications, two equivalent probability measures agree on which values for the business prospect are possible and which values are not. They differ however on their assessments of the likelihood of each possible value happening.

Therefore, the remaining question is how to go from one probability measure \mathbb{P} to another probability measure \mathbb{Q} . Better said, how to model the evolution of a stochastic process in a new synthetic probability measure \mathbb{Q} once it has been calibrated under the historical probability measure \mathbb{P} using historical data? This is done using the Radon-Nikodym derivative for equivalent probability measures also known as the likelihood ratio. The likelihood ratio terminology is easier to understand since it represents the ratio of the likelihood of an event in one probability measure \mathbb{P} over the likelihood of the same event in the other probability measure \mathbb{Q} . It is often defined as expressed in Eq. 9, where ξ is the density of the probability measure \mathbb{Q} with respect to the measure \mathbb{P} .

$$\xi = \frac{d\mathbb{Q}}{d\mathbb{P}} \text{ with the property that } E_{\mathbb{P}}(\xi) = 1 \quad \text{Eq. 9}$$

Real options are however dealing with continuous-time processes instead of finite sequence of random variables. Indeed, it is the continuous evolution of the underlying business prospect value that drives the value of the real option. When dealing with these processes, a further extension of the Radon-Nikodym derivative is necessary. The change

¹ Two probability measures \mathbb{P} and \mathbb{Q} on a sample space S are equivalent if for all events A in S the following is true: $P(A) = 0 \Leftrightarrow Q(A) = 0$

in the dynamics of stochastic processes when probability measures are distorted is described by the Cameron-Martin-Girsanov theorem. For a brief introduction on the mathematics of probability measure, equivalent probability measure, and change of measure in continuous-time, the reader is referred to APPENDIX C: while more thorough texts can be found in Shreve [137], Shreve [138], and Neftci [139].

4.3.1 Risk-neutral measure for some common stochastic processes

In the previous section, the mechanics underpinning the change of probability measure required for real options valuation was introduced. Indeed, in order to use simulations to value a business prospect, the dynamics of the underlying investments must be specified in the risk-neutral measure. Simply said, the risk-neutral measure is a probability measure for which the returns of all assets are exactly the risk-free rate of return. Mathematically, this is equivalent to subtracting the risk-premium from the expected returns which makes investors indifferent towards risk, hence the name of the measure. In this section, the techniques introduced above will be used to derive the expression for the dynamics of some common stochastic processes which can then be used directly in Monte Carlo simulators. The derivations are again presented in APPENDIX C: and only the results are presented in Table 21.

The Cameron-Martin-Girsanov theorem is helpful to perform changes of measures and to express the continuous evolution of stochastic processes under different measures. However, one of the main advantages of Monte Carlo simulations for real options evaluation is their ability to cope with discontinuous processes. In this regards, how can the change of measure be performed when the underlying process is discontinuous and features jumps?

Table 21: Changing the measure of popular stochastic processes for real options

| Stochastic Process | Stochastic Differential Equation using Historical Probability Measure | Comments |
|---|---|--|
| | Stochastic Differential Equation using Risk-Neutral Probability Measure | |
| Arithmetic Brownian Motion | $dS_t = \mu dt + \sigma dW_t$ | Radon-Nikodym derivative given by: $\xi_t = e^{\frac{(\mu-r_f)^2}{2\sigma^2}t + \frac{\mu-r_f}{\sigma}W_t}$ |
| | $dS_t = r_f dt + \sigma dW_t^{\mathbb{Q}}$ | |
| Geometric Brownian Motion | $dS_t = \mu S_t dt + \sigma S_t dW_t$ | Radon-Nikodym derivative given by: $\xi_t = e^{\frac{(\mu-r_f)^2}{2\sigma^2}t + \frac{\mu-r_f}{\sigma}W_t}$ |
| | $dS_t = r_f S_t dt + \sigma S_t dW_t^{\mathbb{Q}}$ | |
| Jump Diffusion Process (Merton) | $dS_t = \mu S_t dt + \sigma S_t dW_t + (J_t - 1)S_t dN_t$ | Jump model with Poisson counting process of intensity λ for occurrence of jumps which are of $J-1$ distributed amplitude |
| | $dS_t = (r - \lambda E(J_t - 1))dt + \sigma S_t dW_t^{\mathbb{Q}} + (J_t - 1)S_t dN_t$ | |
| Mean-Reverting Process (Ornstein-Uhlenbeck) | $dS_t = \eta(\bar{S} - S_t)dt + \sigma dW_t$ | Mean-reverting model with adjustment speed η and long-term average \bar{S} . Risk-neutral model uses the market price of risk λ |
| | $dS_t = \eta\left(\bar{S} - \frac{\sigma\lambda}{\eta} - S_t\right)dt + \sigma dW_t^{\mathbb{Q}}$ | |

4.3.2 Esscher transform for option pricing

An extension of the change of measure technique was proposed in 1994 by Gerber and Shiu [140] to handle a wider variety of processes featuring stationary and independent increments such as Wiener processes, Poisson processes, Gamma processes, and inverse Gaussian processes. Similarly to the previous technique, a transformation is used to induce an equivalent probability measure. The transformation is based on the *Esscher transform* [141], a time-honored tool in actuarial finance pioneered by Swedish

mathematician Fredrik Esscher and later publicized by Kahn [142]. For a probability density function f and a real number h , the Esscher transform f_{Ess} with parameter h is expressed using the moment generating function of f as shown in Eq. 10:

$$f_{Ess}(x, h) = \frac{e^{hx} f(x)}{M(h)}, \text{ with } h \in \mathbb{R} \text{ and } M(h) = \int_{-\infty}^{\infty} e^{hx} f(x) dx \quad \text{Eq. 10}$$

Looking at this definition, the Esscher transform is the product of an exponential function and a density function, normalized by a moment generating function. As a result, this transformation induces an equivalent probability measure as both distributions agree on sets with probability zero. It also becomes clear why the Esscher transform is sometimes called exponential tilting: the transformation distorts the original probability measure using an exponential function. The goal of this technique is to use the free parameter h introduced by the Esscher transform to ensure that the new probability measure is an equivalent martingale measure. Consequently, the parameter h is determined to ensure that the discounted underlying asset price is a martingale or, better said, that the price of the underlying asset is exactly its expected discounted payout.

A sketch of the derivation of the Esscher transform for option pricing is found in APPENDIX D: but the main steps are presented in the following paragraphs. In the first step, the Esscher transform that was defined for a single random variable is modified for the purpose of option valuation by adapting it to stochastic processes and indexing it with the time parameter t . This leads to a new definition of the Esscher transform for stochastic processes given in Eq. 11. This new Esscher transform leads to a new stochastic process. The moment generating function associated with this new probability distribution is given in Eq. 12:

$$f_{Ess}(x, t, h) = \frac{e^{hx} f(x, t)}{M(h, t)}, \text{ with } h \in \mathbb{R} \text{ and } M(h, t) = \int_{-\infty}^{\infty} e^{hx} f(x, t) dx \quad \text{Eq. 11}$$

$$M(z, t; h) = \int_{-\infty}^{\infty} e^{zx} f_{Ess}(x, t, h) dx = \frac{M(z + h, t)}{M(h, t)} \quad \text{Eq. 12}$$

The next step consists in solving for the parameter h that makes the transformed distribution a risk-neutral measure. This parameter, noted h^* , is unique and solves the expectation shown in Eq. 13:

$$S_0 = E_{Ess}^h(e^{-r_f t} S(t)) \quad \text{Eq. 13}$$

Assuming now that the process describing the behavior of the underlying asset can be written as $S(t) = S_0 \cdot e^{X(t)}$ with $X(t)$ a stochastic process with stationary and independent increments starting at zero, then a simplification of the previous equation leads to $1 = E_{Ess}^h(e^{X(t)}) = M(1, t; h^*)$ which, in turn, yields Eq. 14. This defines the risk-neutral Esscher transform of parameter h^* and the corresponding measure is the risk-neutral Esscher measure.

$$r_f = \ln(M(1, 1; h^*)) \quad \text{Eq. 14}$$

To summarize, the original stochastic process has been distorted to yield a new stochastic process. To compute the value of the real option, it now suffices to first simulate the dynamics of the underlying asset under this new risk-neutral Esscher measure, then to estimate the expectation of the payoffs under this new measure, and finally to discount back these payoffs to the present time using the risk-free discount factor.

The Esscher transform is a powerful technique that has been applied to various pricing problems in finance. It presents the advantage of being both a rather straightforward and versatile technique being able to handle many different types of

processes, including some of the most commonly used stochastic processes in finance such as diffusion processes and diffusion processes with jumps. Gerber and Shiu demonstrate that the classical results of Black-Scholes [66] and Merton [67] for a Wiener process, Cox and Ross [92] for a shifted Poisson process and finally Cox, Ross, and Rubinstein [99] for the binomial model can be reproduced with this approach.

When markets are complete, the equivalent martingale measure is unique and therefore the risk-neutral Esscher transform gives the unique arbitrage-free price for the real option. The Marketed Asset Disclaimer (MAD) assumption presented earlier in this thesis ensures that the market is complete and therefore that a unique price for the real option can be found. On the other hand, when the market is incomplete, the claim is not attainable and there is no possibility for the market and its arbitrageurs to *enforce* a no-arbitrage price. Mathematically, there may be many equivalent martingale measures and the practitioner has to select one of them. Several equivalent measures [143] have been proposed such as the minimal martingale measure [144], the minimal entropy martingale measure [145], the utility martingale measure [145], and of course, the Esscher martingale measure. Each of them corresponds to a different attitude towards risk. As a result, some assumptions regarding the preferences and risk attitude of decision-makers must be set to pick which utility function and therefore which equivalent martingale measure is most appropriate.

In fact, in the discussion pertaining to their paper [140], Gerber and Shiu show that the Esscher martingale measure is consistent with investors or decision-makers

exhibiting power utility behaviors¹. Power utility functions, also known as isoelastic utility functions, have the property of Constant Relative Risk Aversion (CRRA) which means that the risk aversion is independent of the level of initial wealth. The reader is referred to APPENDIX D: where the relationship between the measure of risk aversion η and the Esscher parameter h^* is derived. The power utility assumption also has the advantage of being consistent with some other fundamental results of finance and economics (such as the mutual fund theorem in Cass and Stiglitz [146] and Stiglitz [147] for instance).

Over the years, several improvements pioneered by Gerber and Shiu have been made to the change of measure by means of Esscher transform. Buhlmann et al. [148] use conditional Esscher transforms to construct equivalent martingale measures for classes of semi-martingales. Inspired by the work of Duan [149] on locally risk-neutral valuation relationships, Siu, Tong, and Yang [150] also use the conditional Esscher transform to price options with an underlying following the popular Generalized AutoRegressive Conditional Heteroscedasticity (GARCH) model [112]. Following-up on their acclaimed work, Gerber and Shiu [151] use the Esscher transform and the optional sampling theorem [152] to derive the price of perpetual American options. More recently, Goovaerts and Leaven [153] use an axiomatic characterization to define a pricing mechanism that can generate approximate arbitrage-free derivative prices and use a

¹ A power utility function belongs to the class of hyperbolic absolute risk aversion utility functions. It is a special case in that it exhibits a constant relative risk aversion. The power utility function relates the utility U to the level of consumption c using the following formula with η a constant measuring risk-aversion:

$$U(c) = \begin{cases} \frac{c^{1-\eta}-1}{1-\eta} & \eta > 0, \eta \neq 1 \\ \ln(c) & \eta = 1 \end{cases}$$

probability measure transformation closely related to the Esscher transform called the Esscher-Girsanov transform.

Despite their many advantages, Esscher-based valuation techniques still require the explicit formulation of a model describing the dynamics of the underlying asset value which may be an issue for real options valuation purposes. Indeed, unlike financial options for which the underlying asset price is readily available and for which historical price data is available, real options models cannot usually rely on large datasets to both assume a particular type of behavior and then calibrate the parameters of the assumed stochastic process. This may lead to two different kinds of errors for the real options practitioner: model misspecification if the functional form of the model is wrong and model calibration error if the estimation of parameters is skewed.

Surprisingly, the Esscher transformation has never been used for real options analysis to the author's knowledge. This may be due to the lack of exposure of practitioners to the technique or to the modeling issue mentioned above. Still, this technique seems promising enough to warrant further research to adapt it for corporate investment analyses.

Research Question 1.1.3 – Adaptation of Esscher transform for pricing real options

How can option pricing by means of Esscher transform be adapted to a corporate investment analysis within the context of a real options methodology?

4.3.3 Non-parametric Esscher transform and real options pricing

In the previous paragraph, the Esscher transform was introduced as a powerful and efficient tool to price options in both a complete (no arbitrage pricing leads to a single equivalent martingale measure) and incomplete markets (the probability measure

induced by the Esscher transformation is consistent with a representative agent featuring a power utility function) when the underlying asset has some nice properties (stationary and independent increments). This is a promising framework for real options analysis as most real options cannot be hedged by a replicating portfolio traded on the markets.

Adapting the Esscher transformation technique so that it does not require the explicit formulation of the underlying stochastic process would prove particularly useful for real options analysis. Indeed, for many real options analyses, the main source of uncertainty is not the cash flow itself but rather the multitude of uncertainties affecting the cash flow. In the context of technology developments for aerospace applications, the source of market uncertainty driving the value of technologies is usually the price of commodities such as the price of jet fuel or the price of carbon permits. Of course, this uncertainty propagates downstream and affects the operating costs and therefore the value and attractiveness of fuel-saving technology developments. While some classes of stochastic processes are suitable to model these uncertainties and ample historical data is available to calibrate these models, assuming and using stochastic processes to directly model the cash flows is more difficult to substantiate because the lack of historical data prevents a proper estimation of the model parameters. Of course, it is possible to back-engineer a cash flow model by first modeling the driving uncertainties to generate estimates of cash flows which can then be used to assume a cash flow model. This circuitous approach adds however another step and another layer of assumptions that may neither be necessary nor desirable.

Fortunately, a paper presented at a conference in Southern France by Pereira, Epprecht, and Veiga [154] proposed a non-parametric method that inspired this research

endeavor. The method is an extension of the Esscher transform presented earlier and uses a non-parametric, model-free, Esscher transform to simulate the behavior of the underlying asset from the physical probability measure to the risk-neutral probability measure. The technique is geared towards the pricing of financial options and therefore will need to be adapted for the economic evaluation of corporate investments featuring flexibility. Detailed derivations of this non-parametric technique may be found in the last section of APPENDIX D: devoted to the Esscher transform. Still, the main steps are presented in the following paragraphs for the sake of completeness.

The first step of the analysis starts with the collection of the data (and eventually its reduction if required) regarding the price of the underlying asset S_t . This data may have either one of two origins: it can be directly observable and available (such as the market price of the underlying asset) or it can be generated by the practitioner if the underlying asset is synthetic or not publicly traded. These prices are used to estimate the continuously compounded rate of return x_t at time t . In the first case, there is only one rate of return at each time-increment: indeed, there is a single price observation since there is a single realization of the uncertainty during that time increment. A bootstrapping technique is therefore used to generate many possible realizations at each step. This enables an approximation of the unknown distribution of rate of returns at each time t . In the second case, the prices are generated by the practitioner using one or more stochastic processes. A Monte Carlo simulation is therefore sufficient to generate a distribution of returns for each time step. As shown in Eq. 15, let's now call \widehat{X}_t the vector of size n containing these n rates of return sampled from the unknown probability distribution at time t .

$$\widehat{X}_t = [x_t^1, x_t^2, x_t^3 \dots x_t^n] = \left[\ln\left(\frac{S_t^1}{S_{t-1}^1}\right), \ln\left(\frac{S_t^2}{S_{t-1}^2}\right), \ln\left(\frac{S_t^3}{S_{t-1}^3}\right) \dots \ln\left(\frac{S_t^n}{S_{t-1}^n}\right) \right] \quad \text{Eq. 15}$$

The second step of the analysis consists in the computation of the moment generating function which is estimated using its empirical counterpart, the empirical moment generating function given in Eq. 16:

$$\widehat{M}_t(h, t) = \frac{1}{n} \sum_{i=1}^n e^{hx_t^i} \quad \text{Eq. 16}$$

The third step of the analysis is directly inspired by the work of Gerber and Shiu in that it solves for the specific value of the parameter h such that the asset price is a martingale under the new, to be constructed, probability measure induced by the Esscher transform. The parameter h^* must solve Eq. 17 and, in a complete market with no arbitrage, the fundamental theorem of asset pricing [102] ensures that this solution is unique.

$$e^{r_f} = \frac{\sum_{i=1}^n e^{(h^*+1)x_t^i}}{\sum_{i=1}^n e^{h^*x_t^i}} \quad \text{Eq. 17}$$

With the proper value h^* of the Esscher transform parameter, the final step consists in constructing the new probability measure. This is done by reweighting each of the observation and ensuring that these probabilities sum to one. The risk-neutral probability vector providing the probability of each observation is given by Eq. 18. This is the set of probabilities that is used for the pricing of options and for the computation of expectations.

$$\mathbb{Q}_t^{h^*} = \left[\frac{e^{h^*x_t^1}}{\sum_{i=1}^n e^{h^*x_t^i}}, \frac{e^{h^*x_t^2}}{\sum_{i=1}^n e^{h^*x_t^i}} \dots \frac{e^{h^*x_t^n}}{\sum_{i=1}^n e^{h^*x_t^i}} \right] \quad \text{Eq. 18}$$

In summary, the non-parametric Esscher transform enables practitioners to distort an unknown probability measure into a risk-neutral probability measure. This transformation is done on a sample of representative observations or a sample of simulated observations for option pricing purposes and leads to a new sample. Next, this new sample is used for the estimation of the option payoffs which are then discounted back to the present time using the risk-free interest rate to estimate the option value. The algorithm of the non-parametric Esscher transform is depicted in Figure 17.




Figure 17: Non-parametric Esscher transform for change of probability measure

All in all, this technique tremendously simplifies the analyses of practitioners who no longer need to estimate, calibrate, and substantiate the choice of one particular stochastic process for the underlying asset, provided some mild conditions of stationary and independent increments are satisfied. This leads to the following research hypothesis regarding the change of probability measure.

Hypothesis 1.1.2 – Non parametric Esscher transform for pricing real options

Real options valuation via non-parametric Esscher transforms is a promising framework for staggered investment analyses. It is based on rigorous foundations, offers a clear and transparent methodology for practitioners, and uses probabilistic techniques widely accepted within companies.

4.4 Path-dependent options

So far, most of the discussion has revolved around enabling methods for option analysis. However, little has been said about the type of options that can be useful for real options analyses. The most widely studied options are European options which give the option holder the right but not the obligation to undertake an investment at one pre-specified point in time. These options are so common that they are referred to as “plain vanilla options” and are usually simpler to analyze. Let’s pause momentarily and remember that one goal of real options analyses is to leverage the upside potential created by the identification of precursors of successful program developments. European types of options with set exercise dates may not be the most appropriate type to use. In fact, two other types of options may be more useful for corporate investment applications: American options and Bermudan options.

4.4.1 Managerial flexibility and trigger events

Indeed, managerial flexibility represents the opportunities offered to management to react in real-time to the unfolding of an uncertain future. In this context, decision-makers can exercise their managing privileges to alter substantially the course of development programs. In particular, following the detection of trigger events

announcing an unfavorable future, managers may decide to shrink, delay, or abandon investments. Following the detection of trigger events announcing a favorable future, managers may decide to expand or rush an investment. How can real options analyses be framed to handle these real-time decisions? This leads to the following research question.

Research Question 1.2 – Managerial flexibility and timing of investments

How can the flexibility offered to management to optimally time the launch of new investments be accounted for in a real options framework?

4.4.2 American and Bermudan real options

An American option gives the holder the right but not the obligation to undertake an investment at any time prior to a pre-specified deadline. This is strikingly in line with the decision-makers ability to undertake an investment whenever they feel the market is ready and the conditions are optimal. A Bermudan option is similar to an American option but exercising the option can only be done at several pre-specified dates up to the expiration of the option. In the context of pricing options via simulations, the time-discretization introduced for the generation of trajectories basically transforms any continuous-time American option into a Bermudan option with exercise possibilities at each time-step. In the following sections, the algorithms presented for the pricing of American options are in fact pricing Bermudan options with as many possible exercise dates as there are time-steps in the simulation. As the number of time-steps in the simulation grows, the Bermudan option tends to be more and more similar to an American option and its price converges to the price of its American counterpart.

The striking similarity between American and Bermudan types of derivative contracts and the flexibility offered to management and decision-makers to invest whenever conditions become optimal lead to the following assertion: practitioners could leverage some of the techniques developed for the evaluation of path-dependent options to analyze corporate investments featuring flexibility.

Hypothesis 1.1 – Path-dependent options to model managerial flexibility

As uncertainty unfolds, programmatic, technological, and market opportunities emerge and disappear. Flexible management and flexible timing of investment decisions allow the maximization of the upside potential of these opportunities. Path-dependant options such as American options present a means to model the flexibility offered to management in timing these investment decisions.

4.4.3 Early-exercise boundary

American options and their Bermudan approximations are special in that these contracts can be exercised at almost any time prior to the expiration of the options. Quoting Glasserman [131], “*the value of an American option is the value achieved by exercising optimally.*” In fact, if this was not the case, arbitrageurs would actually kick-in and enforce a price that is in agreement with an optimally enforced option. Valuing this type of option is therefore equivalent to finding the optimal exercise rule and then computing the expected discounted payoff using this rule to decide whether the option is exercised early or not. Defining the optimal exercise rule is however not a trivial affair. The optimal exercise rule is a function of several parameters and can be interpreted as a multi-dimensional surface. Heuristically, it has to be a function of the current asset price and the remaining time before expiration of the option. On the one hand, if the current

price of the underlying asset takes extreme values, it might become profitable to exercise early in-the-money options so as to pocket the payoff with certainty. On the other hand, if a significant amount of time remains before expiration, it might not be worth exercising early an option that is barely in-the-money as better opportunities might arise later. Two extra parameters enter into the equation for defining the early-exercise boundary. The first one is the risk-free interest rate which indicates how the option payoff will earn interest after early-exercise. The second one is the underlying asset volatility which indicates how likely the underlying asset is to move significantly in the future.

Up to this point, the early-exercise boundary is defined as a function of at least four variables: current price, time to expiration, risk-free interest rate, and underlying volatility. A notional early-exercise boundary is given in exhibit (a) of Figure 18 for an American put option. As the American option gets more complex, some new parameters may nonetheless affect the shape of the early-exercise boundary. Let's assume for instance that the underlying asset is paying a dividend. On the ex-dividend date, the holder of the stock will reap a positive benefit while the holder of the option will experience a downward pressure on the underlying price. This must affect the early-exercise boundary as an astute investor will ensure not to exercise an American put option right before the ex-dividend date in order to collect the extra benefit. This is shown in exhibit (b) of Figure 18 which displays the early-exercise boundary of an American put option on a stock with one dividend payment: the early-exercise boundary disappear suddenly before the ex-dividend date.

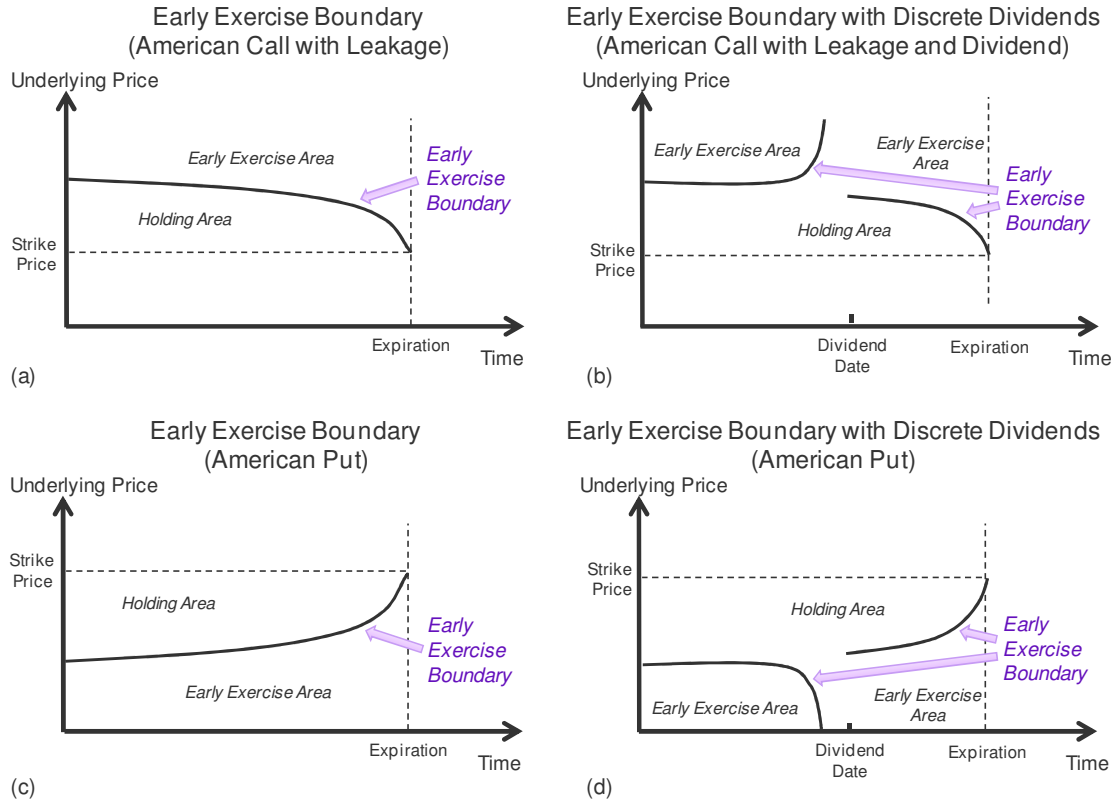


Figure 18: Early-exercise boundaries for American put options

For a real options application, modeling dividend may seem useless at first sight: after all, a real development program usually does not pay any dividend to the company. This is obviously true, but dividend-like payments may be useful to model some other aspects that are very relevant in corporate finance. This is for instance the case for modeling the cost of delay or the entrance of a new competitor in the market which both reduce the expected value of the development program.

Now, the main question remains: how can one approximate the early-exercise boundary which corresponds to the optimal investment boundary for real options and corporate investments? This leads to the following research question:

Research Question 1.2.1 – Early-investment boundaries to detect trigger events

How can early-investment boundaries be identified and how can they be used for the identification of precursors of successful development programs?

4.4.4 Monte Carlo simulations for pricing American options

Previously, Monte Carlo simulations were introduced as a clear and transparent technique for pricing options. For real options applications, Monte Carlo simulations are even more useful as they enable the capture of a multitude of uncertainties and their interdependencies. However, pricing real options using Monte Carlo simulations has long been hindered by its perceived inability to correctly handle path-dependant options. Indeed, in 1993, Hull writes in the second edition of his book [155] that “*Monte Carlo simulation can only be used for European-style options.*”

The main reason for this difficulty is that simulations will yield an estimate of the option value at a single point defined by the current time and the current asset price. The Monte Carlo simulations technique does not yield information regarding the option value at future times and for different asset prices which is problematic. How then to ensure that the early-exercise policy is not violated? In other words, when moving along a simulated trajectory, one needs to ensure that the optimal early-exercise policy is followed. This means that, while marching forward in time, one has to compare the value of holding the option for at least one extra step to the payoff earned by an immediate exercise. Mathematically, the value of the American option at the k^{th} time-step t_k denoted V_{t_k} on an asset S with observed value S_{t_k} and with payoff function P can be expressed as the maximum between exercising immediately and holding the option as shown in Eq. 19. The issue is that there is yet no estimate of the present value of the one-period-ahead

option value $V_{t_{k+1}}$. One could imagine computing this value by performing a nested simulation (simulation within the simulation) but these nested simulations lead to an exponential number of computations which is not a feasible solution.

$$V_{t_k} = \max[P(S_{t_k}), e^{-r_f(t_{k+1}-t_k)} E_{\mathbb{Q}}(V_{t_{k+1}} | S_{t_k})] \quad \text{Eq. 19}$$

Fortunately, this paradigm has evolved starting in 1993 with the paper of Tilley [156] which aim was to dispel the belief that American-style options could not be valued using simulations. Tilley proposed an algorithm where the holding value is determined by first partitioning the price distribution at each time-step and bundling together prices (and associated trajectories) that are similar. In the following step, the bundles are used to estimate the expected one-period-ahead option value by using the trajectories belonging to the asset price realizations contained in the bundles. In other words, a bundling of similar prices is performed at each time-step to compute a conditional expectation by pretending that the corresponding trajectories constitute a new sampling of the underlying asset price trajectories. Despite providing reasonable results, this algorithm lacks justification regarding why it works and how the bundling shall be performed (what choice of criteria for identifying *similar* prices).

4.4.5 Pricing American options using Longstaff-Schwartz algorithm

A significant improvement came in 1996 with the work of Carriere [157] regarding the valuation of options with early-exercise properties. Faced with the same problem of estimating the one-period-ahead option value for subsequent comparison with the immediate exercise payoff, Carriere suggests the use of non-parametric regressions to regress the conditional expectation and therefore to estimate the value of holding the

option. As noted by Stentoft [158], the reason for this regression is that a conditional expectation is a function and “any function belonging to a separable Hilbert space may be represented as a countable linear combination of basis-functions for the space.” Consequently, let’s introduce $\{\phi_i\}_1^\infty$ as a family of basis-functions for that space. The expectation may be rewritten and approximated using the first M basis-functions $\{\phi_i\}_{i=1}^M$ as shown in Eq. 20:

$$E_{\mathbb{Q}}(V_{t_{k+1}}|S_{t_k}) = \sum_{i=1}^{\infty} \alpha_i(t_k) \cdot \phi_i(S_{t_k}) \sim \sum_{i=1}^M \alpha_i(t_k) \cdot \phi_i(S_{t_k}) \quad \text{Eq. 20}$$

Any family of basis-functions should work, but Carriere suggests using either splines or a polynomial smoother. The next task consists in estimating the coefficients α_i of the linear combination. This is done by marching backward, starting at expiration and moving back time-step by time-step until the present time. At expiration, the value of the option is exactly the payoff. For all preceding time-steps, denoted t_k , a regression is performed using the observations of the asset price for the N simulated trajectories at that time t_k , denoted by S_{t_k} , as well as the option value $V_{t_{k+1}}$ at the following time-step t_{k+1} . The regression objective is to select a family of coefficients $\{\alpha_i\}_1^M$ that minimizes the error between the regressed conditional expectations and the option value across all simulated trajectories. This error is defined in Eq. 21:

$$\min_{\{\alpha_i\}_1^M} \sum_{n=1}^N \left(\sum_{i=1}^M \alpha_i(t_k) \cdot \phi_i(S_{t_k}^n) - V_{t_{k+1}}^n \right)^2 \quad \text{Eq. 21}$$

The immediate exercise value at time t_k , denoted $P(S_{t_k})$, is then compared to the discounted regressed conditional expectation to find the option value, defined by Eq. 22.

The procedure is repeated for each time step and for each trajectory until the present time to find the value of the American option.

$$V_{t_k} = \max \left[P(S_{t_k}), e^{-r_f(t_{k+1}-t_k)} \sum_{i=1}^M \alpha_i(t_k) \cdot \phi_i(S_{t_k}) \right] \quad \text{Eq. 22}$$

The algorithm for American option valuation through regression is depicted in Figure 19. Another popular enhancement to this work is the Least-Squares Monte Carlo approach of Longstaff and Schwartz. Dating back to 2001, this approach is very similar to the method of Carriere except for two facts: the algorithm uses a least-squares regression and the regression is made using only in-the-money paths.

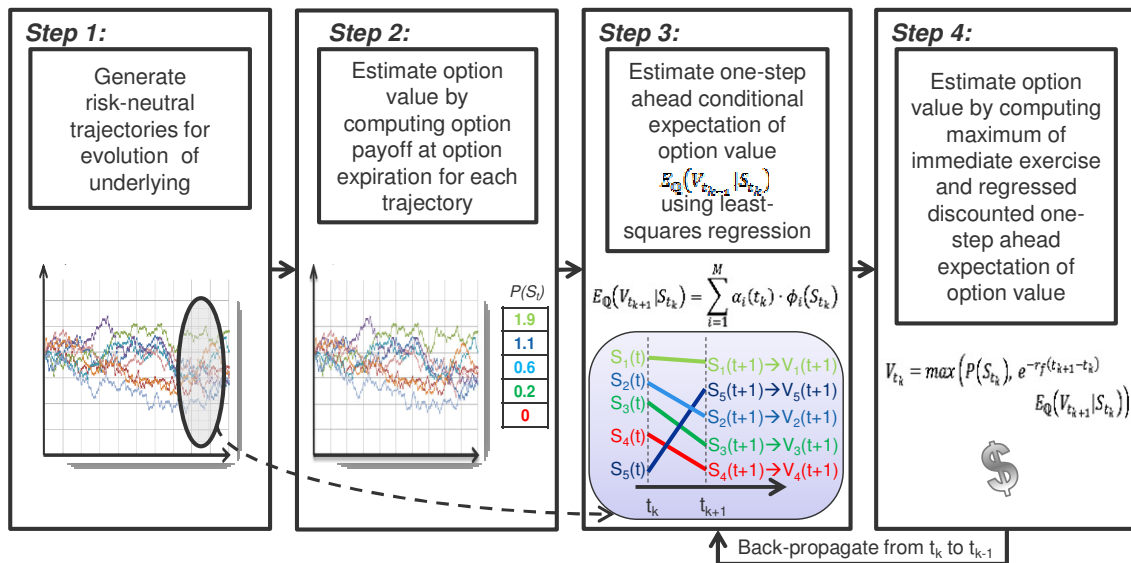


Figure 19: American option valuation with regression

Indeed, the proposed regression uses an ordinary least-squares technique to regress the conditional expectation $E_{\mathbb{Q}}(V_{t_{k+1}} | S_{t_k})$ against a set of explanatory variables. The set of explanatory variables is a family of basis-function $\{\phi_i\}_1^M$ valued using the conditioning underlying asset price S_{t_k} . The families of basis-functions proposed by

Longstaff and Schwartz include the simplest possible monomial family $\{\phi_i: X \rightarrow X^{i-1}\}_{i=1}^M$ as well as the family of weighted Laguerre polynomials defined as: $\{\phi_i: X \rightarrow e^{-X/2} \cdot \frac{e^X}{i!} \cdot \frac{d^i}{dX^i} (X^i e^{-X})\}_{i=1}^M$ Other types of basis-functions that could be used include the Hermite, Legendre, and Chebyshev polynomials. As in Carriere's method, only the first M basis-functions are used to perform the conditional expectation calculation.

Besides, the regression is reduced by using only paths that are in-the-money since the decision to exercise or not the option is only relevant whenever the option is in-the-money. According to Longstaff and Schwartz, *"by focusing on the in-the-money paths, [... they...] limit the region over which the conditional expectation must be estimated, and far fewer basis functions are needed to obtain an accurate approximation to the conditional expectation function."* All in all, this significantly improves the efficiency of the algorithm as the dimensionality of the regressions is reduced and the regressions must be estimated many times over the course of the simulations.

Another subtle difference with the works of Carriere is the choice of realized payoffs as dependent variables for the regression instead of using previously computed conditional expectations. These realized payoffs may be resulting from an early-exercise at the subsequent time-step t_{k+1} or from an early-exercise several steps down the trajectory, for instance at t_{k+j} with j greater than one. According to Longstaff and Schwartz, this precludes *"an upward bias in the value of the option"*. In other words, the conditional expectation at time t_k denoted by $E_{\mathbb{Q}}(V_{t_{k+1}} | S_{t_k})$ is used only once in the entire algorithm: to check whether the value of holding the option is greater than the value of immediate exercise. For all other purposes, such as the estimation of the option

value at time t_k denoted as V_{t_k} , or the regression of the conditional expectation at time t_{k-1} denoted by $E_{\mathbb{Q}}(V_{t_k}|S_{t_{k-1}})$, the conditional expectation at time t_k is not used. This leads to the following exercise rule and option value in Eq. 23. Let's notice the subtle difference with Eq. 22 in what the option value really is (the exercise rule remains the same).

$$V_{t_k} = \begin{cases} P(S_{t_k}) & , \text{if } P(S_{t_k}) \geq e^{-r_f(t_{k+1}-t_k)} \sum_{i=1}^M \alpha_i(t_k) \cdot \phi_i(S_{t_k}) \\ e^{-r_f(t_{k+1}-t_k)} \cdot V_{t_{k+1}} & , \text{if } P(S_{t_k}) < e^{-r_f(t_{k+1}-t_k)} \sum_{i=1}^M \alpha_i(t_k) \cdot \phi_i(S_{t_k}) \end{cases} \quad \text{Eq. 23}$$

With the least-squares Monte Carlo approach used to approximate the early-investment boundary, one question remains: how can these regressions be made when observations, both explanatory and dependent variables, are weighted? This leads to the following research question:

Research Question 1.2.2 – Approximation of early-investment boundaries

How can algorithms approximating the early-investment boundaries be used in conjunction with the non-parametric Esscher risk-neutralization?

4.4.6 Bootstrapping for American and Bermudan options

This section might be slightly misleading as the bootstrap technique is not used *per se* to value options. However, it was deliberately put at this place because it details an enabling technique to value American and Bermudan options. In this section, three possible alternatives are reviewed and some emphasis is given on a re-sampling technique also known as bootstrapping.

The non-parametric Esscher transform was previously presented as an efficient method to perform a change of measure and find the equivalent risk-neutral measure. By doing so, the technique is changing the mean of the terminal distribution of the underlying asset return by reweighting the different outcomes simulated. This procedure is however acting only on the terminal distribution of the underlying asset price. What about the price distributions for intermediate steps? As much as the procedure is sufficient for valuing European types of options whose values depend only on the distribution of the underlying asset prices at expiration, valuing an American or a Bermudan option requires the knowledge of the underlying asset price distribution at each and every intermediate steps under the risk-neutral measure \mathbb{Q} in order to perform the conditional expectation regressions of the Longstaff-Schwartz algorithm.

Having identified this issue, there are at least three natural ways to proceed forward:

- 1) perform the risk-neutralization for the underlying asset distribution at each time-step independently, resulting in a time-indexed vector of Esscher transform parameters;
- 2) perform the risk neutralization across all time-steps concurrently and optimize to find a single Esscher transform parameter that would reasonably work for each time-step; or
- 3) use a re-sampling technique to generate new underlying asset trajectories using the risk-neutralized terminal distribution.

- Independent risk-neutralization at each time step:

This first technique consists in taking a cross-section of the simulated trajectories at each time-step and then performing the Esscher transformation for each of these cross-sections. This would yield a risk-neutralized distribution at each time-step, each using a different Esscher parameter. Despite being appealing for its

simplicity, this solution would not work well in conjunction with the least-squares Monte-Carlo method of Longstaff and Schwartz. Indeed, the non-parametric Esscher transform tilts the distribution by assigning a set of weights for each observation in the simulation. Therefore, risk-neutralizing at each time-step would yield different sets of weights for each cross-section of the trajectories. As a result, realizations along a single asset price trajectory in the simulation would carry different weights at different times. How then to perform the conditional expectation regression if the explanatory variables and the dependent variables do not have the same weight?

- Concurrent risk-neutralization for all time steps:

The second technique is a bit more appealing since the concurrent risk-neutralization would ensure that a single Esscher parameter is used and therefore a single set of weights is used to reweigh distributions for all cross-sections of the asset price trajectories. As a result, entire trajectories are reweighted at once. This method works better with the Longstaff-Schwartz algorithm since explanatory variables and dependent variables have the same weights. The least-squares regression simply needs to be adjusted to account for the weights associated with each trajectory. The problem with this approach is that the search for the single Esscher parameter that would risk-neutralize many distributions is over-constrained. Indeed, this is equivalent to searching for a single parameter to solve several equations. At the very best, an optimization algorithm may be able to find a solution that yields approximate risk-neutralization at some time-steps; more likely the optimization algorithm will simply fail.

- Re-sampling:

The third technique is probably the most promising as it uses the risk-neutralized terminal distribution to generate new trajectories that will be risk-neutralized by construction. This technique is quite popular in statistics and finance where it is usually called bootstrapping. Bootstrapping is a statistical method used for estimating the precision of a sample statistics by drawing randomly with replacement from a set of data points. In other words, it creates new distributions from observed distributions by sampling with replacement directly from the set of observed distributions.

Bootstrapping is a term first coined by Efron in his 1979 Rietz Lecture [159] to describe a re-sampling technique used to estimate the properties of some statistics such as the mean, median, or standard deviation of a distribution. In these cases, bootstrap samples are constructed by sampling with replacement a subset of the original distribution. The statistics of interest is then computed for each bootstrap sample and the variability between the results (standard error) can be analyzed to derive some confidence intervals for the statistics. For the problem under investigation, the essence of the bootstrap method is retained but the application is totally different. Similarly to the original application, the bootstrap method is used to sample with replacement from an original distribution but what is new is that the bootstrap sample is then used to generate asset price trajectories. In other words, the distribution of asset prices under the historical probability measure \mathbb{P} is first risk-neutralized using the non-parametric Esscher transform yielding a new re-weighted probability distribution. In turn, this risk-neutral distribution under probability measure \mathbb{Q} is sampled with replacement to yield bootstrap subsamples

which are used directly to construct trajectories under the risk-neutral probability measure \mathbb{Q} . These trajectories will all carry the exact same weight and therefore the Longstaff-Schwartz least-squares Monte Carlo method can be applied. The method is illustrated in Figure 20.

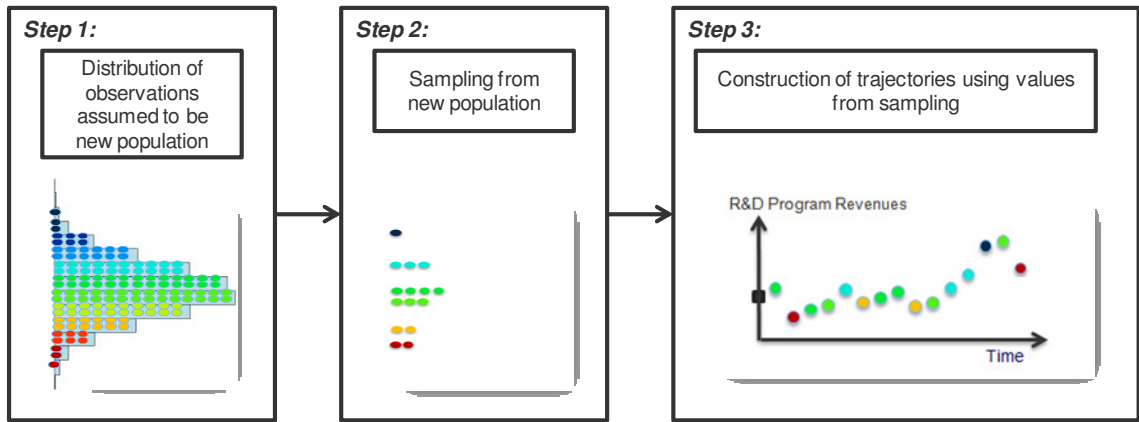


Figure 20: Bootstrap method to generate trajectories

For the pricing of options, special attention needs to be paid as to when it is appropriate to use a bootstrapping technique. Indeed, instead of simply generating distributions, the bootstrapping technique will be applied to simulate the realization of a stochastic process. In fact, bootstrapping will no longer generate simple distributions but rather trajectories or time-indexed distributions. If the original process to be simulated has some serial correlation properties, these would need to be accounted for in the bootstrapping method since a naïve bootstrapping does not induce any serial correlation. In order not to complicate matter further, the simplifying assumption of lack of serial correlation will be imposed, at least initially.

Provided that no serial correlation exists, another update to the method is required for the purpose of risk-neutral trajectory generation. Indeed, the bootstrapping method

usually starts with a sample of observations for which each individual observation carries a similar weight. In other words, there is a uniform distribution of each observation used as the original distribution. However, for the purpose of valuing options, the observations have already been reweighted during the risk-neutralization process using the non-parametric Esscher transform. As a result, the individual weights associated with each observation have to be accounted for when re-sampling to ensure that the risk-neutral property is preserved and carried over to the trajectories to be generated.

4.4.7 Summary

In summary, the bootstrap technique enables practitioners to re-sample a non-parametric risk-neutralized distribution for the purpose of generating trajectories for the underlying asset price. The risk-neutral property of the original distribution will be carried over to the sample of trajectories thus constructed. This transparent and intuitive method enables the use of the least-squares Monte Carlo technique to value path-dependent options which are particularly common in real options applications, and to approximate the early-exercise boundary which is useful to decision-makers to substantiate the need to act now, delay, or abandon an investment. This leads to the following set of research hypotheses regarding the selection of a technique to evaluate path-dependent real options as well as an enabling technique to do so.

Hypothesis 1.1.3 – Least-squares Monte Carlo approach for real options

Real options with early-exercise properties may be analyzed using a least-squares Monte Carlo approach to both estimate their value and derive the early-investment policy and

the optimal-investment boundary. The optimal-investment boundary may be used by decision-makers to substantiate the need to invest now, delay, or abandon an investment.

Hypothesis 1.1.4 – Bootstrapping for trajectory generation

Real options with early-exercise properties may be analyzed using a bootstrapping technique to generate risk-neutral trajectories for the evolution of the research and development program values via simulations and re-sampling.

4.5 Meeting the Georgetown challenge and more

Having navigated the waters of financial engineering and corporate finance, it is now time to summarize the results and provide a way forward that would both meet the challenges identified earlier in this thesis regarding the assessment of long-term staggered investments in aircraft development programs while meeting the requirements set forth by the Georgetown challenge mentioned in section 4.2.2 (Copeland and Antikarov [106]). So far, real options analysis has been advocated as being a superior technique to assess the economic viability of long-term corporate investments featuring flexibility while being exposed to many market risks. Out of the techniques surveyed, real options analysis was the only one that could account for the flexibility offered to management to revise program development roadmaps given the realization of the value-driving uncertainties.

Within the context of real options analysis, Monte Carlo techniques have been identified as superior because they can handle many different types of uncertainties while accounting for correlations between uncertain metrics. The major advantage of Monte Carlo techniques for real options pricing is that they are versatile enough to eliminate the

need for major revisions whenever the surrounding conditions change. Monte Carlo techniques also enable the analysis of staggered corporate investments with many different dynamics and with flexible decision tollgates, meaning that subsequent investments might be rushed forward or delayed depending on the evolution of market uncertainties.

Besides, within the context of Monte Carlo simulations for real options analyses, the use of the Esscher transform and its non-parametric counterpart for risk-neutralization provides a working framework sufficiently simple and versatile for wide acceptance amongst practitioners. The parametric Esscher transform can be used whenever the dynamics of the staggered investment value are well known and can be modeled using some classic stochastic processes. The non-parametric counterpart can be used whenever data is not sufficient to both assume and calibrate an underlying model or whenever the dynamics of the staggered investment value are too complex or too peculiar to be properly fitted with standard models.

Finally, the popular marketed asset disclaimer assumption underpins the whole methodology and claims that the analyst *subjective* evaluation of the staggered investment is the best *objective* assessment of said investment, and that its intrinsic value without flexibility can play the role of a synthetic asset allowing to make the market complete. In case this assumption is not accepted, the Esscher transform and its non-parametric counterpart actually select a specific probability measure that is consistent with the preferences of economic agents best described by a power or iso-elastic utility function. Table 22 maps the requirements of the Georgetown challenges to the Monte Carlo-based and Esscher-transformed valuation of staggered investments.

Table 22: Addressing the challenges facing the analysis of long-term corporate investment analyses featuring flexibility

| | | Monte Carlo-based and Esscher-transformed real options approach |
|---|---|--|
| Georgetown Challenge Requirements (Adapted from Copeland and Antikarov [106]) | Intuitively dominates other decision-making methods | <ul style="list-style-type: none"> • Ability to capture the flexibility in decision making • Recognize the value created by active and astute management |
| | Capture the reality of the problem | <ul style="list-style-type: none"> • Ability to handle optimum timing issues related to decision-making using American-type options • Ability to handle staggered investment programs with decision gates using compound options |
| | Use mathematics that everyone can understand | <ul style="list-style-type: none"> • Esscher transform ensures that risk-neutralization is performed in a transparent and tractable way • Non-parametric Esscher transform removes the requirement to calibrate complex models |
| | Rule out the possibility of mispricing by eliminating arbitrage | <ul style="list-style-type: none"> • Esscher transform provides the price that would be enforced by arbitrageurs in a complete market • Esscher transform provides the price corresponding to the preference of economic agents with iso-elastic utility functions in the case of incomplete markets |
| | Be empirically testable | <ul style="list-style-type: none"> • Tough requirements as there are no published transacted prices for these investments • Only heuristic argumentation can substantiate whether the method provides acceptable solutions |
| | Appropriately incorporate risk | <ul style="list-style-type: none"> • Handling of technical and market risk separately with technical risk analyzed with decision trees • Possibly difficult to estimate volatilities of some particular risks if no prior history exists |
| | Use as much market information as possible | <ul style="list-style-type: none"> • Ability to use market information whenever possible to model the dynamics of the uncertainties driving the development program value |
| Additional requirements | Ability to capture a complex reality with intertwined uncertainties | <ul style="list-style-type: none"> • Monte Carlo simulations allow the use of many different stochastic behaviors for uncertainties • Monte Carlo simulations allow the modeling of correlations between some sources of uncertainties |
| | Ability to visualize uncertainties and the decision process | <ul style="list-style-type: none"> • Visualization of the evolution of uncertainties affecting the decision process • Visualization of the evolution of the development program value over time |
| | Ability to handle corporate investments featuring exotic options | <ul style="list-style-type: none"> • Recent Monte Carlo methods allow analyses of programs with potentially moving decision tollgates and therefore the search for optimum investment timeframes |
| | Ability to converge to a solution in a timely manner | <ul style="list-style-type: none"> • Use of bootstrapping methods allow a reduction in computational time to generate trajectories of program values used for Monte Carlo simulations |

CHAPTER 5: SCOPING THE PROBLEM: RESEARCH QUESTIONS AND HYPOTHESES

5.1 Revisiting the problem

Following the review of the techniques most appropriate for the evaluation of staggered investments in research and development programs facing uncertainties and competitive pressure, it is now appropriate to come back to the original problem statement, synthesize briefly the identified issues, restate the research questions that are to be addressed, and identify how these map to the set of hypotheses proposed so far. In summary, the aim of this research is to analyze and evaluate investments made in the aircraft and engine manufacturing industry. Let's now review the different parts of this sentence and what they entail.

The aircraft manufacturing industry

The aircraft manufacturing industry is best currently described as an oligopoly under assault. The term oligopoly refers to the fact that there are very few competitors and each of them is specialized in just one or two segments of the industry. For commercial aviation, there are generally no more than two major competitors within one single market segment. *Airbus* and *Boeing* dominate the upper end of the industry with aircraft over 100 passengers; *Embraer* and *Bombardier* dominate the regional jet segment with jets seating between 50 and 100 passengers; and finally *ATR* and *Bombardier* dominate the regional turboprop segment with turboprop seating between 40 and 80

passengers. For general aviation, the market is a bit fuzzier but still dominated by no more than three or four major competitors within each industry segment. *Cessna*, *Gulfstream*, *Bombardier*, and *Dassault* dominate the upper end of the industry with business jets; *Beechcraft* and *Piper* dominate the business propeller aircraft segment; *Cessna*, *Cirrus*, and *Diamond* dominate the small recreational aircraft segment; and finally *Cessna*, *Cubcraft*, *Flight Design*, and *Remos* dominate the light sport aircraft segment.

If one extends the aircraft manufacturing industry to some of its main suppliers such as engine manufacturers, a similar type of concentration is observed: *General Electric* in partnership with *Snecma*, *Rolls Royce*, and *Pratt & Whitney* dominate the turbofan manufacturing business segment while *Lycoming*, *Continental*, and *Rotax* dominate the internal combustion engine business segment.

These oligopolies are under attack and have recently attracted interest. Whether it be for the consistent profitability of oligopolists over the years or for an eagerness to become less reliant on foreign industries, many countries have recently pushed for the development of a local homegrown aerospace industry. Even if the barriers to entry are high due to the demands of airlines for efficient and reliable aircraft at entry into service, the political push may prove successful in the long-term and some of these new manufacturers may emerge to become formidable competitors to incumbents. As a result, an increase in competitive pressure is observed and needs to be factored in by established manufacturers. In fact, it is no longer business as usual in a duopoly: incumbents must now ensure that their designs are competitive with all these new offerings. With this

description of the aircraft manufacturing industry, the first observation was stated as follows:

Observation 1:

Consistent profitability and politics stir up the interest for a homegrown aircraft development industry which leads to a substantial increase in the competitive pressure. Without a corresponding growth in the aircraft demand, manufacturers will need to account for the competitive environment early-on in the design to ensure the business plan is sound and the product portfolio is both competitive and well positioned in the market.

Investments in the aircraft manufacturing industry

Investments in the aircraft manufacturing industry can be described by a single word: *substantial*. In fact, these investments are so large that a common saying within the aerospace community is that *aircraft and engine manufacturers are betting the future of their companies whenever a new product is launched* [160]. Once an engine or aircraft development program is launched and manufacturing starts, there are billions of dollars at stake and very few opportunities for redemption if the business case is ill-founded. Under these circumstances, the development of a sound business case should include the study of sensitivities, robustness, and contingency plans. This led to another observation:

Observation 2:

In a competitive industry with long development cycles, there are few opportunities in the later part of the development process for manufacturers to change course as the uncertain environment unfolds. In this context, robust design simulation must be coupled

with extensive competitive scenario investigations to ensure that the realization of uncertainty does not undermine a design that otherwise meets all customer requirements.

Another characteristic of investments in the aircraft manufacturing industry is the uncertainty surrounding them. As aircraft developments are long and production runs often exceed a decade, there are ample opportunities for the surrounding environment to change and evolve while aircraft designs remain frozen due to the costs of certifying even minor upgrades and improvements. This led to a third observation:

Observation 3:

Aircraft and engine developments are characterized by longer and longer development cycles and are therefore subject to significant risk due to the uncertain and volatile business environment. Design methods and design processes must evolve accordingly to provide enough flexibility to managers to steer programs into profitable directions as the uncertainty unfolds.

Besides, aircraft operations are affected significantly by factors outside the control of manufacturers and operators such as new noise regulations, new emission regulations, and the constantly evolving price of jet-fuel. In this context, it is of paramount importance to both recognize and use the flexibility offered to management during the early phase of development to alter course and steer projects into profitable directions as uncertainties unfold. However, traditional investment evaluation methods are inadequate to capture this flexibility offered early in the design and to handle long-

term research and development programs. New methods must therefore be developed and this led to the fourth observation stated:

Observation 4:

In a volatile industry sensitive to business cycles, uncertain energy prices and evolving customer requirements, managerial flexibility defined as the ability of management to actively steer research and development programs into profitable directions is valuable and must be accounted for when business plans are laid-out. Traditional capital budgeting methods do not usually account for this flexibility and consequently undervalue significantly long-term aircraft and engine developments.

Analyses and evaluation of investments in the aircraft manufacturing industry

Analyzing and evaluating these investments entail many different kinds of investigations, from market investigations to technical analysis and financial assessments. Investigating all these aspects would be beyond the scope of a research thesis. Emphasis was consequently put on one aspect that is both specific and relevant to this industry: the timing of development programs to optimize profits and the associated performance evaluation metrics. Indeed, because of the length of development cycles, timing becomes critical to ensure a continuum of development programs while still meeting budget constraints. Windows of opportunities appear for the infusion of new technologies into new designs as well as for the launch of new products, and these windows must be recognized to maximize profits. This led to the final observation stated below:

Observation 5:

In an industry where manufacturers can neither afford to have a gap in their development pipeline (to retain skilled workforce) nor develop two clean-sheet designs concurrently

(due to limited funding and limited engineering capabilities), the time at which technologies become mature enough for commercial application becomes crucial. These timing issues need to be anticipated with both the company and the competition product development pipelines in mind.

5.2 Research questions and hypotheses

Some interesting aspects resulting from several observations of the aircraft and engine manufacturing industry and its evolution over the past decade have been presented. These aspects were highlighted for two reasons. First, they may substantially impact the state of the business for the coming years. Second, because these aspects will probably alter the current state of the business, a rethink of some of the methods and processes currently used by the industry may be warranted. In this context, the overarching research question was formulated as follows:

Overarching Research Question

In the context of aerospace research and development optimization, how can value-based design methods be improved to identify precursors of technological and market opportunities while reflecting the specific challenges associated with long-term and uncertainty-plagued developments, and while accounting for the competitive nature of the business?

Drawing on the literature review presented earlier, a set of three high-level hypotheses named “*method hypotheses*”, are formulated to answer the overarching research question. The first hypothesis deals with a proposed improvement to current value-driven methodologies to handle long-term and uncertain development programs and to account for the flexibility offered to decision-makers to take advantage of

opportunities. This hypothesis proposes a way forward to improve current state-of-the-art economic assessment methods. The second hypothesis deals with a proposed improvement to current viability assessments by the introduction of competitive aspects early-on during the economic analysis of future concepts. The third hypothesis proposes a concurrent use of these two improvements to yield better evaluation of long-term and uncertain research and development programs with staggered investments.

Hypothesis 1 — Real options for valuation with flexibility under uncertainty

Within the context of aerospace research and development programs, real options methods enable the development of value-based design frameworks accounting for the staggered nature of investments and the value created by managerial flexibility in uncertain environments.

Hypothesis 2 — Game theory for investigation of economic robustness with competition

Within the context of aerospace research and development programs, game theoretic methodologies enable transparent and traceable analyses that allow decision-makers to better investigate the economic robustness of selected technology and product development streams in a competitive environment characterized by uncertain moves by competitors.

Hypothesis 3 — Combined real options and game theoretic analyses

Real options methodologies combined with game theoretic methodologies allow the identification of windows of opportunities and yield analyses superior in term of robustness to either of these two analyses performed independently.

Having introduced three method hypotheses advocating the use of both real options and game theoretic analyses, a thorough literature review leads to some further questioning: how can these methods be adapted to the problem under review? In fact, they lead to a set of three second-level research questions, “*modeling research questions*” and their associated second-level hypotheses, “*modeling hypotheses*”.

Indeed, assuming that real options inspired methodologies present the best framework for the analyses of long-term and highly uncertain research and development programs, the first modeling research question is related to the pertinence of these models to evaluate non-traded investments in the aerospace industry. The second modeling research question pertains to the modeling of flexibility offered to decision-makers to time investments optimally since investments are usually made whenever decision-makers feel like the conditions are optimal. Finally, assuming that game theoretic methodologies present the best framework for the analyses of competitive issues, the final modeling research question asks how to model the competitive scenarios that are prevalent in the aerospace industry.

Research Question 1.1 — Creation of an option-thinking framework

Within the context of uncertain product and technology investment analysis, how can state-of-the-art option-based valuation methods be improved upon to ensure their domain of application is consistent with their underpinning assumptions?

Research Question 1.2 – Managerial flexibility and timing of investments

How can the flexibility offered to management to optimally time the launch of new product and technology developments be accounted for in a real options framework?

Research Question 1.3 – Competitive scenario modeling

How can game theoretic analyses be used to adequately model competition in the aerospace industry and how can they be used to identify profitable product and technology development strategies?

A review of the existing literature on financial options and real options points to the use of one particular type of option, which is suitable for the analysis of investments in the aerospace community. It is able to relax one assumption of the more popular models while accounting for the flexibility to optimally time investment decisions. Similarly, reviewing the literature and observing the nature of competition in the aerospace industry oligopolies lead to a specific type of competitive scenario to investigate. Indeed, while research and developments are usually made somewhat simultaneously, actual decisions to launch new aircraft and engine programs are made in a sequential fashion, with one competitor being the leader and other competitors waiting to see what happens before launching their own products. This yields the following set of modeling hypotheses:

Hypothesis 1.1 – Path-dependent options to model managerial flexibility

As uncertainty unfolds, technological and market opportunities emerge and disappear. Flexible management and flexible timing of investment decisions allow the maximization of the upside potential of these opportunities. Path-dependent real options may present a means to model the flexibility offered to management in timing technology development decisions.

Hypothesis 2.1 – Equilibrium in sequential moves for competitive scenarios

Equilibrium-types of solutions in sequential competitive scenarios provide means to quantitatively measure the impact of competing designs on profitability and to identify robust strategies.

Beyond the assumption dealing with the type of real options suitable for corporate investment analyses, popular real options models are also too constrained in the requirements they set for the evolution of the underlying investment. This leads to further reckoning and questioning whether relaxing some of the process-related assumptions would pay in the long-term for the general applicability of real options. This yields the first third-level research question also named “*technical research question.*”

Research Question 1.1.1 – Enlarging the domain of applicability of real options

How can the domain of application of current state-of-the-art real options methods be extended to include product and technology investments with value processes not following classic geometric random walks?

A review of the literature on financial options shows that in order to model processes more complex than classic geometric random walks, simulation becomes a viable, if not necessary, option. Indeed, simulation seems to be the preferred way to handle multi-dimensional options featuring correlations between underlying sources of uncertainty and to handle discontinuous stochastic processes following external shocks in the market. This leads to the following third-level hypothesis also referred to as “*technical hypothesis.*”

Hypothesis 1.1.1 – Monte Carlo methods for real options analyses

Within the context of the aerospace industry, Monte Carlo based real options methods present the best approach to solve for the arbitrage-free value of development featuring flexibility.

As much as simulation for the pricing of options seems pertinent and straightforward, there may still be some obstacles on the way to implementation by practitioners. Thus, a second set of technical research questions is formulated to solve the issue related to the requirement to perform the simulations under the equivalent risk-neutral probability measure:

Research Question 1.1.2 – Improving Monte Carlo methods for real options analyses

Monte Carlo simulations based methods seem appropriate to value corporate investments featuring managerial flexibility and programmatic optionality. With usability by practitioners in mind, how can these methods be modified to alleviate the complexity of finding the proper equivalent probability measure required for the expectation computation while maintaining their rigor?

Research Question 1.1.3 – Adaptation of Esscher transform for pricing real options

How can option pricing by means of Esscher transform be adapted to a product and technology development analysis within the context of a real options methodology?

Drawing on the literature research, the non-parametric Esscher transform seems to be the most appropriate technique to perform the change of measure required for option pricing using simulations. Another technical hypothesis is therefore formulated to answer the two previous technical research questions:

Hypothesis 1.1.2 – Non parametric Esscher transform for pricing real options

Real options valuation via non-parametric Esscher transforms is a promising framework for staggered investment analyses. It is based on rigorous foundations, offers a clear and transparent methodology for practitioners, and uses probabilistic techniques widely accepted within companies.

Coming back to the modeling research questions and hypothesis related to the flexibility offered to managers to time investment decisions, another angle investigated is the detection of optimal times to launch research and development endeavors. This detection study leads to the identification of precursors of successful research and development programs through the definition of the early-investment boundary. This immediately leads to a new set of technical research questions. The first technical research question deals with the ability to define the early-investment policy and therefore approximate the optimal early-investment boundary. The second technical research question deals with the ability to use early-investment boundary approximation schemes in conjunction with the previously proposed Esscher-based risk-neutralization technique.

Research Question 1.2.1 – Trigger boundaries to detect trigger events

How can trigger boundaries be defined and used for the identification of precursors of successful development programs?

Research Question 1.2.2 – Bootstrap for trajectory generation

How can algorithms approximating the early-investment boundaries be used in conjunction with the non-parametric Esscher risk-neutralization?

Drawing on the literature review of techniques developed and used in the quantitative finance industry, the least-squares Monte Carlo approach developed to value American and Bermudan financial options appears to be the most appropriate method to value path-dependent real options. However, because the technique relies on regressions and because the non-parametric Esscher risk-neutralization assigns different weight factors for each and every observation, the two techniques are not directly usable together. Consequently, these techniques need to be updated by changing the way trajectories are generated during the simulation. In fact, trajectories have to be simulated again using a re-sampling method to yield new trajectories that are risk-neutral. A new set of two technical hypotheses is therefore formulated:

Hypothesis 1.1.3 – Least-squares Monte Carlo approach for real options

Real options with early-exercise properties may be analyzed using a least-squares Monte Carlo approach to both estimate their value and derive the optimal trigger boundary. This trigger boundary helps decision-makers substantiate the need to invest now or abandon a development.

Hypothesis 1.1.4 – Bootstrap for trajectory generation

Real options with early-exercise properties may be analyzed using a bootstrap technique to generate, under the equivalent martingale measure, trajectories representing the evolution of the development program values.

5.3 Matching research questions and hypotheses

The diagram in Figure 21 summarizes and maps the different research questions underpinning this research endeavor to their corresponding hypotheses.

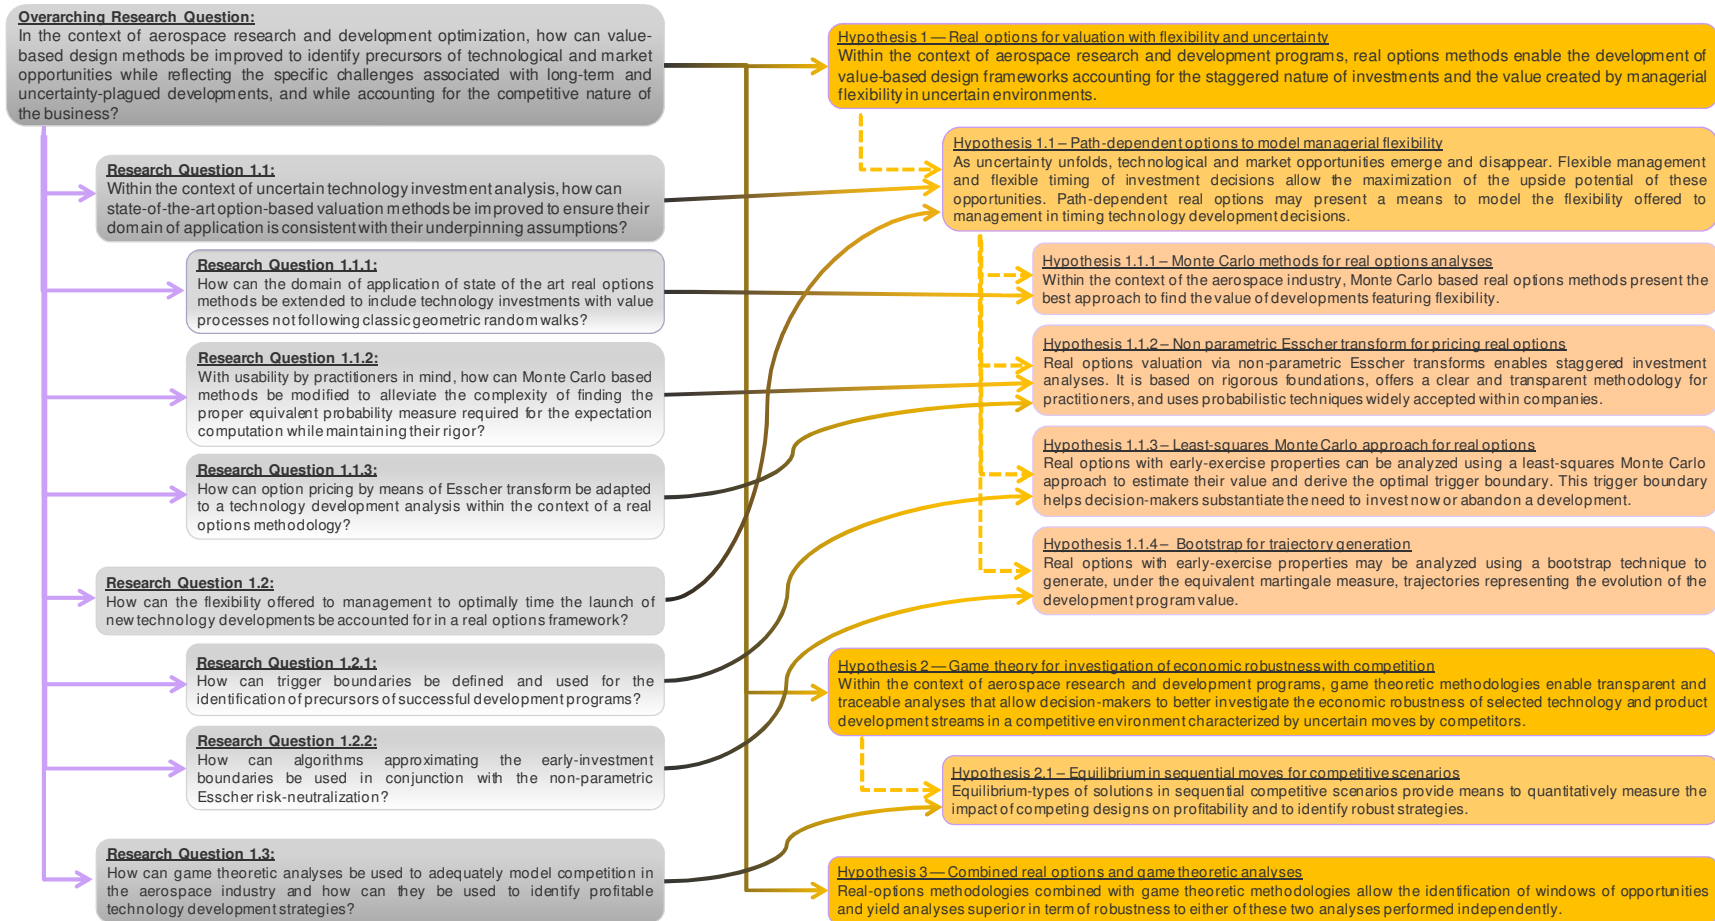


Figure 21: Research questions and hypotheses

CHAPTER 6: METHODOLOGY

In this section, a novel methodology is proposed to tackle the issues presented so far. This methodology is introduced step by step and is progressively contrasted with current traditional methodologies aimed at analyzing research and development programs. However, before the methodology is introduced, a proof-of-concept study is formulated. The aim of this pilot study is to introduce and demonstrate the applicability of the proposed methodology. Later on, this case study is used to verify the approach undertaken and validate some of the hypotheses formulated.

6.1 Brief overview of an industry problem

In this pilot study, a Performance Improvement Package (PIP) is being proposed as a means to improve the operating economics of a currently out of production turbofan engine. The aircraft manufacturer has identified a gap in its development stream which makes it possible to develop, certify, and produce the package. Decision-makers have to identify whether the conditions are currently optimal for the commercial launch of this product and whether it makes sense to commit resources to this development now. If not, there is a wide window to actually launch the development. The manufacturer can then delay its initiation to wait for trigger events that will ensure that the development is a commercial success.

6.1.1 *Windows of possibilities and windows of opportunities*

For the pilot study under investigation, the manufacturer has identified a gap in its development stream between two periods of high activity. The first period of high

activity is related to a previous development requiring substantial engineering resources to complete the detailed design and to get certification. The second period of high activity concerns a future development for the replacement of a current engine design that is getting obsolete. This second program is therefore deemed vital for the profitability of the manufacturer and is projected to tie its engineering resources for several years onwards. In between, there is a development gap during which the manufacturer has no projected development and during which engineering resources might be available. This is an unfortunate situation for aircraft and engine manufacturers as they have to retain the workforce to keep skilled and experienced engineers in-house for future programs. In this context, a window of possibility for the development of the PIP program is defined as the ability to undertake a development program. This situation is depicted in Figure 22.




Figure 22: Timeline of manufacturer development stream

Of interest however are not windows of possibilities but rather windows of opportunities which are defined as the timeframe during which, and the condition for which, launching a new development program is best. If decision-makers invest too early within the window of possibility, they only have limited information and this is risky as the realization of uncertainty might undermine the development program. If decision-makers invest too late, risk also increases since the target market size is reduced as airlines ground older aircraft and become reluctant to invest in an ageing fleet. To be meaningful to decision-makers, a window of opportunity has to be contained within a

window of possibility. Therefore, the largest window of opportunity is the window of possibility.

In addition, windows of opportunities are not static: they morph in real time to adjust to the new reality that unfolds. Increasing energy prices drive the demand for more efficient engines and a low capital expenditure retrofit to reduce fuel consumption looks like an attractive option for airlines. Alternatively, emerging competitors with new engine designs or even competing improvement packages may impact the demand for the PIP and therefore impact the profitability of the program. Combined together, these effects may either stretch or constrict the window of opportunity. This dynamic process is depicted in Figure 23 where the impact of progressive aircraft retirement and the impact of competition on the opportunity window are highlighted.

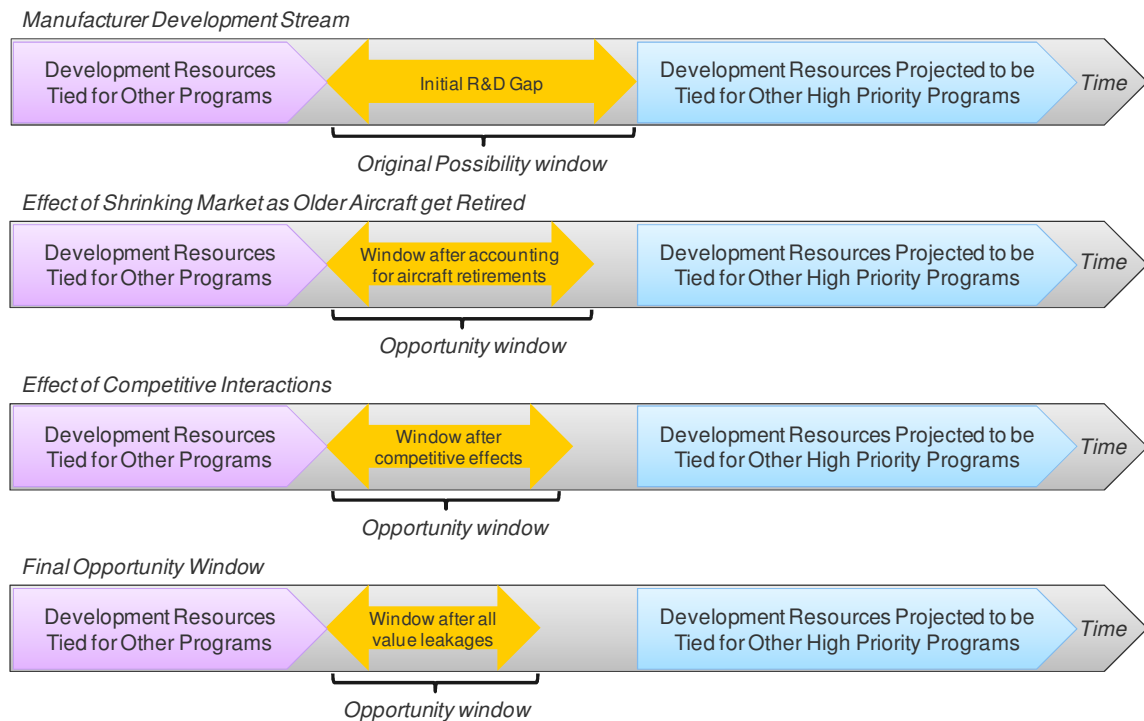


Figure 23: Value leakages and their effects on the opportunity window

6.1.2 Development of the Performance Improvement Package

The development process for the Performance Improvement Package may be described as a staggered development process featuring decision tollgates. It is articulated around four main phases, starting with the initial market research and conceptual design, followed by preliminary and detailed developments, followed by certification and testing, and finally ending with production. Each of these phases is separated by a decision tollgate at which point management can display some flexibility and decide whether to pursue, delay, or abandon the development altogether if the market conditions are not right. This development timeline is shown in Figure 24. If the development program goes ahead, then additional funding is committed and spent during the following phase. All four of these phases do not have the same resource requirements: detailed development and certification and testing are the most expensive. Consequently, it is unlikely that the program will be abandoned at the fourth decision tollgate given that the following phase is relatively cheap and that so much has already been spent during the previous phases. The choice of delaying or abandoning the development is nevertheless very relevant at the second and third decision tollgates if conditions are not favorable.




Figure 24: Development timeline and associated milestones

6.1.3 Identification of decision windows

The decision windows are time windows during which a decision to fund the next phase of the PIP development program must be reached. To do so, the overall window of

possibility as well as the hard constraints regarding the minimum time required to perform each of the four development phases is used to derive four sub-windows of possibility. These four sub-windows of possibility indicate the time-windows during which a decision to fund the initial market research, the detailed development, the certification and testing, and finally the production must be made. They are consequently referred to as decision windows.

In order to derive these decision windows, an investigation is carried out to determine the latest times at which the four decisions need to be made, as well as the earliest times at which the four decisions can be made. The process of figuring out these decision windows is illustrated in Figure 25.

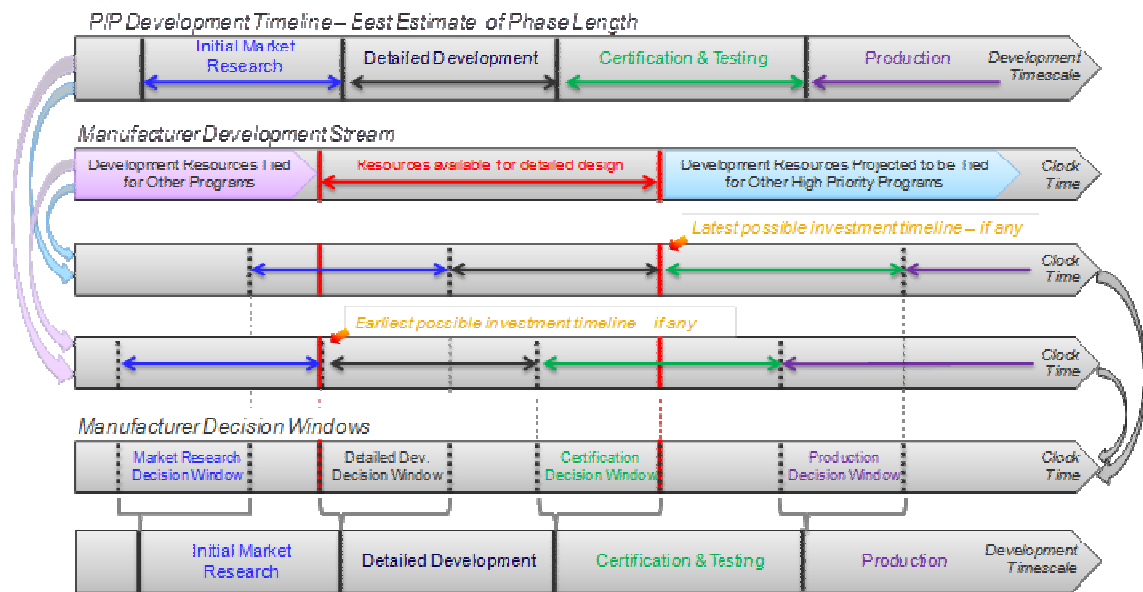


Figure 25: Deriving decision windows

6.1.4 Objectives to be attained and expected results

The objective of this proof-of-concept study is to investigate the optimal conditions for the launch of the development program. This includes finding out the optimal timing of decisions and the corresponding state of uncertainties leading to a

successful development program. To do so, the baseline investment timing is introduced as the latest time at which investment decisions can be made for all four decision windows. Any time a decision is made before this *baseline investment timing*, the decision is called an *early-investment decision*.

The *investment policy* is defined as the policy of timing investments optimally. In other words, it means that the investment policy maximizes the value of the PIP development program for the company. In doing so, the investment policy determines an *early-investment boundary*. The early-investment boundary is the set of external conditions (time and state of uncertainties) that makes investing early optimal. Notional early-investment boundaries are given in Figure 26 for each decision window pertaining to the PIP development program.

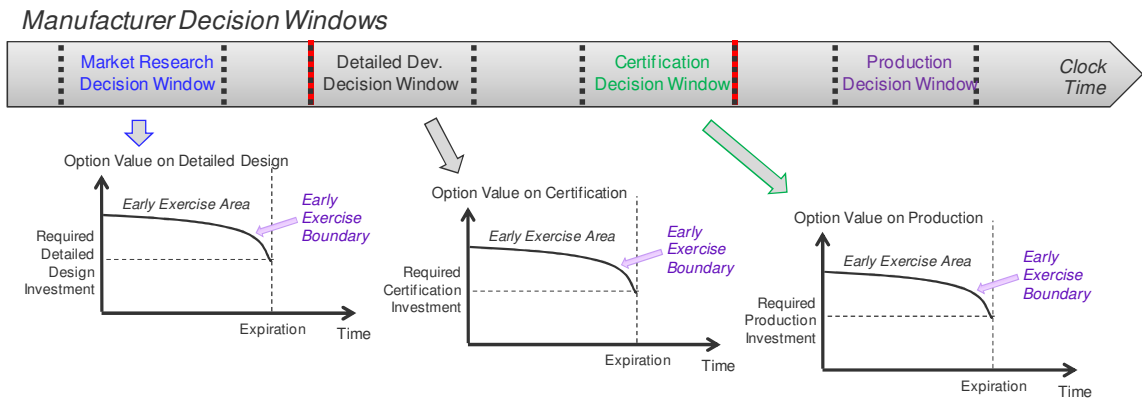


Figure 26: Early-investment boundaries at each decision window

The concept of early-investment boundary is interesting for decision-makers as it allows them to substantiate whether acting now or delaying the decision is optimal. Indeed, by comparing the current state of the business (current time and current observations of the uncertainties) to the early boundaries, decision-makers are able to identify whether they are in an invest-immediately area or whether they get more value

by holding-off and waiting to get more information about the trajectories of the uncertainties. Investigating the shape of the early-investment boundary in a parametric environment yields many interesting observations:

- What is the impact of technical uncertainty on the early-investment boundary?
- How does not meeting the PIP performance targets impact the early-investment boundary?
- How do value leakages impact the shape of the early-investment boundary?
- Which combinations of uncertainties substantially impact the shape of the early-investment boundary?
- How can these combinations be classified to yield a list of precursors or trigger events of potentially successful research and development programs?

6.2 Methodology development

In this section, a typical business case evaluation method is first reviewed in order to highlight gaps and indicate where the proposed improvements take place. Next, the FLexible AViation Investment Analysis method (FLAVIA) is presented in details.

6.2.1 Traditional evaluation method

Following the literature review presented earlier, a research and development methodology that is representing the current approach to business case construction and investment evaluation is set up. It starts with a set of scenarios that are probable and that may be ranked by their likelihood or by their outcome: worst case scenario, best case scenario, and most likely scenario. These scenarios are constructed using potential realization of still uncertain parameters and could be related to uncertain competitor

moves, uncertain technical achievements, and uncertain states of the economy. A morphological decomposition of the different uncertainties and their potential best case, worst case, and most likely realizations leads to a matrix representing the combinations of these uncertain parameters. The subsequent re-composition of these uncertainties and their potential outcomes yields a combinatorial number of scenarios.

Using these scenarios, a market analysis is performed. The goal of the market analysis is to separate the overall market into various market segments, each representing customers with homogeneous preferences. These preferences are built up using different attributes that can be represented by different metrics which may be quantitative or qualitative. The functions linking the level of one attribute to a level of preference are called single attribute utility functions. The tradeoff between these single attribute utilities is captured by weighting these functions and accounting for interactions. When combined together, these functions lead to a multi-attribute utility function which represents the overall preference of customers. In the remainder of this research, the multi-attribute utility function representing the overall preference of customers is referred to as the *overall evaluation criterion*.

Following the market segmentation, a preliminary assessment of the demand is formulated by investigating the market reaction to the new product offering. This is done by looking at how the new product is performing when used by the end-customers. This market reaction analysis is performed at the market-segment level and then aggregated using the size of each market segment. For aerospace research and development applications, the market reaction will be gauged by investigating how the new aircraft or the new technology is performing when operated on a network of flights representing the

typical operations of airlines within the market segment. This will yield a preliminary estimate of the demand for the aircraft or for the technology package to be infused into existing fleets.

In the last step, the research and development program profitability is assessed by looking at the demand for the product and estimating the investments required to fund the research and development program. The timeline of revenues and investments is used to derive a figure of merit representing the profitability of the program. Usually, a discounted cash flow analysis is performed at this point to estimate the net present value or the internal rate of return of the development program. This figure of merit is then used to substantiate whether the development program should or should not be launched. The different steps used to build a business case are illustrated in Figure 27. This process is referred to as the *traditional methodology*.

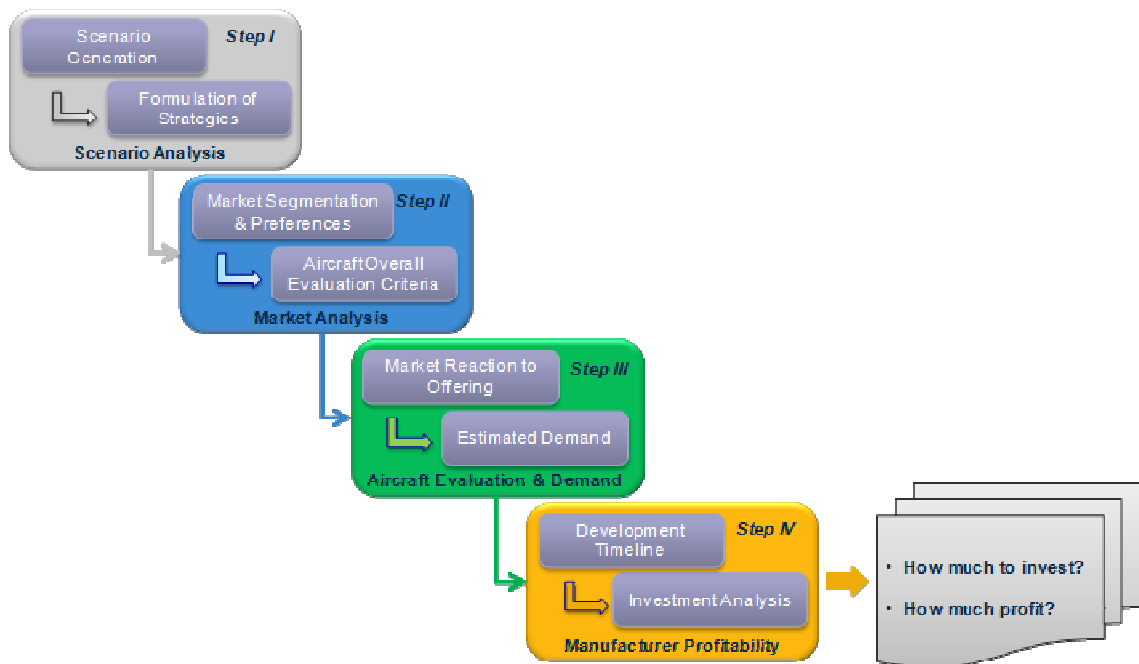


Figure 27: Traditional methodology to build and evaluate a business case

6.2.2 Proposed evaluation method

The proposed methodology builds upon the traditional methodology but makes use of advanced evaluation techniques presented during the literature review which aims at assessing staggered investments featuring flexibility. The proposed methodology is different from the traditional one in three fundamental ways:

- The profitability analysis is performed using a real option-based analysis to account for the value of flexibility offered to management to delay or abandon the research and development program. This is why the profitability analysis is renamed “staggered investment analysis”. A complete reformulation of this assessment is proposed using a novel real option-based approach constructed via cross-fertilization of techniques and methods used in the actuarial sciences, in statistics, and in the quantitative finance industry.
- The profitability analysis is linked back to the competitive scenario analysis and to the aircraft evaluation analysis to dynamically update the demand for the product as uncertainty unfolds in either favorable or unfavorable directions. This affects the payoff structure in the proposed real option-based methodology.
- The result from the profitability analysis is used and analyzed in a game theoretic setting to figure out what course of action is most profitable and most robust with regards to uncertain moves by the competition. This analysis is a game theoretic based analysis which aims at finding profitable and robust strategies.

This leads to a new depiction of the business case construction process highlighted in Figure 28. In this depiction, only high level interactions between the various steps of the proposed methodology are highlighted in order not to clutter the graph.




Figure 28: Proposed methodology to build and evaluate business cases in R&D

6.2.3 Detailed implementation of proposed evaluation method

In this section, the proposed methodology is further decomposed into more detailed steps to present a comprehensive review of the method implementation. It is articulated around nine main steps. The first steps deal with the competitive and market analysis, while the last ones deal with the profitability analysis and the detection of trigger events. More emphasis is given on these later steps since this is where most of the novelty and originality takes place.

Scenario Analysis

The scenario analysis itself is split into two steps with the first step being a morphological decomposition of the competitive environment and the second step being

an identification of the uncertainties driving the value of the research and development program. This is illustrated in Figure 29. In the first step, different scenarios regarding the state of the business are investigated. This includes an investigation of the different strategic options offered to the manufacturer as well as to its competitors, and an investigation of the uncertainties surrounding the development program that have potential to substantially affect its success.

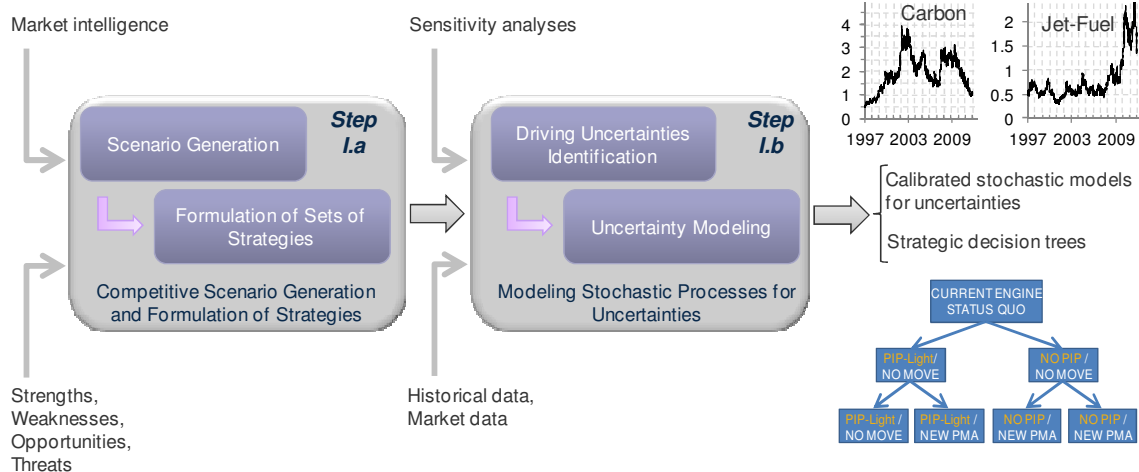


Figure 29: Scenario generation and uncertainty modeling

On the competitive side, the review includes the identification of potential tactical and strategic moves by competitors as well as the timeline associated with these moves. This review results in the formulation of either simultaneous or sequential competitive games for which many strategic moves are studied and an investigation of equilibrium inducing strategies is performed. For simultaneous competitive games, a matrix of strategies is formulated and is populated using results from the yet to be performed profitability analysis and some estimates of the profitability of competitors. For sequential games, decision trees usually referred to as extensive-form representations are used to visualize the timelines of actions made by the various competitors. This

representation is populated using results from the yet to be performed profitability analysis as well as some estimates of the profitability of competitors. Whichever representation is retained, the end-objective is to solve these games to find equilibrium types of solutions. These equilibrium solutions are defined by the set of actions and reactions from which none of the competitors have any incentive to deviate. They therefore define a set of competitively-robust strategies.

On the uncertainty side, the main sources of uncertainties likely to affect the economic success of the development program are reviewed. This review includes both idiosyncratic and market uncertainties although these will be treated differently.

Idiosyncratic uncertainties are uncertainties related to the manufacturer and its ability to deliver the product promised on time and according to guaranteed specifications. They include supply-chain uncertainty, schedule uncertainty, production ramp-up uncertainty, certification uncertainty as well as performance and technical uncertainty. Since there are no historical databases regarding these uncertainties, they cannot be regressed to fit time-series or stochastic models. Instead, these types of uncertainties are best analyzed using decision trees constructed by identifying worst case, best case, and most likely scenarios. Subject matter experts usually provide assessments about the likelihood of each branch in the decision trees as well as ranges for the uncertain parameters to be quantified. Probabilistic analysis can then be performed to assess possible outcomes by using Monte Carlo simulations using distributions for the uncertain parameters (usually triangular distributions using the minimum and maximum ranges supplied by subject matter experts).

Market uncertainties are uncertainties that affect the whole market, and therefore all competitors. These include air-transportation demand uncertainty, energy-price uncertainty, emission taxation, and other regulatory uncertainties. These types of uncertainties usually have a history and this history can be used to fit and calibrate time-series or stochastic models.

The output from these steps is two-folds: a set of scenarios to be investigated that combine both competitive and idiosyncratic uncertainty aspects, as well as a set of market uncertainties with their corresponding models.

Market Analysis

The purpose of the market analysis is to perform an operational decomposition to investigate which customers have similar types of operations to then recombine similar customers into homogeneous market segments for preference analysis. The process is illustrated in Figure 30. To do so, the type of network as well as the type of operations of potential customers is reviewed to find similarities among them. Similarities can be very broad, but in the context of aircraft and engine developments these include the typical flight lengths, typical number of hours flown per year, ease of access to capital, or the cost structure of the airlines. For instance, operators with higher flight-length to flight-cycle ratios (FL:FC) are usually more sensitive to fuel-burn, while operators without easy access to capital are more reluctant to embark in high capital expenditure acquisitions.

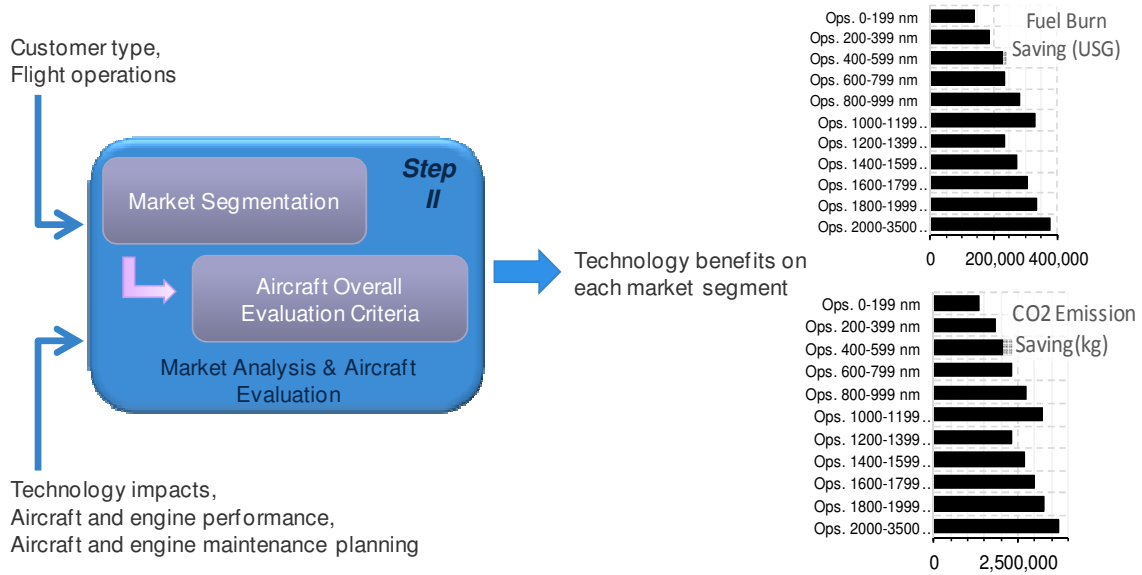


Figure 30: Market analysis and the determination of preferences

These homogeneous market segments enable the determination of a unique set of preferences for the aircraft operators involved. These preferences are built-up using different attributes which can be represented by different metrics. These metrics may be quantitative or qualitative. The function linking a product attribute to a level of preference is called a single attribute utility function. When several attributes must be accounted for, then tradeoffs between these attributes must be captured. This is done by weighting the various single attribute utility functions and accounting for first and second order interactions. When these single attribute utility functions are combined together while accounting for tradeoff preferences and interactions, one gets a multi-attribute utility function. This function represents the overall preference of customers within a market segment given a set of attributes. This function is referred to as the “overall evaluation criterion”.

Market Reaction and Demand Analysis

The market reaction and demand analysis step is performed to model the decision made by potential customers regarding the purchase of the aircraft, engine, or technology package. This analysis is split into two subtasks as depicted in Figure 31.

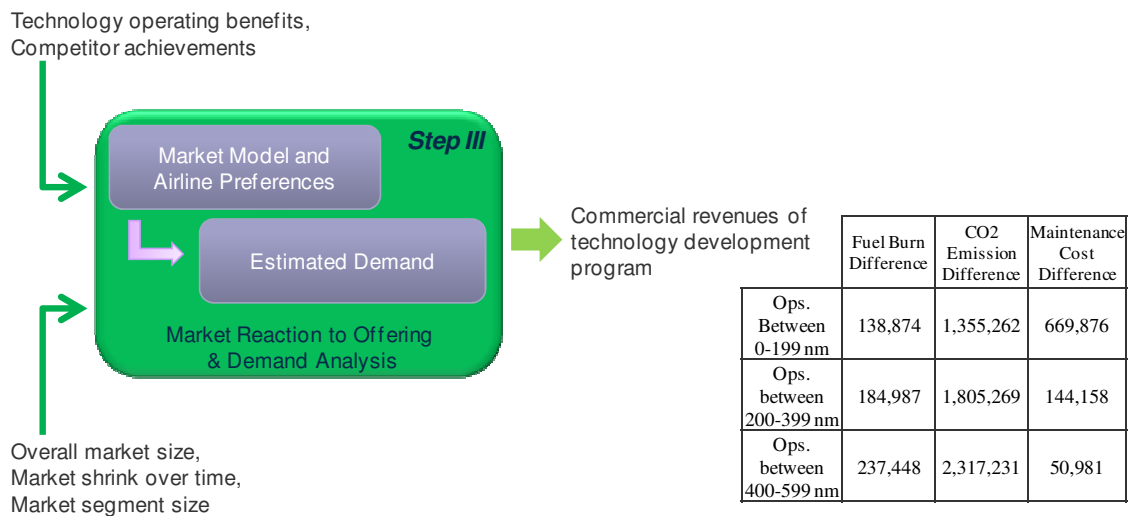


Figure 31: Market reaction and demand estimation

First, a competitive scenario generated during the very first step of the methodology as well as the overall evaluation criterion constructed in the second step of the methodology are used to establish what the competing designs are and to establish the value of each design to end-customers within each of the market segments.

Next, a decision choice model is used to translate these *value-to-the-customers* into purchasing probabilities and purchasing preferences. This step is quite difficult as no realistic calibration can be performed for the choice model due to the lack of historical data. Besides, even if historical data were to be available, it would most probably be contaminated by some noise and other intangible aspects that are not accounted for in this thesis such as political pressure, historical relationship, and loyalty between manufacturers and operators. In this context, a *simpler is better* philosophy is retained

whereby the value to the customer is selected as the exclusive metric to decide whether a market segment, as a whole, is willing to purchase the product.

Manufacturer Profitability

The manufacturer profitability analysis is divided into five tasks and represents the vast majority of the innovations proposed in this research. These tasks follow the different steps required to perform a real options analysis as inspired by the outcomes of the literature research and the resulting cross-fertilization.

First task

The first task illustrated in Figure 32 is to simulate the evolution of the value of the research and development program. These simulations are generated using one scenario obtained in the first step of the proposed methodology as well as the evolution of the market uncertainties. The scenario combines one possible competitive setting and one vector representing the levels of each of the idiosyncratic uncertainties. The market uncertainty evolutions are simulated using the calibrated models from the first step of the proposed methodology.

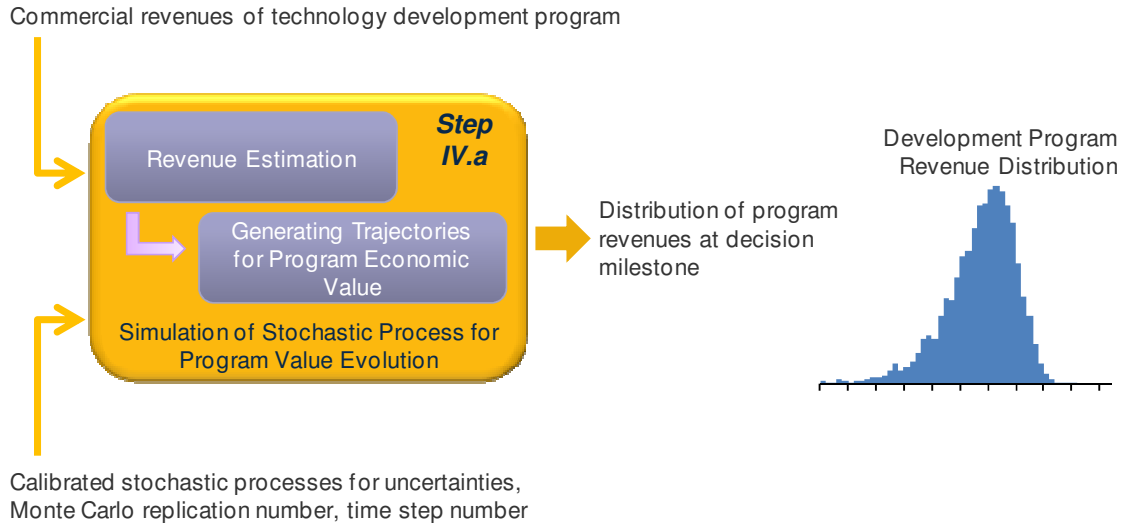


Figure 32: Simulation of the evolution of the R&D program value

This will result in the simulation of trajectories for the value of the research and development program. By using Monte Carlo simulations, the evolution space is sampled many thousands of times to get an approximate distribution of the likelihood of the possible values of the research and development program. Since these evolutions are simulated using stochastic processes calibrated with observed historical data, the evolution is simulated under the historical probability measure.

Second task

The second task is illustrated in Figure 33 and consists in transforming the distribution of the value of the research and development program from the historical probability measure to the equivalent martingale measure also known as the risk-neutral probability measure.

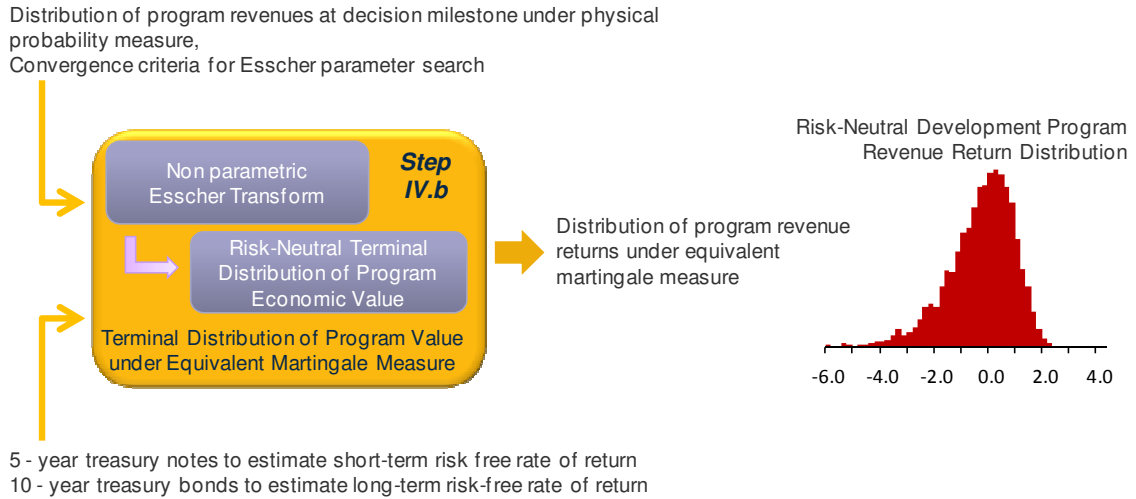


Figure 33: Change of probability measure and risk-neutral terminal value distribution

The non-parametric approximation of the Esscher transform, also known as exponential tilting, is used for this transformation. Provided some mild conditions are satisfied (stationary and independent increments) for the stochastic process generating these terminal distributions, the non-parametric Esscher transformation leads to a new equivalent distribution for the value of the research and development program. This new probability distribution is called the risk-neutral distribution. It has the exact same values as the original distribution but each realization now carries a weight to actually tilt the original distribution.

Third task

The purpose of the third task is to use the risk-neutral distribution generated previously to generate new risk-neutral trajectories for the evolution of the research and development program value. Following the literature review, new trajectories can be generated using a resampling technique known as bootstrapping. This is divided in two separate actions as illustrated in Figure 34.

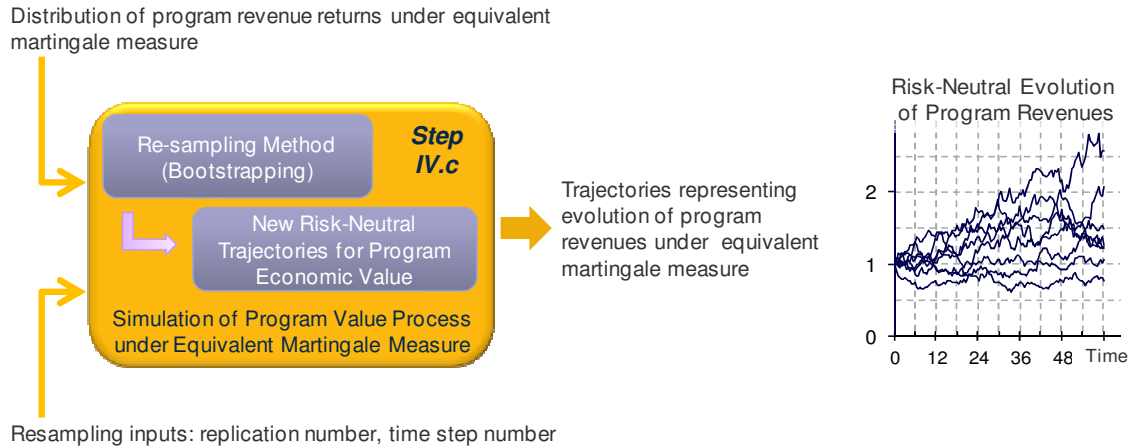


Figure 34: Simulation of R&D program value under risk-neutral probability measure

The first action consists in using the risk-neutral distribution obtained in the previous step and expressing each absolute value for the research and development program in terms of equivalent daily returns if the simulation for the time evolution is done using daily steps. The second action consists in generating new trajectories for the evolution of the R&D program value by repetitively sampling with replacement from the risk-neutral distribution of returns. This resampling is done accounting for the weight assigned to each observation during the previous transformation. This leads to new trajectories that are risk-neutral by construction.

Fourth task

The fourth task is dedicated to the evaluation of research and development investments having timing flexibility. This task investigates the possibility of investing early and defines an early-investment policy to determine when the time is optimal to invest earlier rather than later. This task is subdivided into two actions as illustrated by Figure 35.

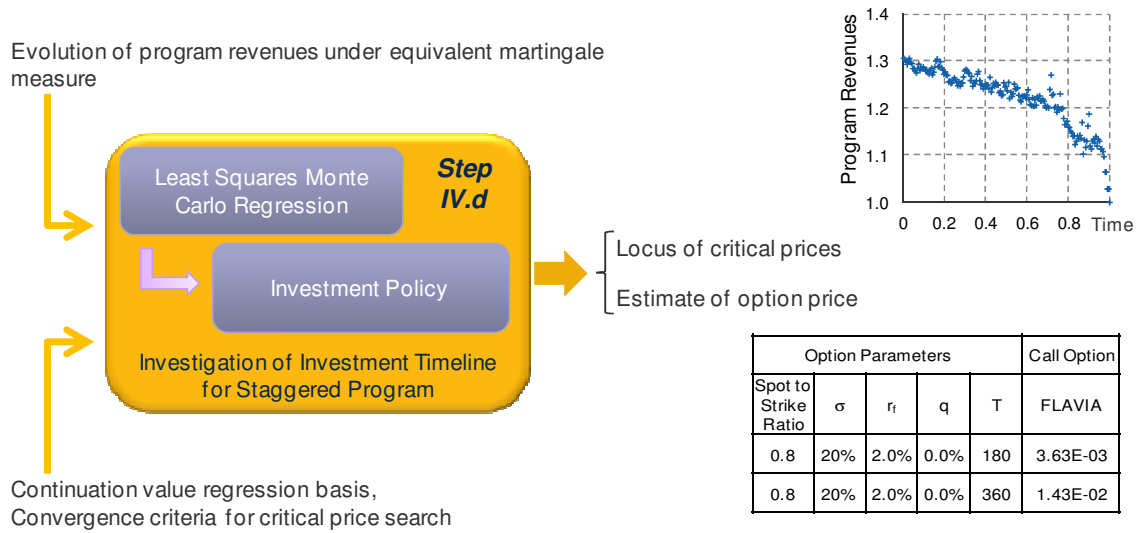


Figure 35: Derivation of the early-investment policy for path-dependent real options

Following the literature review and the hypotheses presented earlier, the algorithm of Longstaff and Schwartz using least-squares Monte Carlo regressions is used to determine whether investing early or delaying the investment by at least one period is optimal. The algorithm approximates the one-step-ahead conditional expectation for the value of the research and development program. By doing so, the algorithm provides two important pieces of information: the first is the value of the R&D program for the company and the second is an approximation of the early-investment policy.

Fifth task

Using the early-investment policy defined in the previous step, the fifth task derives the early-investment boundary and sets the stage for further analyses aiming at comparing the current situation (current realization of the uncertainties) and the early-investment boundary. This task is depicted in Figure 36.

The early-investment boundary is approximated by looking at the times and realizations of the uncertainties that led to an early investment and by looking at those that led to a delayed investment. This way, the analyst is able to detect the edge or trigger

point at which the decision to invest changes. By searching for these trigger points at each time-step in the simulation, the analyst is able to approximate an early-exercise boundary curve (if there is a single source of uncertainty) or an early-exercise boundary surface (if there are two sources of uncertainties).

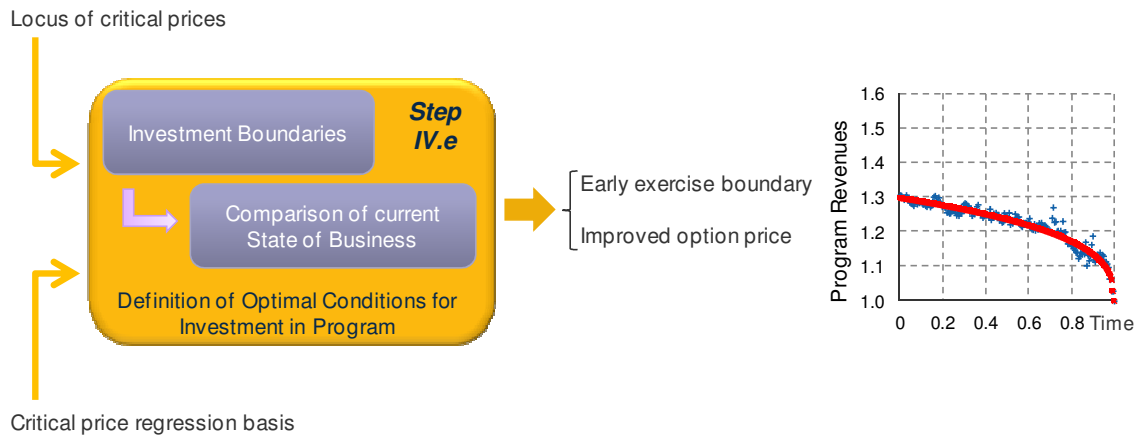


Figure 36: Analysis of the optimal set of conditions to launch R&D programs

This yields a set of conditions that are optimal to launch the research and development program. These can be interpreted as trigger events or precursors of successful development programs. At this point, further analyses are plentiful. In particular, the effect of technical uncertainties can be investigated. In other words, the sensitivity of the early-investment boundary with respect to the expected performance associated with the aircraft, engine, or technology package may be analyzed.

Overview of the entire methodology

Having detailed the different steps of the proposed methodology, all the pieces of the puzzle can now be assembled together to yield the novel real option-based and game theoretic inspired methodology to evaluate research and development programs in the aerospace industry and beyond. This novel methodology is depicted in Figure 37 on the following page.

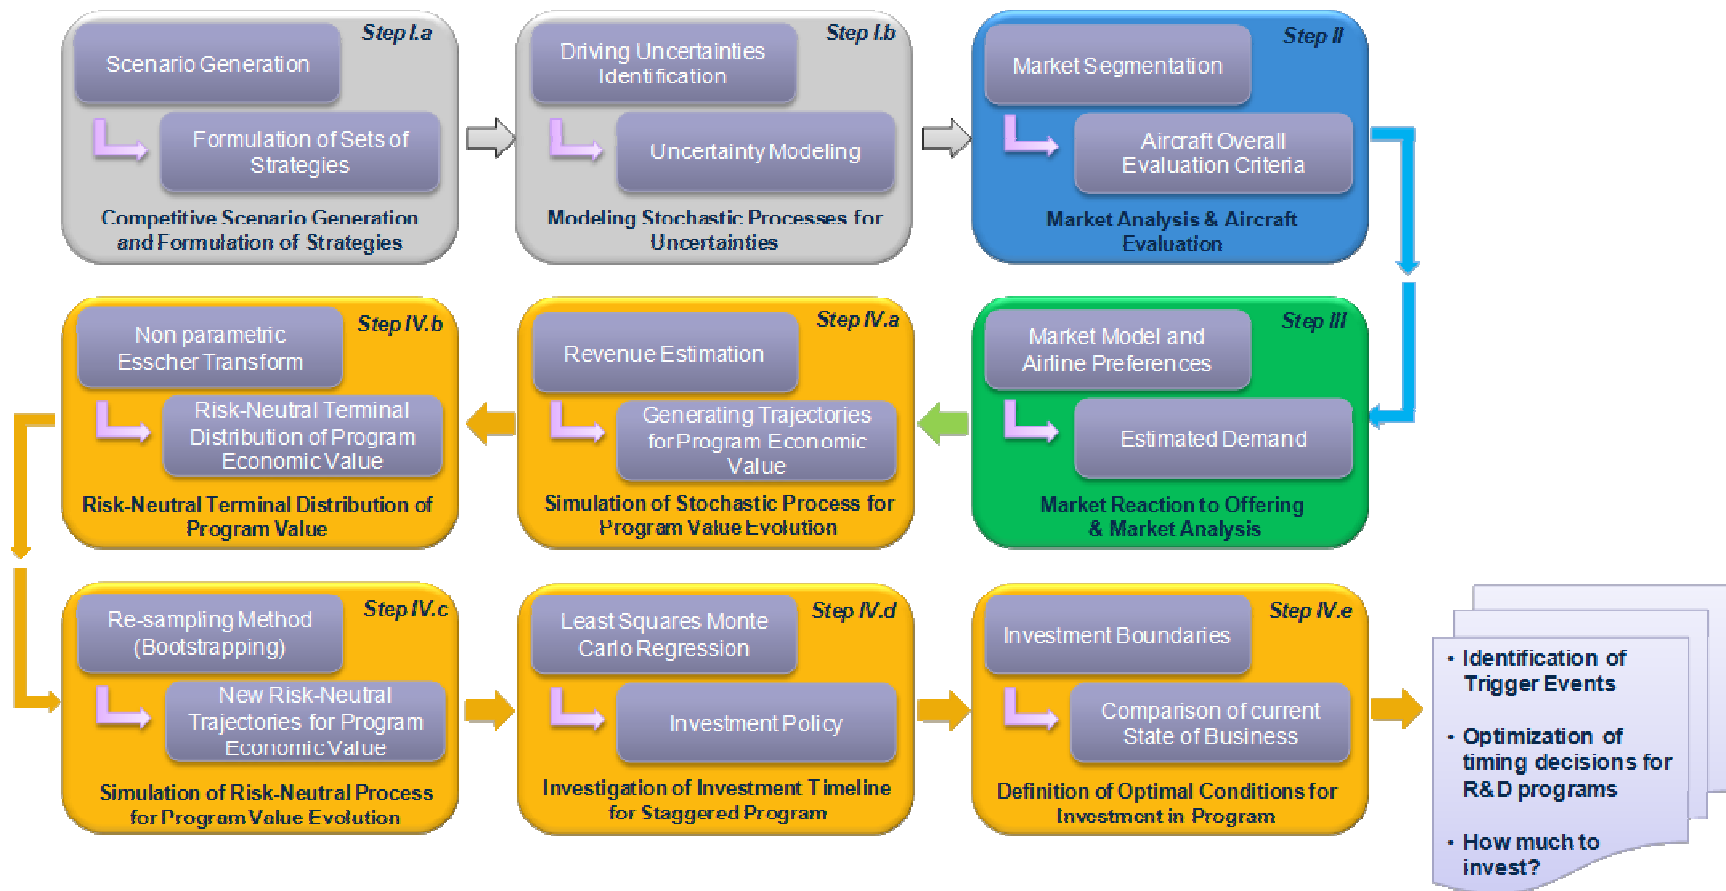


Figure 37: Decomposition of proposed methodology to evaluate R&D business cases

CHAPTER 7: EXPERIMENTAL PLAN - VERIFICATION

Treat with utmost respect your power of forming opinions, for this power alone guards you against making assumptions that are contrary to nature and judgments that overthrow the rule of reason.
Marcus Aurelius, *Meditations*, 3.9

The purpose of the experimental plan is to propose a set of experiments to be carried out to either prove or disprove the hypotheses set forth to answer the research questions formulated in this document. At this stage, it becomes necessary to recognize that there are three types of hypotheses: method hypotheses, modeling hypotheses, and technical hypotheses. Method hypotheses propose a set of ordered procedures to investigate and resolve real-life problems faced by practitioners in the industry. These hypotheses must be validated using an industry relevant problem to ensure they adequately meet the need of practitioners. Modeling hypotheses propose generic mathematical representations of some aspects of real-life. These hypotheses must be verified to ensure that these mathematical representations are correct and that they model all pertinent aspects. Finally, technical hypotheses propose specific mathematical techniques to solve specific mathematical problems. For these hypotheses, a pure mathematical verification is usually sufficient to ensure they properly address and solve the identified problems.

7.1 Preparing the verification for the real options analysis

The purpose of the verification is to check whether the implementation of the real options evaluation methodology yields correct option prices. The similarity between real options and financial options enables the use of financial options to perform the

verification of the option pricer implementation: indeed, the option pricer can evaluate both types indifferently but the necessity of “a context” to price real options, the availability of mathematical models to price financial options, and finally, the prolific literature dealing with the pricing of financial options makes the verification of the latter more straightforward. In fact, real options analyses suffer from the fact that empirical testing is notoriously difficult¹ because of the absence of publicly available data regarding the value of individual research and development programs.

7.1.1 Verification process

The implementation of the real options evaluation methodology proposed in this research is articulated around six successive steps: the Monte Carlo simulation under the physical probability measure, the risk neutralization by means of Esscher transform, the trajectory resampling using bootstrapping under the risk-neutral measure, the least squares regression of conditional expectations, and the early-exercise boundary construction. It is therefore easier to start the verification process by checking that the implementation of each individual step performs adequately in a variety of scenarios before moving on to the verification of the entire implementation.

In this regard, the verification process follows the “bottom-up” approach of the definition-decomposition and verification-validation V-model diagram. The V-model diagram of Forsberg and Mooz [165] is a graphical representation used in systems engineering which depicts the activities related to the development life-cycle of complex systems. Several variants of the V-diagram have been developed over the years [166]

¹ Lenos Trigeorgis during the panel discussion “Real-Options Application: Successes and Impediments” at the 18th International Conference on Real-Options: Theory meets Practice, Medellin, Colombia, July 2014

including the one highlighted in Figure 38 which describes adequately the software development process. The model starts with user needs on the upper left and ends with a user validated system on the upper right. In between, the development process is articulated first in a top-down approach starting with a requirements analysis with increasing granularity as development progresses, followed by the design, and leading to the implementation. The development process follows next a bottom-up approach as higher levels of assemblies and subsystems are successively verified, leading to a system-level verification and finally ending with the actual operation of the system.

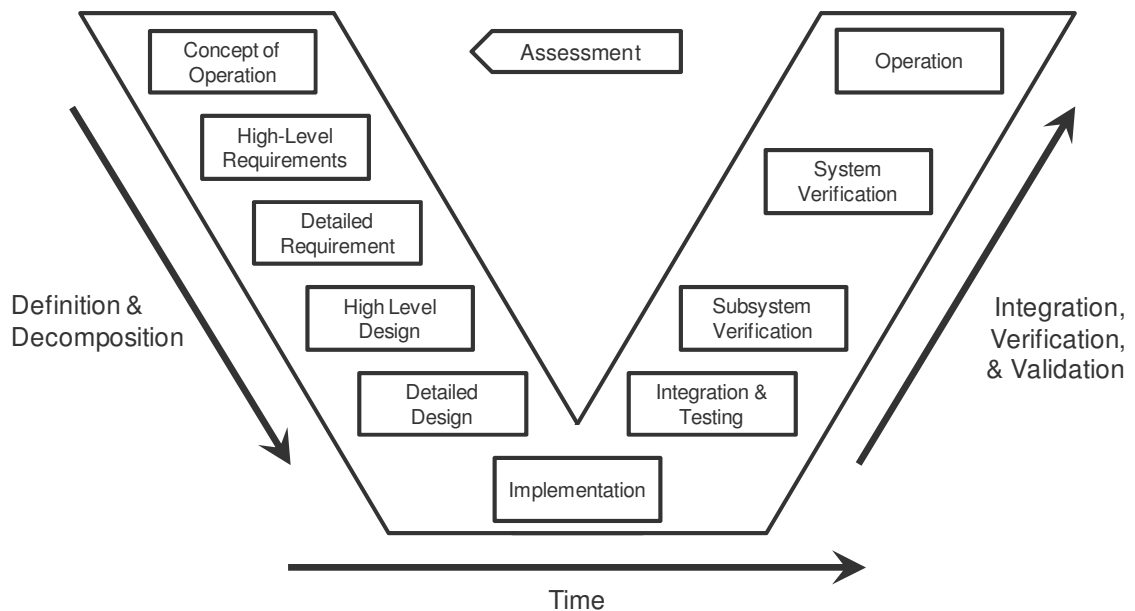


Figure 38: V-Model for systems engineering

In this context, the different steps of the proposed methodology, implemented as modules, are verified independently and a verification capability is thus developed to check their outputs. The verification capability requires different techniques and therefore several testing tools adapted to the module to be checked. Indeed, some modules yield a single number (option price, Esscher parameter), while some modules yield distribution

approximations (risk-neutral distribution), and some others yield two dimensional curved lines (early-exercise boundary). The wide spectrum of tests to be performed can be decomposed into five different categories: visual and graphical methods to check shape of distributions, statistical tests to check properties of distributions, similarity tests to check the shape of curves, numerical comparisons with published results to check quantitative outputs, and finally, numerical comparisons with established techniques to check again quantitative outputs. The verification process is described in Figure 39 with dashed arrows representing verifications of individual modules (subsystem-level) and solid arrows representing verification of the complete implementation (system-level).

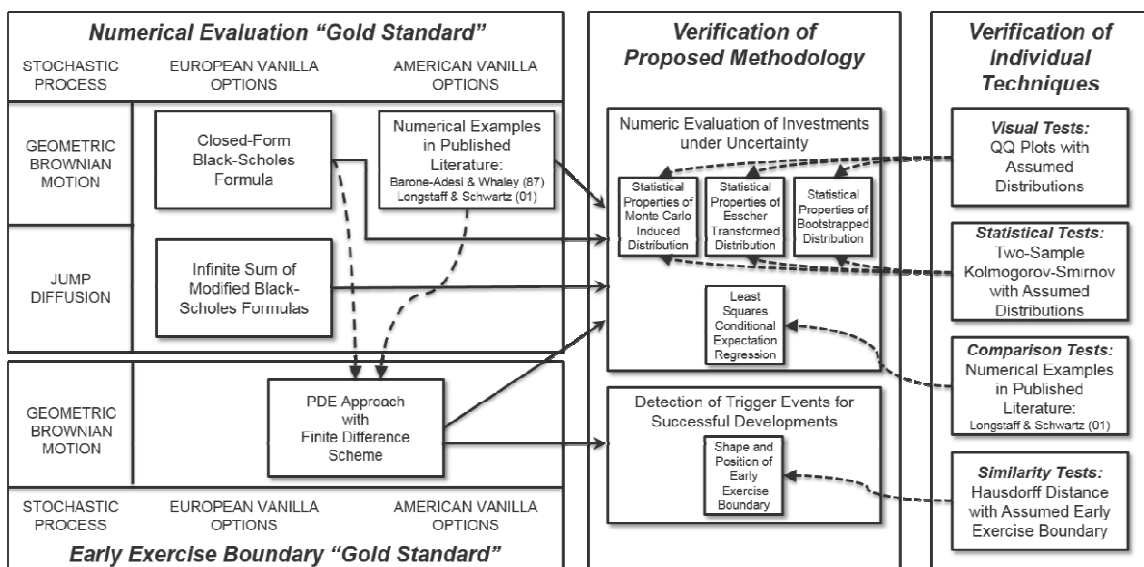


Figure 39: Verification Process

7.1.2 Graphical tests – QQ Plots

One popular technique to visually compare two distributions consists in plotting the quantile of one distribution with respect to the quantile of the other distribution. In such a plot, named a Q-Q plot, the quantiles are values taken at regular intervals from the inverse of the cumulative distribution functions of these distributions. In mathematical

terms, given two cumulative probability distribution functions F and G with associated quantile functions F^{-1} and G^{-1} , the Q-Q plot shows the k^{th} quantile of F against the k^{th} quantile of G for a set of values of k varying between $[0, 1]$. When the two distributions are identical, the Q-Q plot graphs the quantile of one distribution with respect to the quantile of the same distribution which results in a perfectly straight line bisecting the (x, y) plan. It is therefore customary to add this straight bisecting line (following the equation $y=x$) to the Q-Q plot in order to provide a reference for comparisons.

A Q-Q plot presents several characteristics enabling the comparison of distributions. First, this is a visual test providing a graphical representation of how two distributions agree or disagree and enabling a rapid detection of location, scale, dispersion, and skewness differences. Excess dispersion transpires as a plot steeper than the bisecting line. Skewness difference transpires as a plot resembling a curved “S” line. In some cases, location and scale differences can be detected using the intercept and slope of a linear regression between the plotted quantiles. Next, the Q-Q plot is not inherently linked to a specific type of distribution. Provided that the quantiles of the two distributions can be computed, the same Q-Q plot approach can be implemented. This property is used when the proposed methodology is tested for different stochastic processes and therefore different terminal distributions. Finally, another interesting aspect is that Q-Q plots can be used in a non-parametric environment as long as quantiles can be estimated. This is useful if one or both of the distributions do not have any closed-form expression for the inverse of the cumulative probability distribution function. In this case, simulation can be used to generate an empirical distribution function which is then numerically inverted. This property is used during the verification when dealing with

stochastic processes for which the terminal distribution is unknown or difficult to estimate.

One challenge with the implementation of a Q-Q plot is the choice of the plotting positions. It is quite common to use regularly spaced quantiles and a natural choice given a sample of size n is to use k/n with $k=1\dots n$ [167]. However, for probability measures with infinite support, the last quantile which represents the maximum value of the distribution can be infinite. This leads to issues in both plotting and estimating this last quantile. There seems to be little consensus on what is appropriate [168] [169] but $k/(n + 1)$ with $k=1\dots n$ seems to be typically used.

The Q-Q graphical method is implemented in a spreadsheet environment with the help of VBA routines. This enables straightforward communications with the real options evaluation tool which is also implemented in a spreadsheet environment. In the following bullets, some features of the implementation are discussed:

- Q-Q plots are used to compare an empirical distribution function with a known theoretical cumulative probability distribution function. The empirical distribution function results from a sampling algorithm, either Monte Carlo simulation or bootstrapping, and its inverse is never known. As for the theoretical cumulative probability distribution function, its inverse is either computed using published approximations, or simulation is used to generate an empirical distribution function which is then inverted.
- Empirical distribution functions are constructed by sorting outputs from the sampling algorithm in a non-decreasing order (computing the order statistics). Once the ordering is done, the probability associated with each discrete output is


used to estimate the corresponding cumulative probability. For each output, this probability is estimated by adding-up the discrete probabilities associated with all outputs of smaller or equal value.

- Quantiles cannot be computed at regular intervals unless some interpolation is used. To reduce the need for interpolations, quantiles are computed for each and every point of the empirical distribution. The last and “problematic” quantile (the 100% quantile) is replaced by one estimated halfway between the largest two outputs. In other words, if x_{n-1} and x_n are the two largest outputs in a sample of size n , then the n^{th} quantile is given by $F^{-1}\left(\frac{x_{n-1}+x_n}{2}\right)$
- When the inverse of the theoretical cumulative probability distribution function does not have an approximation, simulation is performed to construct an empirical distribution function which is then inverted as explained previously. In this case, the quantiles may need to be interpolated in order to be computed at the same positions as those from the other distribution. A linear interpolation is thus used.


A preliminary “heuristic” verification of the Q-Q plot implementation is provided in Table 23. The intent is to verify that the shapes of these plots are consistent with what is expected. For these tests, two datasets of numbers are generated, their quantiles are estimated, and the corresponding Q-Q plots are graphed. Exhibit (a) displays the Q-Q plot for two datasets from the exact same normal distribution. Exhibit (b) displays the Q-Q plot for two datasets from normal distributions, albeit with a shifted mean. Exhibit (c) displays the Q-Q plot for two datasets from normal distributions, albeit with a different standard deviation. Finally, exhibit (d) displays the Q-Q plot for two different

distributions: one is a normal distribution while the other corresponds to a jump-diffusion stochastic process. The results are as expected: when the simulated and theoretical distributions are identical, the Q-Q plot merges with the bisecting line as shown in (a); when the simulated distribution has a lower mean than the theoretical distribution, the Q-Q plot shift downwards as shown in (b); when the simulated distribution is less dispersed than the theoretical distribution, the Q-Q plot is flatter as shown in (c); and finally when the simulated distribution is platykurtic and positively skewed with respect to the theoretical distribution, the Q-Q plot exhibits a flat “S” shape as in (d).


Table 23: Q-Q plots with 80,000 data points for various simulated and theoretical distributions




(a) Simulated normal (μ, σ) and theoretical normal (μ, σ)
($\mu=3.33e-4, \sigma=1.49e-2$)



(b) Simulated normal (μ, σ) and theoretical normal ($20\mu, \sigma$)
($\mu=3.33e-4, \sigma=1.49e-2$)



(c) Simulated normal (μ, σ) and theoretical normal ($\mu, 2\sigma$)
($\mu=3.33e-4, \sigma=1.49e-2$)



(d) Simulated normal (μ, σ) and theoretical distribution corresponding to jump-diffusion process ($\mu, \sigma, \lambda, \gamma, \delta$)
($\mu=3.33e-4, \sigma=1.49e-2, \lambda=1.0, \gamma=-0.08, \delta=0.4$)

7.1.3 Statistical tests – Kolmogorov-Smirnov test

Another popular technique to test the equality of continuous and one-dimensional probability distributions is the Kolmogorov-Smirnov test. This statistical test may be used to compare a sample with a reference probability distribution or to compare two samples. It quantifies the distance between the empirical distribution function of the sample and the cumulative distribution function of the reference distribution, or between the empirical distribution functions of two samples. If F_n and G represent respectively the empirical distribution function (for the sample of size n) and the cumulative probability distribution function, or if F_n and G_m represent the two empirical distribution functions (for the two samples of sizes n and m), then the distance D_n or $D_{n,m}$ in Eq. 24 is the largest difference between the two functions.

$$\begin{cases} D_n = \sup_{x \in \mathbb{R}} |F_n(x) - G(x)|, & \text{for the one sample case} \\ D_{n,m} = \sup_{x \in \mathbb{R}} |F_n(x) - G_m(x)|, & \text{for the two samples case} \end{cases} \quad \text{Eq. 24}$$

Intuitively, this distance should be small if the two distributions are equal. This distance is used next to compute the Kolmogorov-Smirnov test statistic. The null hypothesis is that the sample is drawn from the reference distribution (in the one-sample case) or that the two samples are drawn from the same distribution (in the two-sample case). If the null hypothesis is true, then the test statistic follows the Kolmogorov-Smirnov distribution. Following traditional hypothesis testing, the null hypothesis is rejected at a significance level α if the Kolmogorov-Smirnov test statistic is greater than the quantity in Eq. 25 with $c(\alpha)$ computed using tables or the Kolmogorov-Smirnov distribution. For large samples, the critical values at the 1% and 5% significance levels are given by $c(0.01) = 1.63$ and $c(0.05) = 1.36$.

$$\begin{cases} KS = D_n \cdot \sqrt{n} > c(\alpha), & \text{for the one sample case} \\ KS = D_{n,m} \cdot \sqrt{\frac{n \cdot m}{n + m}} > c(\alpha), & \text{for the two samples case} \end{cases} \quad \text{Eq. 25}$$

The Kolmogorov-Smirnov statistical test is implemented in a spreadsheet environment with the help of VBA routines which enables straightforward communications with the real options evaluation tool which is also implemented in a spreadsheet environment. In the following bullets, some features of the implementation are discussed:

- When a closed-form expression – or an approximation – of the cumulative probability distribution function of the reference distribution exists, then a one-sample Kolmogorov-Smirnov test is performed.
- When there is no closed-form expression – or approximation – of the cumulative probability distribution function for the reference distribution, then the reference

distribution is constructed using Monte Carlo simulations. A two-sample Kolmogorov-Smirnov test is then performed.

- In some cases, the estimation of the reference cumulative probability distribution function is computationally intensive. In these cases, the sample to be tested is purposefully down-sampled and the resulting empirical distribution function is constructed using the new reduced-size sample. The new down-sampled empirical distribution coincides exactly with the original empirical distribution at each point of the new smaller sample. The Kolmogorov-Smirnov test statistic is computed using a value of n corresponding to the original (larger) sample size so as not to skew the estimated p -value.
- The power of the Kolmogorov-Smirnov test (rejecting the null hypothesis when it is false) increases with the sample size, as suggested by the formulation of the test statistic (square root of sample size factor). Therefore, the Kolmogorov-Smirnov tests are performed on large samples always exceeding 10,000 points, and the critical value retained at a significance level of 5% is given by $c(0.05) = 1.36$
- In order to estimate the p -value associated with these tests, the asymptotic behavior of the statistic is used. It is shown to follow the Kolmogorov-Smirnoff distribution given in Eq. 26. This distribution is expressed as an infinite sum of exponential terms. The contribution of the k^{th} terms is in e^{-2k^2} and thus decreases rapidly. In the numerical implementation, the infinite series is thus approximated by its first thirty terms.

$$P(X \leq x) = 1 - 2 \sum_{k=1}^{\infty} (-1)^{k-1} \cdot e^{-2k^2 x^2} \quad \text{Eq. 26}$$

A preliminary verification of the implementation of the Kolmogorov-Smirnov test is provided in Table 24. The intent is to verify the estimation of the p -values. The implementation is compared to an online calculator for the two-sided two-sample test provided by Wessa [170] and based on the *ks.test* module of the stat package available for the R statistical language. In this study, two datasets of uniformly distributed random numbers are generated and the Kolmogorov-Smirnov test is run to check whether the two datasets are sampled from the same distribution. The first test is performed with two identical datasets ensuring a zero distance. The last test is performed with two shifted datasets ensuring a large distance. All other tests are performed with two distinct datasets sampled from the same uniform distribution. The results are in agreement over the entire range of distances.

Table 24: Kolmogorov-Smirnov statistical test implementation

| | | | | | | |
|-------------------------|------------|-----------|-----------|-----------|-----------|------------|
| Sample size (n, m) | 1000,1,000 | 1000,1000 | 1000,1000 | 1000,1000 | 1000,1000 | 1000,1000 |
| Test distance $D_{n,m}$ | 0 | 0.02700 | 0.02900 | 0.030000 | 0.03600 | 0.11400 |
| Reference p -value | 1 | 0.85929 | 0.79439 | 0.75910 | 0.53605 | 4.5388e-06 |
| Computed p -value | 1 | 0.85929 | 0.79439 | 0.75909 | 0.53605 | 4.5387e-06 |

7.1.4 Statistical tests – Testing the mean using z -tests and t -tests

Unlike the Kolmogorov-Smirnov test which enables the comparison of distributions as a whole, the z -test and the t -test are two popular tests traditionally used to perform statistical inference regarding the mean of a population using the mean of a sample drawn from this population. Both tests are very similar but rely on a slightly different set of assumptions.

When one knows the standard deviation σ of the population from which the sample is drawn, the z -test is used [171]. The null hypothesis for this test is the equality between the population mean and a hypothesized mean. The sample of size n has a mean $\bar{\mu}$ which is standardized to yield the z -statistic defined in Eq. 27. This statistic follows a standard normal distribution under the null hypothesis. Intuitively, the further away the sample mean is from the hypothesized mean, the less likely the null hypothesis is true. Following traditional two-sided hypothesis testing, the null hypothesis is rejected at a significance level α if the absolute value of the z -test statistic is greater than the quantity $z_{\alpha/2}$ computed using tables of the standard normal distribution. The critical values for two-sided tests at the 1% and 5% significance levels are given by $z_{0.005} = 2.575$ and $z_{0.025} = 1.960$.

$$Z = \frac{\bar{\mu} - \mu}{\sigma / \sqrt{n}} \quad \text{Eq. 27}$$

Strictly speaking, the use of the z -test for hypothesis testing requires that the standard deviation of the sample be known. In practice, σ is rarely known and the standard deviation is often replaced by the sample estimate of the standard deviation provided the sample size is large (*large-sample approximation*) [171]. This leads to results (p -value and critical value) that are good, yet approximate. In this case, a large sample size is usually understood to be greater than 30.

When one does not know the standard deviation σ of the population from which the sample is drawn, the t -test is used provided that the population distribution is normal [171]. The null hypothesis for this test is the equality between the population mean and the hypothesized mean. The sample of size n has a mean $\bar{\mu}$ and a standard deviation $\bar{\sigma}$,

which are both used to define the Student's t -statistic defined in Eq. 28. This statistics follows the Student's t -distribution [172] under the null hypothesis. Again, the further away the sample mean is from the hypothesized mean, the less likely the null hypothesis is true. Following traditional two-sided hypothesis testing, the null hypothesis is rejected at a significance level α if the absolute value of the t -test statistic is greater than the critical value $t_{\alpha/2, n-1}$ computed using t -distribution tables.

$$t = \frac{\bar{\mu} - \mu}{s / \sqrt{n}} \quad \text{Eq. 28}$$

Strictly speaking, the use of Student's t -table for hypothesis testing requires that the sample be drawn from a normal distribution. In practice, the t -test yields good yet approximate results (p -value and critical value) even if the normality of the distribution cannot be established, as long as the distribution is symmetric and approximately bell shaped [171]. Probability plots and box-and-whiskers graphs may be used to determine quickly if the distribution satisfies this loose constraint.

In this research, both the z -test and the t -test are used to verify the equality between the average outcome of repeated experiments and the expected theoretical value. Usually, the distributions from which the sample is extracted are unknown and so are the associated standard deviations. Therefore, the verification tests are merely approximations of the z -tests and t -tests using accepted practices. To use both tests, most of the experiments carried out in this research are repeated thirty times to ensure that the large-sample approximation is met and box-and-whiskers plots are used to ensure reasonable symmetry. The z -test and the t -test are implemented in the spreadsheet environment using native Microsoft Excel functions to estimate critical values and p -values.

7.1.5 Similarity tests – Hausdorff distance

The purpose of similarity tests is to check whether the approximation of a curve is close-enough to either the theoretical curve if it is known, or another approximation of the curve obtained through established procedures. Assessing the similarity between two curves seems relatively straightforward when looking at them, but defining an appropriate metric to describe this similarity is more difficult. The field of computer graphics is surveyed in order to find metrics that are used for digital shape recognition and other similar endeavors. It seems that most of the automated shape recognition algorithms revolve around the estimation of a distance between shapes [173]. There are several definitions for this distance and two popular ones in computational geometry are the Hausdorff distance and the Fréchet distance [174] given respectively in Eq. 29 and Eq. 30. The Hausdorff distance δ_H between two curves \mathcal{C}_1 and \mathcal{C}_2 is expressed as the maximum of the two directed Hausdorff distances $\delta_{\mathcal{C}_1, \mathcal{C}_2}$ and $\delta_{\mathcal{C}_2, \mathcal{C}_1}$ (due to the asymmetry of maximin functions), computed using the Euclidian norm. The directed Hausdorff distance is the greatest of all the distances from a point in one curve to the closest point in the other curve. The Fréchet distance δ_F also uses the Euclidian norm but introduces two reparameterizations ϕ and φ of the curves \mathcal{C}_1 and \mathcal{C}_2 respectively. Without loss of generality, let's assume that the reparameterization support is the segment $[0, 1]$. The Fréchet distance is then defined as the infimum over all reparameterizations ϕ and φ of the maximum distance over all $t \in [0, 1]$ measured between $\mathcal{C}_1(\phi(t))$ and $\mathcal{C}_2(\varphi(t))$. A classical interpretation of the Fréchet distance is the minimum length of a leash required to connect a dog and its owner walking two separate paths without backtracking (dubious for a dog).


$$\delta_H = \max(\delta_{C_1, C_2}, \delta_{C_2, C_1}) \text{ with } \delta_{C_1, C_2} = \max_{x \in C_1} \left[\min_{y \in C_2} \|x - y\| \right] \quad \text{Eq. 29}$$

$$\delta_F = \inf_{\phi, \varphi} \left[\max_{t \in [0, 1]} \|C_1(\phi(t)) - C_2(\varphi(t))\| \right] \quad \text{Eq. 30}$$

The Fréchet distance is regarded as a more robust metric but is hard to estimate. On the contrary, the Hausdorff distance is widely used and easier to compute. Alt et al. [175] show that for closed convex curves, the Hausdorff distance equals the Fréchet distance. They also prove that for κ -straight curves, the Fréchet distance between two such curves is bounded by $\kappa + 1$ times their Hausdorff distance. Since this research aims at quantifying the distance between early-exercise boundaries which are continuous and monotonous functions, the κ -straight property is verified (these are curves with increasing chord [176]) and therefore the Fréchet and Hausdorff distance are closely related. For this reason, the simpler Hausdorff distance is retained.

The Hausdorff distance calculation is implemented in a spreadsheet environment with the help of VBA routines which enable direct communications with both the real options evaluation tool and the finite-difference partial differential equation solver to be described next. A preliminary verification of this implementation is provided in Table 25 using a simplified example displayed in the left-most cell. In this example, the Hausdorff distance is computed between a curve C_1 made of two points and a curve C_2 made of three points.

Table 25: Hausdorff distance implementation verification



| | |
|---|--|
| Theoretical Discrete Hausdorff Distance | $\left. \begin{array}{l} \delta_{C_1, C_2} = 1 \\ \delta_{C_2, C_1} = \sqrt{10} \end{array} \right\} \rightarrow \delta_H = \sqrt{10}$ |
| Computed Discrete Hausdorff Distance | 3.162 |

7.1.6 Comparison tests – Solving PDE with finite-difference methods

The purpose of these comparisons is to provide a reference value for real options featuring early-exercise possibilities and to provide a reference for their early-exercise boundaries. Unfortunately, there are only few techniques that can both price American options and locate early-investment boundaries [177] [178]. Solving partial differential equations (and relevant boundary conditions) with a finite-difference scheme is a popular technique achieving these two objectives when the partial differential equations have no known analytical solutions [179].

Finite-difference methods are widely used numerical schemes that enable the pricing of European as well as Bermudan and American options. Finite-difference methods were first proposed by Schwartz [180] and Brennan and Schwartz [181] [182] to solve the Black-Scholes partial differential equation by discretizing the time and asset-price space. Boundary conditions at the extremities of the time and asset-price mesh enable the estimation of the option price which is then propagated throughout the mesh using the finite-difference approximation of the partial differential equation. The boundary conditions are usually set at the maturity of the option, for an extremely large

value of the underlying asset, and for an extremely small value of the underlying asset. Another interesting aspect of solving partial differential equations using a finite-difference scheme is the ability to directly generate the early-exercise boundary. Generating the early-investment boundary is done by checking if the early-exercise privilege is exercised at each and every node in the time and asset-price mesh. The boundary is approximated at each time cross-section by looking at neighboring nodes that have different exercise policies.

The solution of the Black-Scholes partial differential equation using a finite-difference scheme is implemented in a spreadsheet environment with the help of VBA routines which enable direct communications with the real options evaluation tool. In the following bullets, some features of the implementation are discussed:

- There are several ways to express finite-differences in the time and asset-price mesh: forward differences, backward differences, and central differences leading to respectively the explicit, implicit, and Crank-Nicolson finite-difference schemes. Because of potential numerical instabilities with explicit schemes and the extra complexity of Crank-Nicolson schemes, the implicit numerical scheme is used.
- Boundary conditions must be defined at the border of the time and space grid. For option valuation, boundary conditions are usually defined along three boundaries: at the expiration of the option when the payoff is known, whenever the underlying has a value of zero, and whenever the underlying has an infinite value. It is nevertheless impractical to use a boundary at infinity. As a result, it is customary to replace this boundary with an approximate boundary positioned far away from

points of interest, i.e. far away from the near-field where the option valuation will take place in practice. Kangro and Nicolaidis [183] and later Windcliff et al. [184] report that, given a tolerance tol , given a volatility σ , given a maturity T , using a maximum value S_{max} for the underlying asset at least greater than the quantity shown in Eq. 31 is sufficient.

$$S_{max} > Ke^{\sigma\sqrt{2|\ln(tol)|}\sqrt{T}} \text{ often replaced in practice by } S_{max} > Ke^{3\sigma\sqrt{T}} \quad \text{Eq. 31}$$

To compute finite-differences, a discretization of time and space is required. The granularity of this discretization has direct implications on the accuracy of the solutions provided. The proper level of discretization is the one that yields the target accuracy while using the minimum number of mesh-points. APPENDIX I discusses the discretization choice and the final discretization is summarized in Table 26.

Table 26: Grid selection for finite-difference numerical scheme

| | |
|--------------------------------|-----------|
| Space dimension discretization | 500 Steps |
| Time dimension discretization | 400 Steps |

The verification of this implementation is articulated around two steps. First, the finite-difference scheme is used to price European options for which exact analytical solutions are known. The verification is performed for both put and call options, for different spot to strike ratios, different volatilities, different risk-free interest rates, different dividend yields, and different maturities as highlighted in Table 27. A list of 240 test cases is provided in APPENDIX J in a table format while a more synthetic view representing the distribution of the relative difference is provided in Figure 40.

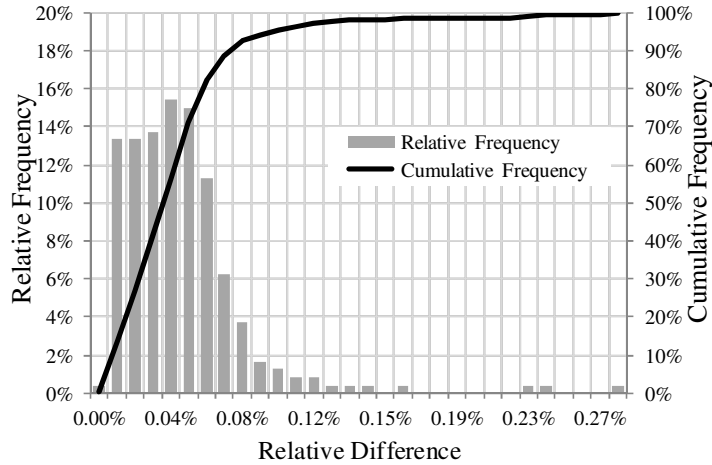


Figure 40: Relative difference distribution for various test cases (finite-difference scheme compared to Black-Scholes solution)

| | Min | Max |
|-----------------|-----|-----|
| Spot to strike | 0.8 | 1.2 |
| Maturity (days) | 180 | 720 |
| Volatility | 20% | 40% |
| Dividend yield | 0% | 4% |
| Risk free rate | 2% | 8% |

Table 27: Range for input parameters

The finite-difference scheme is used next to price American options for which the verification task is more involved since there is no known practical analytical solution. The option prices for different spot prices, different risk-free rates, different volatilities, and different maturities are therefore compared to results published in the literature. Results reported by Barone-Adesi and Whaley [185] are used and comparisons are provided in Table 28.

The results for both European and American options are accurate across the whole spectrum of test-cases. For European options, the relative difference does not exceed a tenth of a percent in ninety five percent of cases. For American options, the magnitude of the absolute error remains minuscule, never exceeding one percent of the option price. Deep out-of-the-money option price calculations seem to be less accurate but these options are of little interest in this research. The implementation of the finite-difference numerical scheme is therefore successful.

Table 28: Finite-difference scheme results for American call and put options

| Comparison ^(a) | | American Call Option | | | American Put Option | | |
|--|------------|--------------------------------------|--------------------------|-------|--------------------------------------|--------------------------|-------|
| Option Parameters ^(b) | Spot Price | Barone-Adesi & Whaley ^(c) | Finite-difference Method | Diff. | Barone-Adesi & Whaley ^(c) | Finite-difference Method | Diff. |
| K=100, r=0.08, $\sigma=0.20$, T=0.25 | 80 | 0.05 | 0.05 | 0.00 | 20.00 | 20 | 0.00 |
| | 90 | 0.85 | 0.85 | 0.00 | 10.22 | 10.223 | 0.00 |
| | 100 | 4.44 | 4.44 | 0.00 | 3.55 | 3.547 | 0.00 |
| | 110 | 11.66 | 11.66 | 0.00 | 0.79 | 0.79 | 0.00 |
| | 120 | 20.90 | 20.90 | 0.00 | 0.11 | 0.114 | 0.00 |
| K=100, r=0.12, $\sigma=0.20$, T=0.25 ^(d) | 80 | 0.07 | 0.07 | 0.00 | 20.00 | 20.00 | 0.00 |
| | 90 | 1.02 | 1.02 | 0.00 | 10.01 | 10.03 | 0.02 |
| | 100 | 4.96 | 4.97 | 0.01 | 3.21 | 3.21 | 0.00 |
| | 110 | 12.50 | 12.50 | 0.00 | 0.68 | 0.66 | -0.02 |
| | 120 | 21.85 | 21.85 | 0.00 | 0.10 | 0.09 | -0.01 |
| K=100, r=0.08, $\sigma=0.40$, T=0.25 | 80 | 1.29 | 1.29 | 0.00 | 20.59 | 20.59 | 0.00 |
| | 90 | 3.82 | 3.82 | 0.00 | 12.95 | 12.96 | 0.01 |
| | 100 | 8.35 | 8.35 | 0.00 | 7.46 | 7.46 | 0.00 |
| | 110 | 14.79 | 14.79 | 0.00 | 3.95 | 3.95 | 0.00 |
| | 120 | 22.71 | 22.71 | 0.00 | 1.94 | 1.94 | 0.00 |
| K=100, r=0.08, $\sigma=0.20$, T=0.50 | 80 | 0.41 | 0.42 | 0.01 | 20.00 | 20.00 | 0.00 |
| | 90 | 2.18 | 2.18 | 0.00 | 10.75 | 10.76 | 0.01 |
| | 100 | 6.50 | 6.49 | -0.01 | 4.77 | 4.77 | 0.00 |
| | 110 | 13.42 | 13.42 | 0.00 | 1.74 | 1.74 | 0.00 |
| | 120 | 22.06 | 22.06 | 0.00 | 0.53 | 0.53 | 0.00 |

^(a) Comparison is made with commodity option prices for which a cost of carry $q = 0.04$ is used
^(b) K = strike; r = riskless rate of interest; σ = standard deviation of returns; T = time to expiration
^(c) Barone-Adesi and Whaley implementation of the finite-difference method [185].
^(d) The author wishes to thank Giovanni Barone-Adesi for providing corrected numbers for this specific case

7.2 Preliminary testing and lessons learned

Preliminary testing is performed prior to entering the detailed verification process for two reasons: the first is to set values for some of the technical parameters used in the proposed methodology; the second is to iron out glitches and determine whether adjustments or improvements to the proposed methodology are warranted.

7.2.1 Variability of results

Preliminary testing indicates that option prices are quite accurate (usually within 5%) but the early-exercise boundary seems to exhibit some quite severe changes of shape

in between repeated experiments. These changes in the shape of the early-exercise boundary in between repeated experiments are due to the variability of the continuation value estimation at intermediate steps in the algorithm. There is therefore a need to decrease this variability and this leads to a new and rather unexpected research question:

Research Question 1.2.2 – Reducing variability of results


How can the variance of results obtained from the least-squares Monte Carlo simulation be reduced to yield consistent real option price estimations and consistent early-exercise boundary shapes?

7.2.2 Pooling sample of returns before bootstrap resampling


Testing is first performed to check the bootstrap resampling using the sampling wheel algorithm. The purpose of the bootstrap resampling is to generate new trajectories from a sample of returns. These trajectories induce empirical distributions of returns at each time step which can be compared to known reference distributions of returns. Preliminary tests using Q-Q plots comparing these distributions indicate that some distortion occurs, especially in the tails of the distributions. When the stochastic process under the physical probability measure is significantly different from the process under the equivalent martingale measure (i.e. the drift rates are significantly different), the non-parametric Esscher transform must heavily tilt the empirical distribution of returns. It is achieved by applying non-uniformly distributed weights to each of the returns within each sample of returns. This means that some of the returns have a large weight attached to them while some others have a small weight attached to them. During the bootstrap resampling, returns with higher weights are more likely to be drawn. Therefore, one

specific return or several returns with relatively larger weights may be drawn over and over again. This translates into the step-like structure apparent in the lower tails of the Q-Q plots in exhibit (b), exhibit (c), and exhibit (d) of Table 29. The lower tails of the empirical distribution under the equivalent martingale measure were indeed heavily reweighted so as to tilt the distribution and achieve a lower drift rate.


Table 29: Q-Q plots for the return distribution induced by bootstrap resampling of two cases of geometric Brownian motions Merton jump diffusion processes




(a) GBM with $r_f=0.05$, $\mu=0.07$, $q=0.02$, $\sigma=0.10$, $T=1$



(b) GBM with $r_f=0.05$, $\mu=0.25$, $q=0.02$, $\sigma=0.30$, $T=1$



(c) JD with $r_f=0.05$, $\mu=0.07$, $q=0.02$, $\sigma=0.10$, $\lambda=4.00$, $\gamma=-0.08$, $\delta=0.40$, $T=1$



(d) JD with $r_f=0.05$, $\mu=0.07$, $q=0.02$, $\sigma=0.10$, $\lambda=4.00$, $\gamma=-0.08$, $\delta=0.40$, $T=1$


To ensure that a specific return is not drawn repetitively, two solutions are available: ensure that no return has a relative weight so large that it is bound to be drawn over and over again; or ensure that many returns have large relative weights therefore diluting the possibility of drawing one particular return over and over again. Implementing the second solution seems relatively straightforward: in order to avoid repetitive sampling of the same return, sampling m returns out of n returns with m significantly smaller than n ensures that the likelihood of repetitive sampling is reduced. In practice, this is implemented by first pooling return samples from different time cross-sections, then by performing the change of measure with the non-parametric Esscher transform over this larger sample, and finally by sampling from this larger risk-neutralized sample of returns. This leads to a new hypothesis formulated below:

Hypothesis 1.1.3.1 – Pooling returns to increase size of sample to bootstrap


Pooling samples of returns from different time cross-sections or increasing the relative size from the original sample with respect to the bootstrap sample limits the repetitive sampling of the same highly-weighted return values.

For the same test case, Table 30 highlights the improvement in the quality of the resampled distribution as the number of pooled samples is increased from a single sample to six samples. It seems that most artifacts of the resampling disappear when four or more samples of returns are pooled together. Therefore, the resampling ratio is set to four for the remaining of this dissertation.


Table 30: Q-Q plots with increasing pooling number




(a) No pooling - Single sample




(b) Pooling - Two samples




(c) Pooling - Three samples



(d) Pooling - Four samples



(e) Pooling - Five samples



(f) Pooling - Six samples

Another technique to achieve the same result is to perform the bootstrap resampling using a larger original sample: instead of sampling n returns from an original

sample of observations of size n , the sampling can also be done using an original sample of size m with m much larger than n . Using the conclusion above, selecting an original sample of size $4n$ seems sufficient to remove artifacts in the new distribution. This has same effective results but will prove useful later when the case of compound options is investigated: indeed, the stationarity assumption invoked to justify the bootstrap resampling is no longer valid for the nested option price process and consequently bootstrapping from several different time cross-sections would prove problematic.

In any case, the two techniques have similar intent and similar results. The benefits are highlighted in Table 30 and a heuristic argumentation is invoked to both justify that resampling from a larger sample leads to fewer repetitions and to verify this hypothesis.

7.2.3 Antithetic variates, moment matching, and control variates

Testing is performed next to check the accuracy of the valuation of American options. The purpose is to quickly check whether the least-squares Monte Carlo algorithm can be used in the proposed methodology without any modification or improvements. Preliminary tests are performed for different American call options on underlying assets following a geometric Brownian motion. Option prices, confidence intervals, as well as early-exercise premia are reported in Table 31. These tests indicate that option prices are reasonably accurate but there is significant variability in the prices obtained from repeated experiments. This yields early-exercise premia (i.e. difference between the American option price and the corresponding European option price) which are not statistically significant in some deep out-of-the money options. Besides, the variability yields wide confidence intervals, particularly for deep out-of-the money options: the relative width of some confidence intervals exceeds seven percent of the

option price. In order to curb some of this variability, several variance reduction techniques have been developed over the years including the use of antithetic variates, moment matching, and control variates which are discussed in the following paragraphs.

Table 31: American call option price and early-exercise premium for underlying assets following a geometric Brownian motion

| Test Case | Finite-Difference Method | | | Proposed Simulation Method | | | | | |
|---|--------------------------|-----------------------|------------------------|----------------------------|----------------|----------------|------------------------|---------------------|-------------------------------|
| | European Option Price | American Option Price | Early-Exercise Premium | American Option Price | Standard Error | Relative Error | Early-Exercise Premium | Confidence Interval | Conf. Interval Relative Width |
| S=0.8 $\mu=0.1$ q=0.05 $\sigma=0.2$ | 0.00843 | 0.00867 | 0.00024 | 0.00864 | 0.00016 | -0.3% | 0.00021 | 0.0083-0.0089 | 7.0% |
| | 0.00843 | 0.00867 | 0.00024 | 0.00886 | 0.00016 | 2.2% | 0.00043* | 0.0086-0.0092 | 7.1% |
| | 0.00843 | 0.00867 | 0.00024 | 0.00861 | 0.00016 | -0.7% | 0.00018 | 0.0083-0.0089 | 7.2% |
| | 0.00843 | 0.00867 | 0.00024 | 0.00837 | 0.00015 | -3.4% | -0.00006 | 0.0081-0.0087 | 6.8% |
| | 0.00843 | 0.00867 | 0.00024 | 0.00869 | 0.00016 | 0.2% | 0.00026 | 0.0084-0.0090 | 7.1% |
| S=0.9 $\mu=0.1$ q=0.05 $\sigma=0.4$ | 0.09163 | 0.09365 | 0.00202 | 0.09267 | 0.00051 | -1.1% | 0.00104* | 0.0917-0.0937 | 2.1% |
| | 0.09163 | 0.09365 | 0.00202 | 0.09345 | 0.00048 | -0.2% | 0.00183* | 0.0925-0.0944 | 2.0% |
| | 0.09163 | 0.09365 | 0.00202 | 0.09371 | 0.00048 | 0.1% | 0.00208* | 0.0928-0.0946 | 2.0% |
| | 0.09163 | 0.09365 | 0.00202 | 0.09375 | 0.00050 | 0.1% | 0.00212* | 0.0928-0.0947 | 2.1% |
| | 0.09163 | 0.09365 | 0.00202 | 0.09403 | 0.00051 | 0.4% | 0.00240* | 0.0930-0.0950 | 2.1% |
| S=1.0 $\mu=0.1$ q=0.05 $\sigma=0.2$ | 0.06330 | 0.06655 | 0.00324 | 0.06729 | 0.00027 | 1.1% | 0.00399* | 0.0668-0.0678 | 1.6% |
| | 0.06330 | 0.06655 | 0.00324 | 0.06651 | 0.00026 | -0.1% | 0.00321* | 0.0660-0.0670 | 1.6% |
| | 0.06330 | 0.06655 | 0.00324 | 0.06652 | 0.00026 | 0.0% | 0.00322* | 0.0660-0.0670 | 1.5% |
| | 0.06330 | 0.06655 | 0.00324 | 0.06790 | 0.00027 | 2.0% | 0.00459* | 0.0674-0.0684 | 1.6% |
| | 0.06330 | 0.06655 | 0.00324 | 0.06731 | 0.00026 | 1.1% | 0.00401* | 0.0668-0.0678 | 1.5% |
| S=1.1 $\mu=0.1$ q=0.05 $\sigma=0.4$ | 0.19581 | 0.20170 | 0.00589 | 0.20043 | 0.00051 | -0.6% | 0.00462* | 0.1994-0.2014 | 1.0% |
| | 0.19581 | 0.20170 | 0.00589 | 0.20117 | 0.00052 | -0.3% | 0.00536* | 0.2001-0.2022 | 1.0% |
| | 0.19581 | 0.20170 | 0.00589 | 0.19945 | 0.00052 | -1.1% | 0.00364* | 0.1984-0.2005 | 1.0% |
| | 0.19581 | 0.20170 | 0.00589 | 0.20185 | 0.00054 | 0.1% | 0.00604* | 0.2008-0.2029 | 1.0% |
| | 0.19581 | 0.20170 | 0.00589 | 0.20019 | 0.00052 | -0.7% | 0.00439* | 0.1992-0.2012 | 1.0% |
| S=1.2 $\mu=0.1$ q=0.05 $\sigma=0.2$ | 0.18839 | 0.20502 | 0.01663 | 0.20555 | 0.00021 | 0.3% | 0.01716* | 0.2051-0.2060 | 0.4% |
| | 0.18839 | 0.20502 | 0.01663 | 0.20387 | 0.00022 | -0.6% | 0.01548* | 0.2034-0.2043 | 0.4% |
| | 0.18839 | 0.20502 | 0.01663 | 0.20510 | 0.00020 | 0.0% | 0.01671* | 0.2047-0.2055 | 0.4% |
| | 0.18839 | 0.20502 | 0.01663 | 0.20500 | 0.00020 | 0.0% | 0.01661* | 0.2046-0.2054 | 0.4% |
| | 0.18839 | 0.20502 | 0.01663 | 0.20468 | 0.00021 | -0.2% | 0.01628* | 0.2043-0.2051 | 0.4% |
| <p>Call option price on an asset following a geometric Brownian motion. Simulation performed with 30,000 original trajectories, and 30,000 resampled trajectories.</p> <p>S = asset price, μ = drift rate, q = dividend yield, σ = volatility, K = strike price = 1, r_f = riskless rate of interest = 2%, T = maturity = 1 year Early-exercise premia with asterisks (*) denote values significantly different from zero</p> | | | | | | | | | |

Antithetic variates

The antithetic variates concept attempts to reduce variance by observing that if a random variable U is uniformly distributed over the unit interval, then the random variable $1-U$ is also uniformly distributed on the unit interval [131]. When performing Monte Carlo simulations, it is customary to use the inverse transform method whereby a random variable Z having probability distribution F is simulated by first sampling a uniform distribution and then by applying the inverse transform F^{-1} . Therefore, if trajectories are constructed from the random variable U and the inverse function F^{-1} , then another trajectory can be constructed with the random variable $1-U$ and the same inverse function F^{-1} . The two are antithetic in that an abnormally large (small) value of U will be immediately counterbalanced with an abnormally small (large) value of $1-U$. In mathematical terms, the antithetic variates method is useful if, for a random variate Z with independent observations Z_i and antithetic observations Z_i^A , the inequality in Eq. 32 holds:

$$\text{Var}\left(\frac{Z_1 + Z_1^A}{2}\right) < \text{Var}\left(\frac{Z_1 + Z_2}{2}\right) \quad \text{Eq. 32}$$

Which by virtue of the independence of the observations Z_i yields Eq. 33:

$$\text{Var}(Z_1 + Z_1^A) < 2\text{Var}(Z_1) \quad \text{Eq. 33}$$

Decomposing the variance yields Eq. 34:

$$\begin{aligned} \text{Var}(Z_1 + Z_1^A) &= \text{Var}(Z_1) + \text{Var}(Z_1^A) + 2\text{Cov}(Z_1, Z_1^A) \\ \text{Var}(Z_1 + Z_1^A) &= 2 \text{Var}(Z_1) + 2\text{Cov}(Z_1, Z_1^A) \end{aligned} \quad \text{Eq. 34}$$

And the useful condition is expressed as shown in Eq. 35:

$$\text{Cov}(Z_1, Z_1^A) < 0 \Leftrightarrow \text{Cov}(F^{-1}(u), F^{-1}(1-u)) < 0 \quad \text{Eq. 35}$$

In other words, if the negative dependence between the inputs produces negative correlation between the outputs, then the antithetic variates method will reduce the variance of the estimator to be found. The antithetic variates method is implemented in the proposed methodology and yields improvement in the reduction of the variance of the option price. This is however not sufficient. Besides, antithetic variates do not usually work well with quasi-Monte Carlo simulations and this might become problematic if quasi-random numbers are used in lieu of pseudo-random numbers.

Moment matching

Starting from the observation that derivative pricing consists in determining the value of a derivative with respect to the value of the underlying asset, the moment matching method aims at reducing the variance of the price estimator by improving the simulation of the underlying asset evolution. This is achieved by ensuring that experimental samples produced at various time cross-sections exhibit accurate statistical moments and if not, by adjusting trajectories so that the empirical moments perfectly match their theoretical values. Indeed, if $Z_{i=1..n}$ are independent random variables used to drive the simulations, there is little chance that the sample moments resulting from the sampling of the random variables $Z_{i=1..n}$ at each time step exactly matches the expected value \bar{Z} . Focusing only on the first moment, let's introduce \hat{Z} as the experimental first moment (mean) as shown in Eq. 36 (time dependency has been omitted for clarity):

$$\hat{Z} = \frac{1}{n} \sum_{i=1}^n Z_i \quad \text{Eq. 36}$$

Trajectories can then be corrected at each time step by adjusting the realization of the random variable Z_i and creating a new random variable \tilde{Z}_i as shown in Eq. 37.

$$\tilde{Z}_i = Z_i - (\hat{Z} - \bar{Z}) \quad \text{Eq. 37}$$

The first moment is now matched as highlighted in Eq. 38:

$$E(\tilde{Z}_i) = E(Z_i - (\hat{Z} - \bar{Z})) = \bar{Z} - E(Z_i) + \bar{Z} = \bar{Z} \quad \text{Eq. 38}$$

In the context of this research, the moment matching technique is used differently in order to accommodate the specificity of the problem to be solved. First, it is applied to the underlying asset return instead of being applied to the underlying asset price: using the stationarity assumption, it is indeed easier to track a constant expected return than an ever-changing expected price. Second, the moment matching is not used during the initial Monte Carlo simulation but rather during the bootstrap resampling performed according to the sampling wheel algorithm. Indeed, the underlying business prospect is subject to many uncertainties, each following a specific and possibly correlated stochastic process. This means that the stochastic process driving the evolution of the business prospect value is unknown and the expected value of its return process is also unknown. Still, there is hope: when the bootstrap resampling is performed, the pooling of several samples of returns under the equivalent martingale measure (as described previously in section 7.2.2) yields a very large sample of returns. Owing to the large size of the pooled sample, the estimation of the sample average return under the equivalent martingale measure is assumed to be accurate. This estimator is henceforth named the *reference return*. The stationarity assumption is then invoked to require that this *reference return* be matched at each time step of the resampling process. The match is achieved using the moment matching technique just described.

Control variates

Another technique to reduce the variance of estimates obtained with Monte Carlo simulations and therefore to reduce the computational effort is the use of control variates. Control variates exploit errors in the estimates of known quantities to reduce the error in estimates of an unknown quantity. For instance, it is customary to use the price of European options as control variate during the pricing of American options: in this case, the European option price is estimated using the same set of trajectories as those used for the pricing of the American option and the European option price estimate is compared to the corresponding closed-form solution to compute the estimation error. This error is used next to correct the American option price estimate.

Following Glasserman [131], if the objective is to estimate the average of n discounted payoffs denoted $X_{i=1..n}$ using Monte Carlo simulations, then another output of the simulation with known expected value can be tracked. For instance, the discounted underlying asset price denoted $S_{i=1..n}$ is a martingale under the equivalent martingale measure and therefore its expected value is its current value. Therefore, for any fixed value of b , the quantity in Eq. 39 can be estimated for the n trajectories in the simulation (i denotes the trajectory index).

$$X_i(b) = X_i - b(S_i - E(S)) \quad \text{Eq. 39}$$

Averaging over all trajectories, the sample average of the new random variable called the control variate estimator $\hat{X}(b)$, is given in Eq. 40 as a function of the sample mean discounted payoff \hat{X} and the sample mean discounted underlying asset price \hat{S} :

$$\hat{X}(b) = \hat{X} - b(\hat{S} - E(S)) = \frac{1}{n} \sum_{i=1}^n (X_i - b(S_i - E(S))) \quad \text{Eq. 40}$$

The control variate estimator is unbiased as shown in Eq. 41:

$$E(\hat{X}(b)) = E(\hat{X} - b(\hat{S} - E(S))) = E(\hat{X}) = E(X) \quad \text{Eq. 41}$$

And the variance of the control variate is given in Eq. 42:

$$\text{var}(\hat{X}(b)) = \frac{1}{n}(\text{var}(X) - 2b \cdot \text{cov}(X, S) + b^2 \cdot \text{var}(S)) \quad \text{Eq. 42}$$

The purpose of the control variate technique is to perform a reduction in variance of the estimator to be computed. Minimizing the variance of the estimator $\hat{X}(b)$ consists in minimizing the above quadratic function in b . It is achieved for an optimal value of b , identified as b^* , and leads to the variance reduction factor given in Eq. 43:

$$b^* = \frac{\text{cov}(X, S)}{\text{var}(S)} \quad \text{and} \quad \frac{\text{var}(X)}{\text{var}(X - b^*(S - E(S)))} = \frac{1}{1 - \rho_{X,S}^2} \quad \text{Eq. 43}$$

The variance of X , the variance of S , and the covariance between X and S are usually unknown and only an estimate \hat{b}^* of b^* can be used. The estimate \hat{b}^* is computed by replacing the variance and covariance with their sample estimates as shown in Eq. 44:

$$\hat{b}^* = \frac{\sum_{i=1}^n (S_i - \hat{S})(X_i - \hat{X})}{\sum_{i=1}^n (S_i - \hat{S})^2} \quad \text{Eq. 44}$$

With the variance reduction ratio given in Eq. 43, there are several comments that can be made. First, the higher the correlation between the quantity to be estimated and the control variate, the higher the variance is reduced and therefore the more efficient the technique is. The high correlation between European option prices and American option prices explains why European options are often used as control variates for the pricing of American options.

Second, the variance reduction increases sharply with the correlation between the control variate and the quantity of interest as shown in Figure 41. Thus, for a control variate to be efficient, the correlation must be high. Otherwise, the extra complexity introduced by the tracking of control variates outweighs the benefits.




Figure 41: Variance reduction factor versus of control variate correlation

Let's now try to understand the source of the variance of the control variate corrected estimator of an American option price. For this, let's introduce the stopping time τ , the maturity of the option denoted by T , and let's decompose the discounted option payoff as shown in Eq. 45. The variance of the first term on the right hand side is linked to the time interval between expiration and exercise ($T - \tau$). The variance of the second term depends on how close the control variate is to the discounted payoff. The variance of the last term is null. All in all, this means that the variance of the control variate corrected estimator of the American option price increases when the time between expiration and exercise is long, i.e. for in-the-money and deep in-the-money options when exercise occurs early during the life of the option.

$$X_\tau - b^*(S_T - E(S_T)) = (X_\tau - b^* \cdot X_T) + b^* \cdot (X_T - S_T) + b^* E(S_T) \quad \text{Eq. 45}$$

Control variate sampled at exercise

The previous observation highlights one lingering issue with the use of control variates for pricing American options: discounted payoffs used for the pricing of the

option do not always occur at expiration of the option while control variates are usually sampled at expiration. In fact, the timing discrepancy between the sampling of the discounted payoffs used for option pricing and the sampling of control variate reduces the correlation between these two quantities and thus the efficiency of the control variate technique. To improve this, Rasmussen [186] suggests a different sampling scheme for the control variates: instead of sampling the control variates at maturity, the control variates are sampled for each and every simulation trajectory individually at the time of exercise of the American option so as to maximize correlation with the American option discounted payoffs. The new sampling process for the control variate is highlighted in Figure 42.




Figure 42: Sampling control variates at maturity (left graph) is less correlated with option payoffs than sampling control variates at exercise (right graph)

Under this new sampling scheme, the control variate corrected discounted option payoff can be decomposed as shown in Eq. 46. According to this equation, the variance of the control variate corrected estimator of the payoff no longer depends on the time interval between early-exercise of the option and expiration of the option, but only on the time to exercise τ .

$$X_\tau - b^*(S_\tau - E(S_\tau)) = (X_\tau - b^* \cdot X_\tau) + b^* \cdot (X_\tau - S_\tau) + b^*E(S_\tau) \quad \text{Eq. 46}$$

Control variate improved regressions

To further reduce the variance of the American option price, Rasmussen [187] propose another improvement which consists in reducing the variance of the conditional expectation regressions of the continuation value. Similarly to why control variates reduce the variance of Monte Carlo estimations, the control variate improved regressions replace the regression of the discounted payoffs sampled at exercise by the regression of a random variable with same conditional expectation but smaller conditional variance. This helps produce a more efficient estimator of the conditional expectation of the continuation value and therefore a more accurate definition of the early-exercise policy. Indeed, Rasmussen [188] states that *“if there is correlation in the discrepancies between the projection estimates and the true conditional values of [both] the continuation value and the control variate, [...] the latter [can be used] to improve the former”*. In other words, the error between the regression and the conditional expectation of a known quantity is used to improve the regression of the unknown continuation value.

As usual, Rasmussen selects European option prices as control variates for several reasons. First, there is a high correlation between European and American option payoffs. Then, the European option price process is a martingale under the equivalent martingale measure. Finally, a closed-form analytical formula is available to estimate European option price under certain assumptions. Unfortunately, the stochastic process governing the evolution of the underlying business prospect value is unknown in the current research and therefore European option prices cannot be used as control variate. Instead, the discounted business prospect value $\{S_t\}_{0 \leq t \leq T}$ is suggested in this research as

control variate to improve the regressions. Indeed, it is strongly correlated with the American option payoffs for in-the-money options for which an optimal early-exercise policy is important. Therefore, the discounted business prospect value is a good approximation of the conditional expectation of the continuation value. Besides, it is a martingale under the equivalent martingale measure and its conditional expectation is known at each and every time step. Finally, its evolution over time is already simulated during the Monte Carlo simulations which alleviate some of the computational complexity of introducing control variates. The control variate improved regression technique is illustrated in Figure 43.

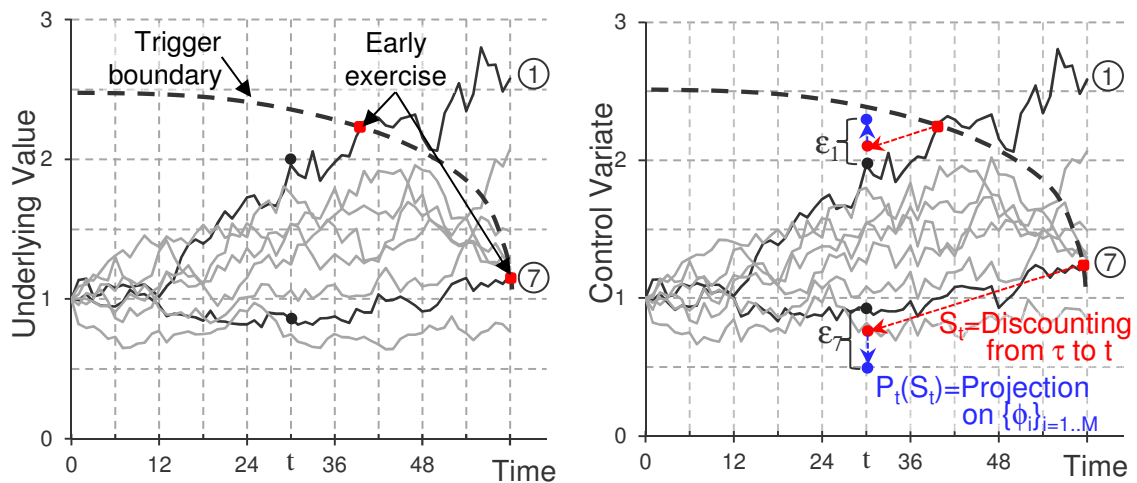


Figure 43: Illustration of the control variate improved regressions on two trajectories. The discounted underlying is used as control variate, sampled at exercise of the option, and projected onto the set of basis-functions.

In mathematical terms, let's now denote by τ the stopping time consistent with the early-exercise policy, by S_τ the control variate sampled at the stopping time τ (i.e. the discounted business prospect value sampled at the stopping time), by X_τ the discounted payoff sampled at the stopping time τ , by Z_τ the control variate corrected discounted

payoff sampled at the stopping time τ , and by b_t a time-indexed control variate correction factor. Similarly to what was expressed in Eq. 47, the control variate corrected discounted payoff is given in Eq. 47:

$$Z_\tau = X_\tau - b_t(S_\tau - E_t(S_\tau)) \quad \text{Eq. 47}$$

Let's now introduce P_t as the projection (or regression) of Eq. 47 onto a set of basis-functions. This is the same set of basis-functions as the one used in the regressions of the conditional expectation continuation value in the least-squares Monte Carlo algorithm. This leads to Eq. 48:

$$P_t(Z_\tau) = P_t(X_\tau) - b_t(P_t(S_\tau) - P_t(E_t(S_\tau))) \quad \text{Eq. 48}$$

The discounted business prospect value is a martingale under the equivalent martingale measure. Invoking the optional stopping theorem [189] at the stopping time τ yields $E_t(S_\tau) = S_t$ which leads to Eq. 49. The first term on the right of Eq. 49 is the usual projection of the discounted payoffs onto the set of basis-functions, while the second term is a correction factor for the error between the projection of the discounted control variate sampled at exercise and the conditional expectation.

$$P_t(Z_\tau) = P_t(X_\tau) - b_t(P_t(S_\tau) - S_t) \quad \text{Eq. 49}$$

The optimal factor b_t^* was previously introduced as the optimal ratio of covariance and variance that minimizes the variability of Monte Carlo estimators. In this improved setting, b_t^* is introduced as a function using the projections of the covariance and variance onto the set of basis-functions already used. This is highlighted in Eq. 50.

$$b_t^* = \frac{\text{cov}(X_\tau, S_\tau)}{\text{var}(S_\tau)} \approx \frac{P_t(E(X_\tau \cdot S_\tau)) - P_t(E(X_\tau)) \cdot P_t(S_t)}{P_t(E(S_\tau^2)) - P_t(S_t)^2} \quad \text{Eq. 50}$$

Combined approach

The variance reduction techniques presented hitherto enable the formulation of a new hypothesis stated below:

Hypothesis 1.1.3.2 – Moment matching and control variates

By using the moment matching technique during the generation of trajectories and by sampling control variates at exercise of the option, the variability of the option prices estimate is reduced.

The variance reduction techniques discussed previously are implemented and a comparison between the original approach and the improved approach is made. For this, several test cases of call options on geometric Brownian motion are investigated and each test case is repeated fifteen times. The results are reported in Table 32 where the left column displays charts for the original approach while the right column displays charts for the improved approach using moment matching and control variates. For each approach, five test cases with different parameterizations of geometric Brownian motions are investigated and each test case is repeated fifteen times leading to an experiment featuring seventy five trials for the two approaches.

Table 32: Comparison between original approach (left graphs) and control-variate improved approach (right graphs)


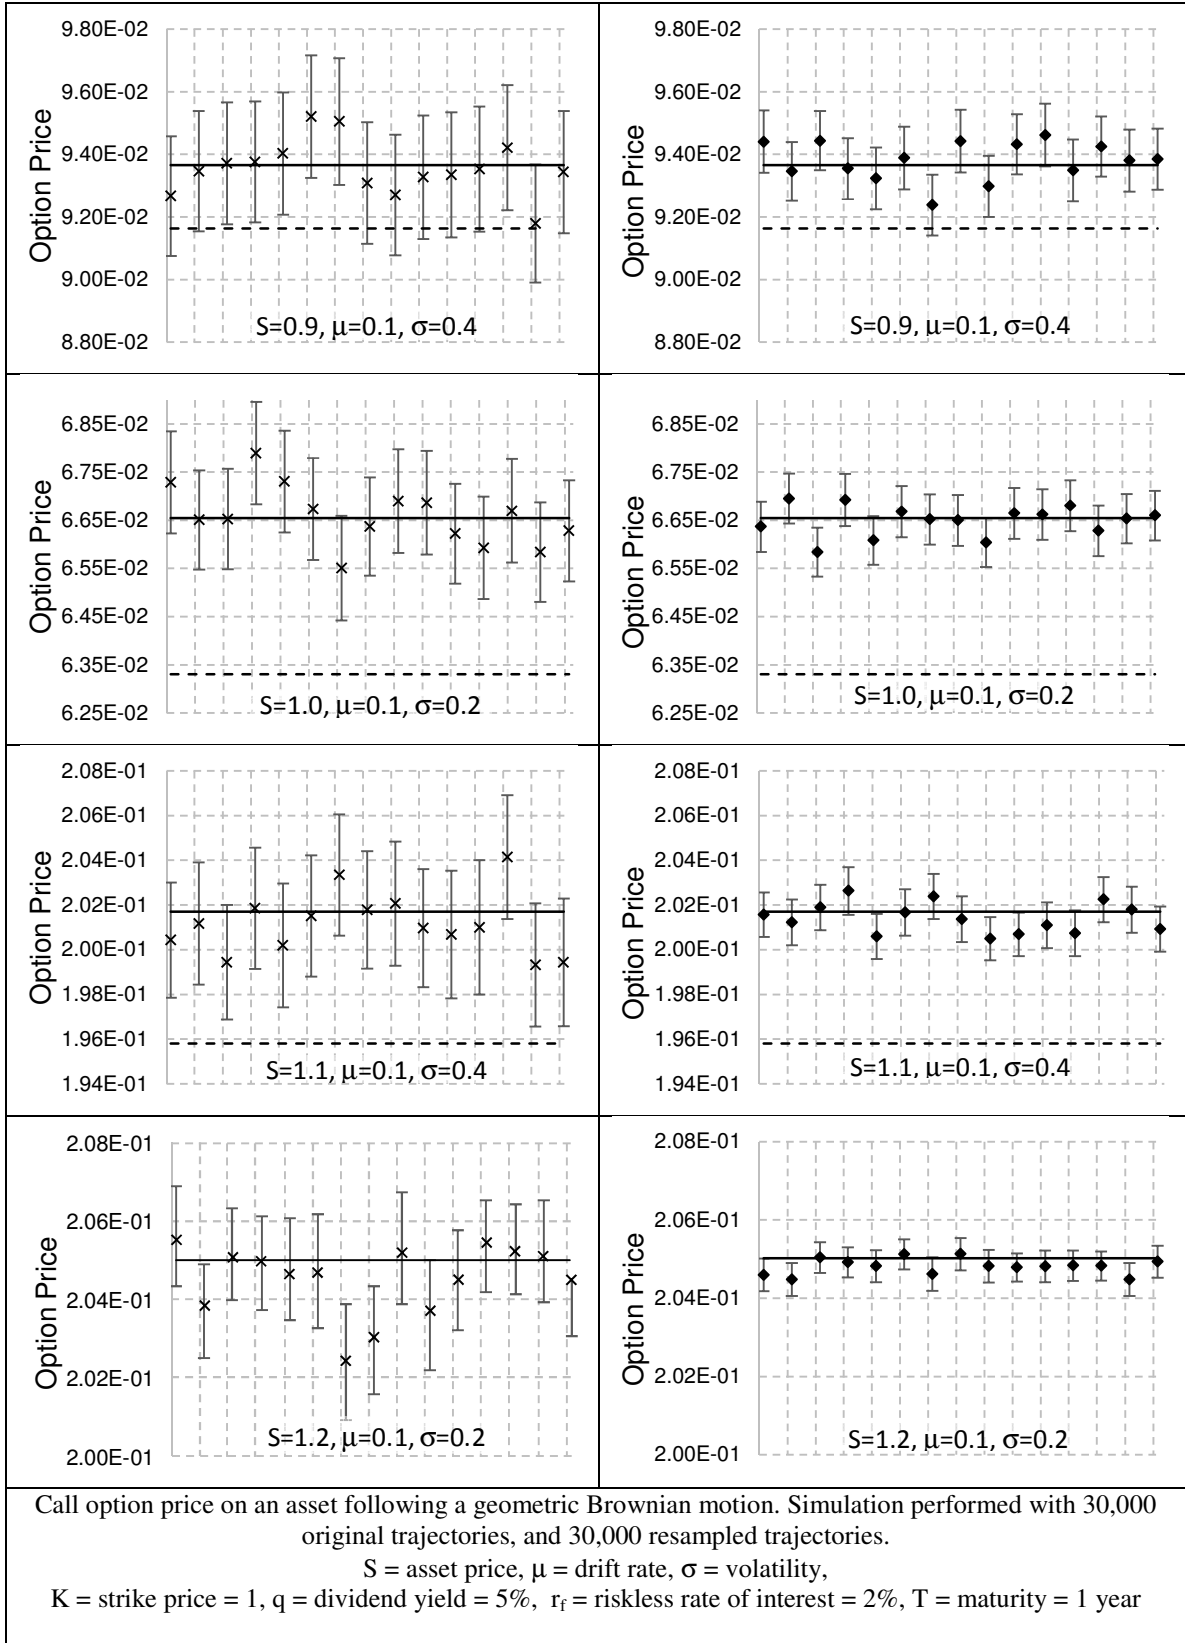


Table 32 Continued



The charts in Table 32 highlights many interesting aspects since each chart features the true American option price represented by the straight line, the true European option price represented by the dotted line, as well as the American option price computed by the proposed methodology (cross signs for baseline method results and rhombus signs for improved method results), and finally the 95% confidence interval around the American option price estimate represented by the vertical bars.

One objective is to be as close as possible to the true American option price. In this regards, it seems that the variance-reduction improved technique yields better results as the black rhombuses (representing the new approach) are on average much closer to the straight line than the black crosses (representing the unmodified approach). Another objective is to reduce as much as possible the width of the 95% confidence interval. Indeed, one objective of the proposed method is to estimate the value of the timing flexibility offered to decision-makers. Mathematically speaking, this value is the early-exercise premium of American options when compared to European options (i.e. the distance between the straight line and the dotted line). If the 95% confidence interval is so wide that it encompasses the price of the equivalent European option, then the early-exercise premium is not statistically significant and the proposed method fails to establish the value of the timing flexibility. Ensuring there is as little overlap as possible between the confidence interval surrounding the American option price and the European option price is therefore of paramount importance. Here again, the improved method outshines the baseline method with a drastic reduction in the width of the 95% confidence interval: simple computations using the standard error of the option prices indicate that the widths of confidence intervals are reduced on average by 20%, reaching up to 70% for deep in-

the-money options. Better results for in-the-money options are expected as the discounted underlying asset price, which serves as control variate, is more correlated with the option payoff when options are in-the-money. Owing to the better accuracy and the reduction in variance offered by the improved method, the hypothesis is successfully verified.

7.2.4 Quasi-Monte Carlo simulations

In the quest to lower the variability of the option prices obtained via Monte Carlo simulations, the pseudo-random number generator is also investigated. Indeed, repeated experiments yield slightly different option prices and slightly different trigger boundaries owing to the changing seeds used by the pseudo-random number generators. The change of seeds results in slight variations in the quality of the sequence of pseudo-random numbers used (non-uniformity, serial correlations) and may give rise to a bias in the simulation.

Jackel [190] argues that low-discrepancy sequences provide superior performance when trying to generate uniformly distributed numbers for the purpose of inverse transform sampling¹. The discrepancy of a sequence measures whether the amount of numbers in a sequence that is in an arbitrary set is proportional to the measure of this arbitrary set. Therefore, a low-discrepancy sequence of numbers on the unit interval ensures that numbers are uniformly distributed on the unit interval and that the amount of numbers in any segment contained within the unit interval is in proportion to the length of this segment. As a result, low-discrepancy sequences exhibit no gap or clustering and this is illustrated in Figure 44 which compares a two-dimensional set of uniformly

¹ Inverse transform sampling is a method to generate pseudo-random number from any probability distribution by generating first uniformly distributed numbers on the unit segment and by using next the inverse of the cumulative distribution function. Also discussed in section 7.2.3.

distributed numbers generated with the linear congruential generator implemented in the Microsoft VBA environment and a two-dimensional Sobol sequence. Exhibit (a) corresponds to the pseudo-random number generator and features gaps and clustering. Exhibit (b) corresponds to a two dimensional Sobol sequence and seems evenly distributed. As a result, one natural refinement to the proposed methodology consists in using low-discrepancy sequences instead of pseudo-random numbers for the simulations.




Figure 44: 20,000 uniformly distributed numbers across two dimensions: (a) represents pseudo-random numbers from a linear congruential generator while (b) represents a Sobol sequence

Unlike pseudo-random numbers, low-discrepancy sequences are not serially uncorrelated: by design, low-discrepancy sequences take into account points that have already been sampled when drawing new points in order to avoid clusters and gaps. This means that the numbers are not independent and the autocorrelation plot in Figure 45 illustrates the serial correlation for lags up to fifty. As may be observed, several autocorrelation bars are outside of the 99% confidence bands indicating that serial correlation is significant. This is prevalent for even number lags. The serial correlation requires care when using low-discrepancy sequences for the purpose of quasi-Monte

Carlo simulations: to avoid the introduction of unwanted correlations when constructing trajectories, one properly initialized low-discrepancy sequence must be used for each dimension of the problem (i.e. for each time cross-sections in the simulation).




Figure 45: Serial correlation at various lags for a sequence of 1,000 Sobol numbers

Several low discrepancy sequences have been proposed for quasi-Monte Carlo simulations: Van-der-Corput sequence, Halton sequence [191], Niederreiter sequences [192], and the Sobol sequences [193]. However, these sequences are not all well suited for high-dimensional applications such as the pricing of path-dependent options. As mentioned earlier, each time step represents one dimension and path-dependent options require a fine time discretization which may lead to several hundred dimensions. An American option with one year to maturity requires at least 250 dimensions if one number is required for each simulation time step and the discretization is done for each trading day of the year. Nevertheless, Sobol et al. [194] argue that properly initialized Sobol sequences may be used in high dimension applications. Jackel [190] indicates that the rate of convergence of quasi-Monte Carlo simulations is not one over the square root of the number of samples (as in traditional Monte Carlo simulations) but closer to one over the number of samples which leads to a substantial gain in computational efficiency. The exact formulation of the convergence rate is a bit more complicated and introduces

the number of dimensions; this is however beyond the scope of this research and an interested reader is referred to Jackel [190] for more details. As a consequence of these favorable properties, Sobol low-discrepancy sequences are used for preliminary testing. However, several hurdles have been encountered and are discussed in the following paragraphs.

Interactions have been observed between the Sobol's low-discrepancy sequences and the linear congruential generator of pseudo-random numbers. These interactions are particularly obvious when the number of replications in the simulation is a power of two (i.e. 16384, 32768, 65536...). No definite answer to explain this observation has been found, if not for the fact that using non independent (i.e. correlated) low-discrepancy sequences to construct the empirical distribution from which the bootstrap resampling is performed may not be the wisest thing to do. Indeed, the significant serial correlation in Sobol's sequences previously highlighted may collide with the supposedly random, yet non perfect, numbers generated by the linear congruential generator.

This has motivated further research in the field of pseudo-random number generators. Indeed, it is widely believed [195] that the Microsoft VBA RND() function uses a method based on linear congruential generators at least in the older versions of Excel, while the newer versions of Excel may use a more reliable Mersenne Twister. This is however not officially documented. In this context, a local implementation of a good pseudo-random number generator is warranted to reduce the risk of using a potentially unreliable random number generator when the real options analysis implementation is ported from one computer machine to another. The Mersenne Twister is a popular pseudo-random number generator developed by Matsumoto and Nishimura [196] which

has several desirable properties: it has a very long period of $2^{19937}-1$ which is so large that the likelihood of sampling the same sequence of numbers twice during a typical Monte Carlo simulation is null, and it passes the Diehard reference tests [197] for statistical randomness. For all these reasons, the Mersenne Twister has been implemented in the option analysis tool using the *mt19937ar.dll*¹ subroutines. This enables the user to benefit from the execution speed of subroutine coded in C from within the Excel VBA environment.

Uniform numbers are required twice in the proposed methodology. The first time is when the uncertainties are simulated under the physical probability measure. The second time is during the bootstrap resampling when using the sampling wheel algorithm. Low-discrepancy sequences can therefore be used in the first step, in the second step, in both steps, or in none of these two steps. The quasi-Monte Carlo method is therefore tested for these different cases which are named as follows: PP MC for pseudo-random numbers used in both the simulation and resampling, QQ MC for Sobol low-discrepancy sequences employed in both simulation and resampling, PQ MC for pseudo-random numbers applied in the simulation and Sobol low-discrepancy sequence used in the resampling, QP MC for Sobol low-discrepancy sequences employed in the simulation and pseudo-random numbers applied in the resampling, MT MC for Mersenne Twister pseudo-random numbers used in both the simulation and resampling, and finally, MT RND for VBA generated pseudo-random numbers employed in both the simulation and resampling.

¹ Retrieved August 2015: www.math.sci.hiroshima-u.ac.jp/~m-mat/MT/VERSIONS/ASSEM-DLL/assem-dll.html

Hypothesis 1.1.3.3 – Quasi-Monte Carlo simulations

Using Sobol's low discrepancy sequence in lieu of pseudo-random numbers increases the convergence of the least-squares Monte Carlo method

Results of convergence tests are given in Figure 46 for the evolution of the relative error as a function of the number of simulations and resampled trajectories. There is no obvious pattern and no obvious better solution. This is somewhat surprising as the Sobol low-discrepancy sequences were expected to perform better. Results from another convergence test are given in Figure 47 where the evolution of the standard error as a function of the number of simulations and resampled trajectories is reported in a log-log graph. The slope of the curve in the log-log graph enables a quick estimation of the convergence rate and surprisingly, the low-discrepancy sequences do not outperform the pseudo-random numbers. In fact, the plots of Figure 47 are confounded and the rates of convergence are identical for all cases. The convergence rates, measured by the slope of the line in a log-log scale or by the exponent value in a power regression are close to -0.5. This means that the convergence rate is close to the inverse of the square root of the simulation number ($1/\sqrt{n}$) which is, in fact, typical of regular Monte Carlo simulations. It is surprising that the use of low-discrepancy sequences does not improve the convergence rate. This is possibly due to the high dimensionality of the problem for which low-discrepancy sequences are known to be struggling. Another explanation may be unwanted interactions between the initial simulation of trajectories and the subsequent resampling of these trajectories which might introduce some side-effects owing to the autocorrelation structure of low-discrepancy sequences. In any case, the hypothesis is

proven wrong and the use of quasi-Monte Carlo simulation does not help in this application.




Figure 46: Relative error as a function of the number of simulations for different pseudo-random and quasi-random number generators

Call option price on an asset following a geometric Brownian motion
 $S = 0.8, \mu = 10\%, \sigma = 20\%, K = 1, q = 5\%, r_f = 2\%, T = 1$ year




Figure 47: Standard error as a function of the number of simulations for different pseudo-random and quasi-random number generators

Call option price on an asset following a geometric Brownian motion
 $S = 0.8, \mu = 10\%, \sigma = 20\%, K = 1, q = 5\%, r_f = 2\%, T = 1$ year

7.2.5 Scoping conditional expectation regressions

For American options, the exercise policy defines the time and price of the underlying asset for which early-exercise maximizes the option value. Thus, the exercise policy yields an early-exercise boundary, which is a locus of time-indexed critical prices maximizing the value of the option. Critical prices lay at the edge between the immediate exercise region and the continuation region: they are the underlying asset prices for which the payoffs from immediate exercise exactly match the continuation values. In the least-squares Monte Carlo algorithm, the continuation value is computed at each time step using a conditional expectation regression. Improving the conditional expectation regressions should help refine the generation of the early-exercise boundary. In turn, this optimizes the exercise policy and maximizes of the option value.

Hypothesis 1.1.3.4 – Scoping the regression domain

Reducing the domain over which the continuation value conditional expectation is regressed facilitates the search for the critical price

Natural boundary

Longstaff and Schwartz [136] argue that the regression of the conditional expectation is improved by using only in-the-money paths “*since it allows [...] to better estimate the conditional expectation function in the region where exercise is relevant and significantly improves the efficiency of the algorithm*”. Restricting the regression domain even further may improve the quality of the continuation value regressions. To do this, the concept of *natural boundary*, introduced by Rasmussen [186] for other purposes, is used. Unlike the early-exercise boundary which is defined as the locus of points for

which the holding value exactly matches the immediate exercise value, the natural boundary is defined as the locus of points for which the value of holding the option until maturity exactly matches the immediate exercise value. The difference between the two boundaries is that the trigger boundary is constructed using the holding value with possibility of exercise at any time until maturity (an American real option), while the natural boundary is constructed using the holding value with no possibility of intermediate exercise before maturity (a European real option). Since European options have less value than American options, the natural boundary yields a locus of points less in-the-money than the early-exercise boundary. An improvement to the least-squares Monte Carlo method is therefore suggested and consists in restricting the regression domain of the continuation value by identifying the natural boundary and then regressing the continuation value using trajectories deeper in-the-money than the natural boundary. This reduction of the regression domain is illustrated in Figure 48.




Figure 48: Removing points inside the natural boundary (right graph) scopes down the conditional expectation regression domain and improves the estimation of critical prices for American and Bermudan real options (call option depicted)

The natural boundary is straightforward to define: at each time step t_k starting from the next to last one, a bisection algorithm is used to search for the value of the business prospect S_{t_k} such that the European option price $V_{t_k}^{EMC}$ equals the option immediate payoff P as shown in Eq. 51.

$$V_{t_k}^{EMC}(S_{t_k}) = P(S_{t_k}) \quad \text{Eq. 51}$$

The bisection algorithm is iterative and therefore several identical European options with different spot prices must be priced at each time step. To speed-up the computations, the pricing is carried out with the same set of returns but with different simulation starting points: right before expiration, the European options have a one-step maturity and therefore only returns associated with the first time step of the trajectories are used; for the preceding step, the European options have a two-step maturity and therefore only returns associated with the first two time steps of the trajectories are used. The search for critical prices featuring the bisection algorithm with efficient use of simulated data is illustrated in Figure 49. In the depiction, the position of the natural boundary is estimated eighteen steps before expiration and to speed-up the computation, the returns associated with the first eighteen steps of one original trajectory are used to estimate the price of three different options with underlying spot prices at (1) $S_{(42)}=2.5$, (2) $S_{(42)}=0.5$, and (3) $S_{(42)}=1.5$




Figure 49: Bisection algorithm is used to search for critical prices using a single set of returns

Imposed boundary

The natural boundary provides a lower bound for the critical prices of an American call option, while it provides an upper bound for the critical prices of an American put option. To better scope the regression domain, another boundary is required. For American put options, a lower boundary can be imposed with the null value of the underlying asset. For American call options, there is unfortunately no simple upper boundary and an ad-hoc boundary has to be found. It is customary in option pricing via finite-difference schemes to assume that the upper boundary of the grid mesh is at three times the strike price and that this upper boundary represents an infinite underlying asset price. The same assumption is proposed for the upper bound of the continuation value regression domain: only trajectories less in-the-money than a fixed upper boundary at three strike prices are used for the continuation value regression. The boundaries of the regression domain for a notional call option and a notional put option are depicted in Figure 50 where the left graph depicts the continuation value regression domain for a call

option, while the right graph depicts the continuation value regression domain for a put option.





Figure 50: Scoping the regression domain for the conditional expectation continuation value

Combined boundaries

The natural boundary and the imposed boundary are combined together to reduce the size of the domain used for the regression of the continuation value. The results are given in Table 33 for three different types of geometric Brownian motions. Each test is repeated fifteen times and the red (+) indicate the early-exercise boundary position error using traditional continuation value regressions while the black (x) indicate position errors using the scoped continuation value regressions.

Table 33: Comparison between traditional continuation value regressions and scoped continuation value regressions



In almost all cases, the position errors with the scoped regressions are better with error decreased by 50% on average. As a result the hypothesis stating that scoped regressions enable a better evaluation of critical prices is verified.

7.2.6 Projection basis selection for continuation value regressions

The generation of the trigger boundary using Monte Carlo simulations is a notoriously difficult task. At each time step, it relies on searching for the critical price, price for which the continuation value equals the immediate exercise payoff. Consequently, an accurate trigger boundary requires an accurate way of finding the critical price and an accurate regression of the continuation value. There are nonetheless several issues hindering the ability to properly regress the continuation value.

One issue is related to the “noise” in the data used for the regression: as the time to expiration increases, the effect of diffusion becomes more and more prevalent, and the relationship between the explanatory variable (i.e. the underlying asset price) and the dependent variable (i.e. the continuation value) becomes less and less obvious. This is observed in Figure 51 where dependent and explanatory variables are plotted far from expiration on the leftmost graph and right before expiration in the rightmost graph. Far from expiration, the diffusion effect is prevalent for lower underlying prices whereas close to expiration, the diffusion effects are curbed and the continuation values are close to the payoff values. The simulation is done with 20,000 trajectories, with a drift of 10%, a dividend yield of 5%, a volatility of 20%, and a risk free rate of 2%.




Figure 51: Continuation values for a call option on a geometric Brownian motion

This set of graphs also raises the question of the choice of a projection basis for regressions both close and far from expiration. The regressions are usually done with orthogonal polynomial families such as the simple monomials, the Laguerre polynomials, the Legendre polynomials, and the Chebyshev polynomials as reported by Stentoft [158] [198]. However, polynomial regressions often introduce artificial wiggles at the edge of the regression domain, known as Runge's phenomenon. These wiggles can lead to multiple solutions to the critical price equation, a feature exacerbated by the proximity and collinearity between the continuation value and the immediate payoff as shown in Figure 52. In Figure 52, the objective function represents the difference between the continuation value and the immediate payoff: the regression looks fine unless zoomed-in where oscillatory behavior becomes more apparent and precludes the selection of a single critical price. Indeed, these oscillations are a significant problem because numerical solvers based on the Newton-Raphson method or on the bisection method are unable to discriminate the spurious solutions from the correct critical price solution. Black rhombus represents critical price solution obtained with Newton-Raphson while green square represents the reference critical price obtained with the finite-difference method. The simulation is done with 20,000 trajectories, with a drift of 10%, a dividend yield of 5%, a volatility of 20%, a risk free rate of 2%, and a time to expiration of 30 days.

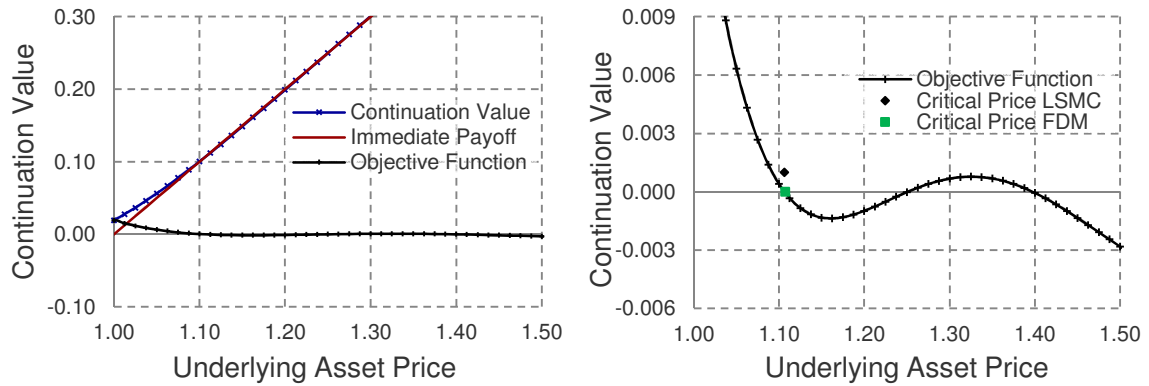



Figure 52: Continuation value fitted with the first four Legendre polynomials for a call option on a geometric Brownian motion.

A quick literature review indicates that pricing American options usually features regressions with up to five polynomials (Longstaff and Schwartz [136], Stentoft [198]), even though some analyses feature regressions with bases of up to twenty polynomials as described in Moreno and Navas [199]. Large projection bases usually include polynomials of higher degrees which produce more wiggles and more spurious solutions. Besides, large projection bases come at an increased computational complexity cost.

In this context, a projection basis featuring fewer but adequately chosen terms is proposed as a way to curb the wiggles and mitigate the risk of spurious critical price selection. After investigation of the shape of the continuation value, a projection basis consisting of four elements is selected: a constant, the payoff itself, a decreasing exponential in the payoff, and another decreasing exponential in the square root of the payoff. This is motivated by several observations: the conditional expectation of the continuation value is very similar to the option payoff right before expiration which motivates the use of the payoff as a basis element; far from expiration, the conditional expectation of the continuation value retains some of the payoff appearance but is smoother and exhibit some convexity (at least for put and call options) therefore

motivating the two exponentials in the payoff; for very large values of the payoff, the continuation value is almost the payoff itself meaning that the convexity terms must vanish with very large payoffs. Besides these heuristic arguments, there is no theoretical reason for this particular choice of projection basis. The example in Figure 53 highlights the reduced oscillations of the objective function (regressed continuation value minus immediate exercise payoff) using this LSMC projection base. The experiment is performed with 20,000 trajectories, a drift of 10%, a dividend yield of 5%, a volatility of 20%, a risk free rate of 2%, and a time to expiration of 30 days.



$$\{\phi_i\}_{i=0..3} \begin{cases} \phi_0 = 1 \\ \phi_1 = \text{Payoff} \\ \phi_2 = \exp(-\text{Payoff}) \\ \phi_3 = \exp\left(\frac{-\sqrt{\text{Payoff}}}{2}\right) \end{cases} \quad \text{Eq. 52}$$

Figure 53: Critical price search with ad-hoc set of basis functions $\{\phi_i\}$ exhibits fewer wiggles

7.2.7 Multi-start Monte Carlo simulations

In the relentless quest to improve the trigger boundary generation, another issue investigated is the erratic behavior of regressions during extrapolations. In fact, if no underlying asset trajectory crosses the still unknown trigger boundary, the regression of the continuation value is performed with no data point close to the critical price. Thus, the search for the critical price leads to predictions outside the range of data points used for the regression and therefore to extrapolations. Extrapolations, especially when using polynomial regressions, are notoriously bad. This is highlighted in Figure 54 where the

critical price solutions are very noisy in the extrapolation region close to the beginning of the simulation.




Figure 54: Solving for the critical price at each time step involves extrapolations close to the starting point of the simulation since none of the trajectories cross the trigger boundary (left graph). This results in noisy and approximate critical price solutions (right graph) close to the starting point.

To address this problem, a significant improvement to the least-squares Monte Carlo methodology is proposed and consists in performing a simulation for which several adequately chosen starting points are retained to initialize the trajectories. This requires splitting the least-squares Monte Carlo method into two distinctive parts: the generation of the early-exercise boundary using the multi-start Monte Carlo simulations on the one hand, and the pricing of the option using a known early-exercise boundary on the other hand. The proposed multi-start Monte Carlo improvement stems from the observation that the quality of the least-squares regressions improves as more points and therefore more trajectories lie “in-the-money”. Indeed, with more trajectories “in-the-money”, the continuation value regressions as well as the critical price estimations become more accurate since more trajectories are likely to be in the neighborhood of the trigger boundary.

In fact, even when the trigger boundary is reasonably well approximated as a whole, the approximation deteriorates close to the initial starting time in a traditional Monte Carlo simulation. This is mostly an extrapolation problem when using numerical solvers to search for the critical price with continuation value regressions. At the initial steps of the simulation, the effects of diffusion are limited and the simulated business prospect values used for the continuation value regressions are not dispersed enough to encompass the neighborhood of the critical price. Rasmussen [30] suggests starting the simulation prior to the current time (i.e. back in time) in order to let the diffusion artificially disperse the trajectories and “*to provide sufficient in-the-money observations to estimate the exercise boundary*”. Even though this is a step in the right direction, this solution does not go far enough and several improvements are suggested. First, the objective should not be to provide a sufficient number of in-the-money observations but rather to provide a sufficient number of observations close to the unknown early-exercise boundary so as to avoid extrapolations during the critical price estimation. Next, the Rasmussen technique would not be efficient in the proposed real options methodology from a computational point of view as a longer clock-time must be simulated to accommodate the back-in-time starting point. Finally, the Rasmussen technique does not guarantee that the dispersion is sufficient to provide observations close to the critical price. In fact, for at-the-money call options with low risk-free rates and large volatilities, the drift of a geometric Brownian motion under the equivalent martingale measure is likely to be negative and the proposed approach tends to drive trajectories away from the initial critical prices of the trigger boundary. These observations lead to the following hypothesis.

Hypothesis 1.1.3.5 – Multi-start Monte Carlo simulations

Using several starting points for the generation of trajectories facilitates the construction of the early-exercise boundary by improving the quality of the least-squares regression in the neighborhood of the critical price.

Instead, the suggested multi-start Monte Carlo simulation is based on the fact that the position of the trigger boundary is not affected by the initial business prospect value. Therefore, using different starting points for the simulations will yield the same boundary. From a computational point of view, no extra complexity is introduced: instead of generating $n \cdot m$ replicating trajectories in a traditional setting, a multi-start Monte Carlo featuring m starting points will have n replicating trajectories attached to each starting point. As such, the technique illustrated in Figure 55 does not increase the computational burden.




Figure 55: Simulation with multiple starting points leads to fewer extrapolations and less noisy critical price estimates, even far from expiration

The remaining question concerns the choice of the starting points. These are chosen so as to maximize the likelihood that trajectories “encompass” the early-exercise boundary while minimizing the likelihood of sampling the domain where early-exercise is not optimal. The strike price and a multiple of the strike price can almost always be used to select two extreme starting points. For instance, in a preceding section dealing with regressions, the upper bound is set as three times the strike price for call options, while the lower bound is set with a price close to zero for put options. The initial point of the natural boundary derived previously provides another excellent lowest starting point for call options or another highest starting point for put options. Finally, the domain in between these extreme starting points can be evenly discretized to get evenly-spaced simulation starting points. One range of starting points found by considerable trial and errors and used in the remaining of this research is illustrated in Figure 56: the range starts at 85% of the terminal value of the natural boundary and ends at the strike price plus 300% of the difference between the initial value of the natural boundary and the strike price.



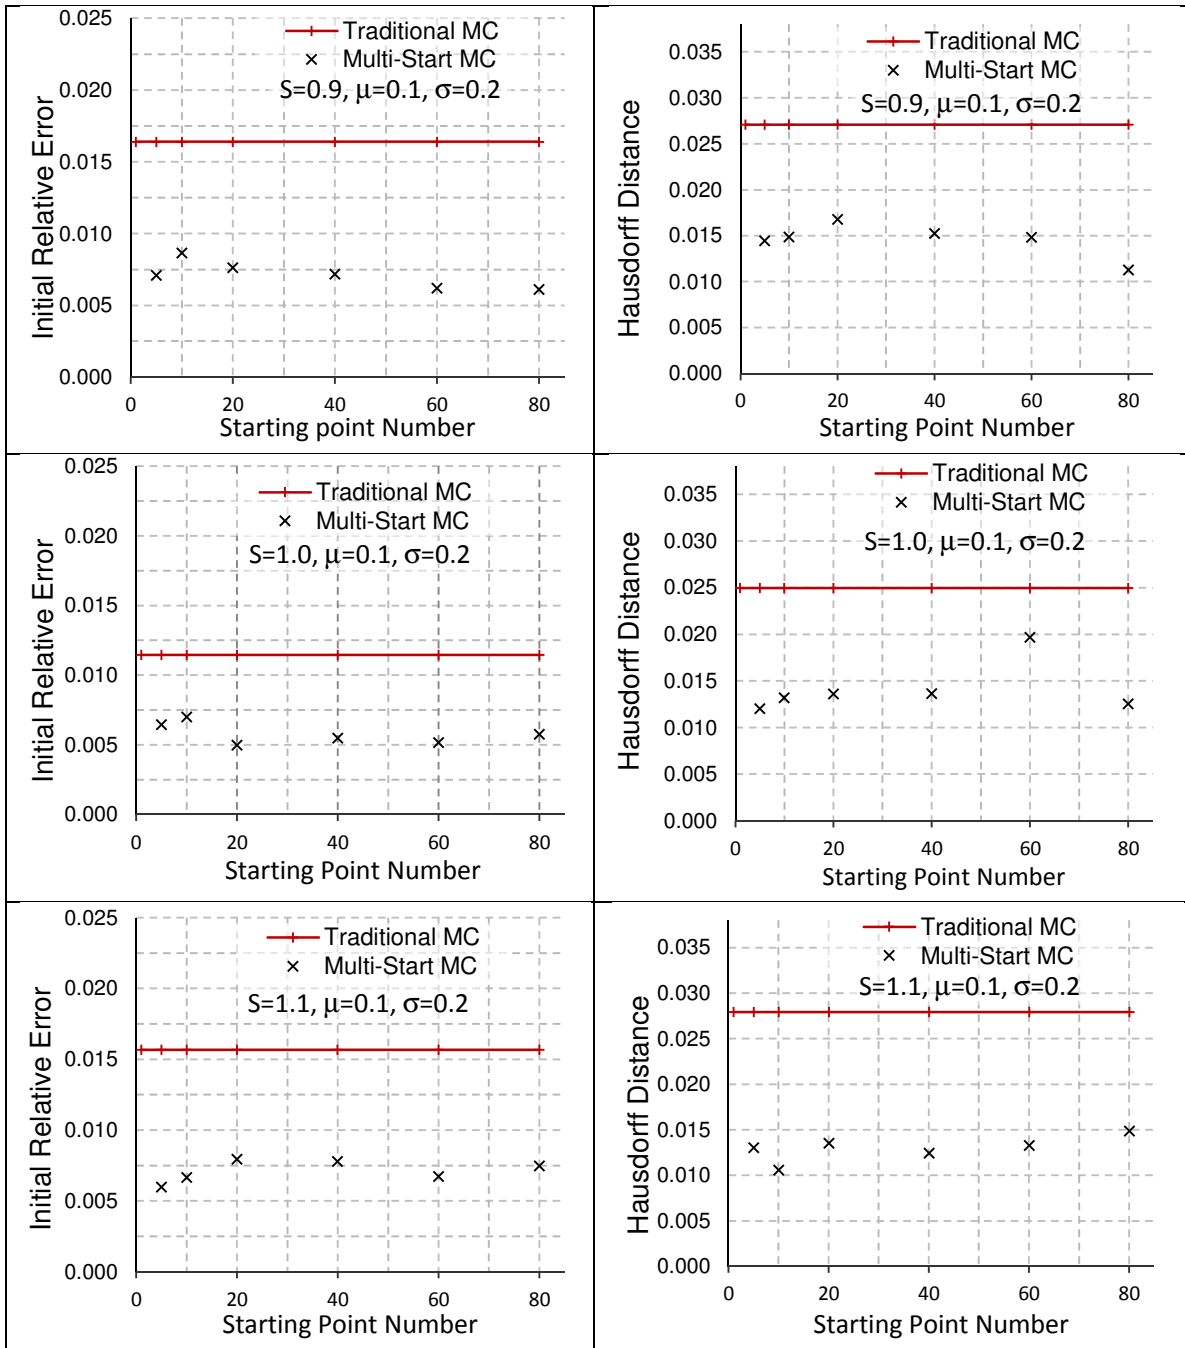


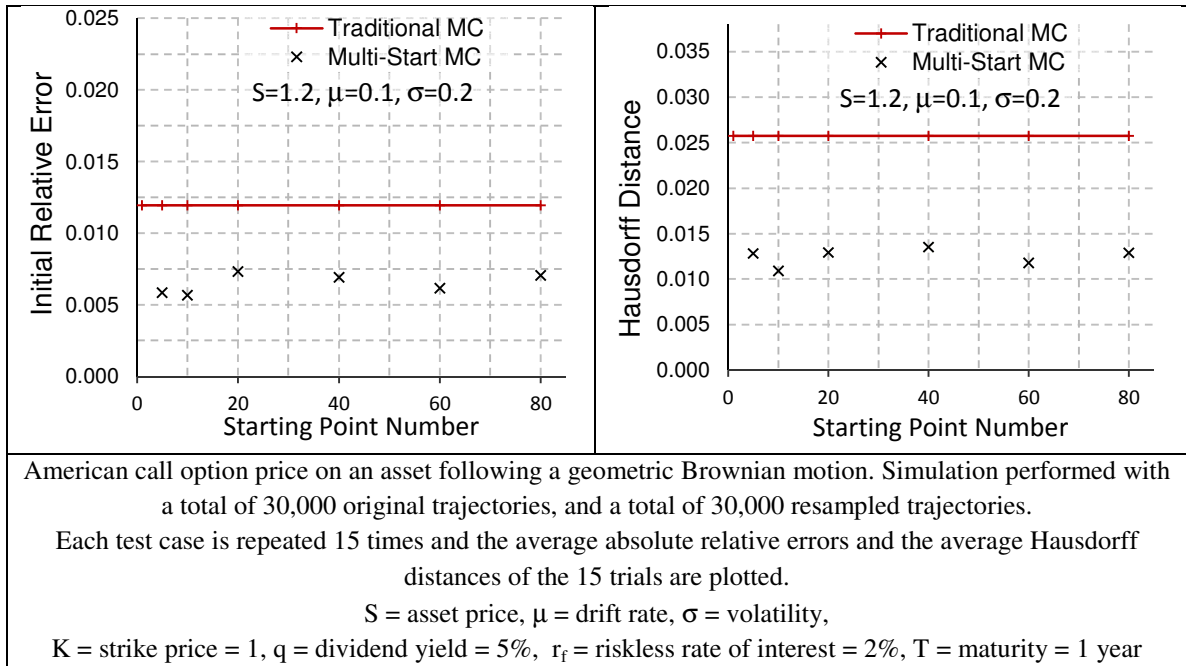
Figure 56: Range of values for starting points in Multi-Start Monte Carlo

Table 34 describes the position error of the early-exercise boundary for five American call options on a geometric Brownian motion when varying the number of starting points. The red line corresponds to a traditional least-squares Monte Carlo algorithm (with a single starting point at the current spot price), while the black crosses correspond to the proposed multi-start Monte Carlo algorithm with five to eighty starting points. To account for the variability of results when performing Monte Carlo simulations, each test case is repeated fifteen times. Two metrics are used for these comparisons: the initial relative error and the Hausdorff distance. The initial relative error is of interest because it indicates how the experimental trigger boundary differs from the reference boundary at the very time when investment decisions must be made. The Hausdorff distance indicates how close the shapes of the experimental and reference boundaries are. Averages of absolute initial relative errors for these fifteen trials are plotted in the charts of the left column, while averages Hausdorff distances for the fifteen trials are reported in the charts of the right column.

Table 34: Comparison of the average error of early-exercise boundaries using traditional Monte Carlo simulations (red line) and using multi-start simulations (x) for five different call options







Results in Table 34 indicate a significant reduction in the early-exercise boundary position error when using Multi-start Monte Carlo simulations. On average, the absolute values of the initial relative errors are reduced by 50% while the Hausdorff distances are reduced by about 45%. It is worth mentioning that this is achieved without additional computation cost, and execution times are comparable to a typical two-stage least-squares Monte Carlo simulation (one stage for early-exercise boundary generation and one stage for option pricing). Out of interest, the improvements are not monotonous: when going from one starting point to two starting points, results degrade because the two starting points are encompassing, yet very far, from the early-exercise boundary. Improvements only appear after five starting points are selected, become very noticeable when going from ten to twenty starting points, slow down when going to forty, and after fifty starting points no extra benefits are observed. Therefore, fifty starting points are selected in the remaining of this research. The hypothesis that using Multi-start Monte Carlo simulations improves the generation of the early-exercise boundary is thus verified.

7.2.8 Regression and filtering of critical prices to form trigger boundary

The refinements of the least-squares Monte Carlo algorithm described up to this point aim at improving the estimation of the critical price at each and every time step in the simulation. The next step is to actually construct the trigger boundary from these improved critical price estimates. The straightforward approach of Longstaff and Schwartz [136] consists in using this locus of critical prices as the early-exercise boundary. Yet, the construction of the trigger boundary may benefit from additional theoretical results available in the literature. First, the trigger boundary is a monotonous function of time for usual put and call options with continuous dividend yields as proved by Jacka [200] and Pham [201]. Second, Kim [178] and Pham [201] show that the trigger boundary is smooth and continuous except at maturity. At maturity, a discontinuity may occur depending on the relative values of the continuous dividend yield and the risk-free rate of return (Kwok [202]). One may then formulate the following hypothesis.

Hypothesis 1.1.3.6 – Regression and filtering of the set of critical prices

Regressing the set of critical prices and removing critical price outliers improves the construction of the early-exercise boundary.

Regression of critical prices to generate trigger boundary


Based on the previous observations, it may be stated that the critical prices obtained with the least-squares Monte Carlo simulations can be regressed to yield a continuous and monotonous function and therefore remove the jaggies present in the locus of critical prices. Indeed, the proposed continuation value regressions yield critical price predictions that are noisy and that induce a non-monotonous early-exercise

boundary. A specific treatment is however required at expiration of the option to handle the possibility of discontinuity. This is done by simply removing the expiration point from the regression: in fact, at expiration, the trigger boundary is exactly at the strike price and therefore the true position of the boundary is known with certainty and does not need to be regressed. Similarly, at the time step immediately preceding expiration, the American option is a European option with no possibility of intermediate exercise. Consequently, the early-exercise boundary and the natural boundary are confounded at this time step. Because the natural boundary has already been constructed with good accuracy, this point does not need to be regressed and may serve as an anchor for the trigger boundary.

Selection of a regression basis for the trigger boundary

Once a decision to regress the critical price is made, the question of the selection of the projection basis arises. Any type of orthogonal projection basis should work but knowledge of the general shape of the trigger boundary can help select an adequate basis which in turn can improve the regression. For instance, Carr et al. [203] indicate that put and call options exhibit an infinite slope at expiration which means that a projection basis made of a finite number of polynomials is probably not adequate. Instead, a set of ad-hoc functions is suggested: a constant to account for the possible jump of the trigger boundary at expiration, a square-root function in the time to expiration to approximate the boundary behavior close to maturity, and finally an exponential in the time to expiration to model a vanishing contribution far from maturity. Besides these heuristic arguments, there is no theoretical reason for this particular choice. A notional example of critical price estimates

regression is depicted in Figure 57 and the basis-functions used in this regression are mentioned in Eq. 53.



$$\{\varphi_i\}_{i=0..3} \begin{cases} \varphi_0 = 1 \\ \varphi_1 = e^{-|T-t|^{0.33}} \\ \varphi_2 = \sqrt{|T-t|} \end{cases} \quad \text{Eq. 53}$$

Figure 57: Regression of critical price estimates (red) and true boundary (black)

Filtering and removal of outliers

Because the trigger boundary is “well behaved”, it becomes possible to perform much more than simply regressing the critical prices. For instance, it was previously mentioned that the locus of critical prices is extremely noisy and therefore some form of filtering might improve drastically the regression of the trigger boundary, provided it does not slow down the execution time. A two-step regression is suggested to perform this filtering: a multiple linear regression is first done on the crude set of critical prices to generate a crude trigger boundary reference which enables the estimation of residuals which helps filter out outliers; another multiple linear regression is performed next on the reduced set of critical prices to generate an improved trigger boundary.

According to Pope [204], the detection of outliers is typically done by estimating the Studentized residuals of a regression. The Studentized residuals recognize that the standard deviation of residuals changes from data point to data point in a regression. Although this is the correct way to detect outliers, this is cumbersome as the Studentized

residuals require the computation of the influence matrix (hat matrix denoted \hat{h} which maps the vector of dependent variables denoted y , to the vector of fitted values denoted \hat{y} , with terms $h_{i,j} = cov(\hat{y}_i, y_j) / var(y_j)$). This considerably slows down the procedure, and absent the influence matrix, the semi-Studentized residuals are used for filtering purposes. Semi-Studentized residuals do not account for the variability of the standard deviation of residuals from one data point to another which makes the computation faster as residuals are simply normalized by the sample average standard deviation of residuals. If the regression model has $m+1$ parameters to be estimated from n data points, then the semi-Studentized residuals are given by Eq. 54.

$$\varepsilon_i^* = \frac{y_i - \hat{y}_i}{\sqrt{\frac{1}{n-m-1} \sum_{i=1}^n (y_i - \hat{y}_i)^2}} \quad \text{Eq. 54}$$

The normalization of the residuals creates a scale-free t -like statistic following a t -distribution which provides information regarding the size of the residual [171]. If the Studentized residual is large, then the error is larger than what is expected by chance and some further checking is warranted. In the simplified approach proposed, the use of semi-Studentized residuals somewhat complicates the conclusion as the t -like statistic does not follow exactly a t -distribution. Furthermore, a large semi-Studentized residual is assumed to suggest that the corresponding data point is an outlier. Given that the number of critical price data points is much larger than the number of parameters to be estimated during the regression, the approximate t -distribution is replaced with the standard normal and the threshold for outlier identification is set at ± 1.96 corresponding to the 2.5 and 97.5 percentile of the standard normal distribution. An example of filtering critical prices is given in Figure 58 where the outliers represented in red are filtered out.

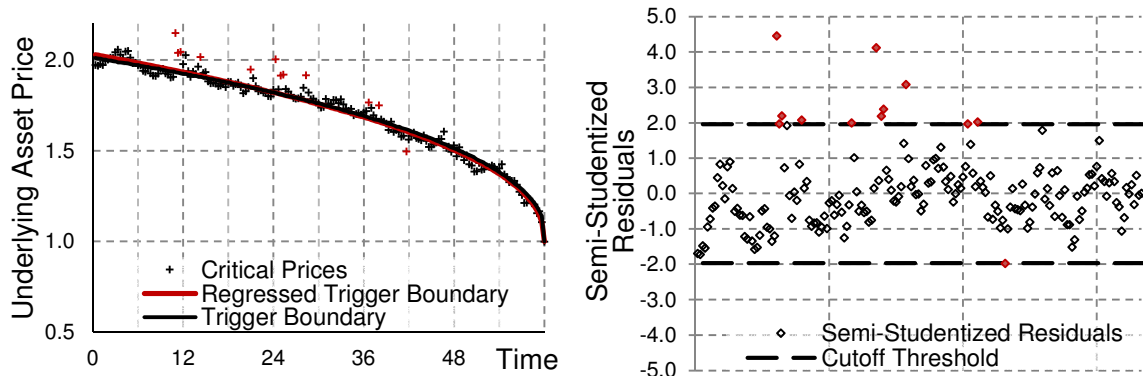

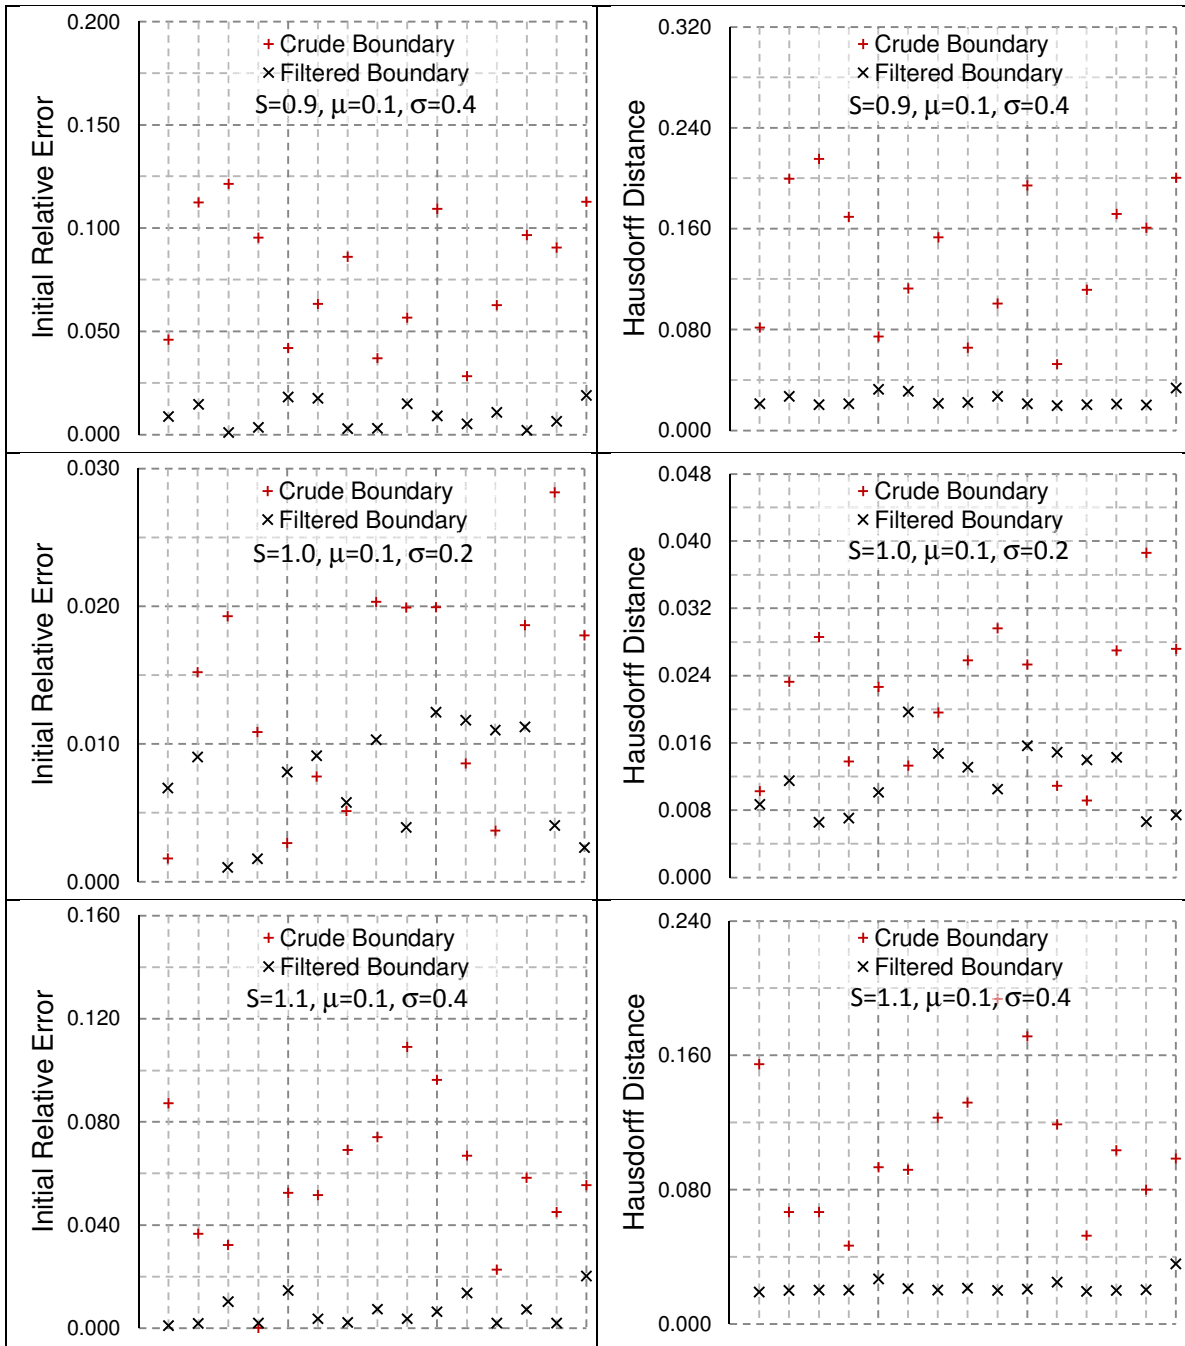


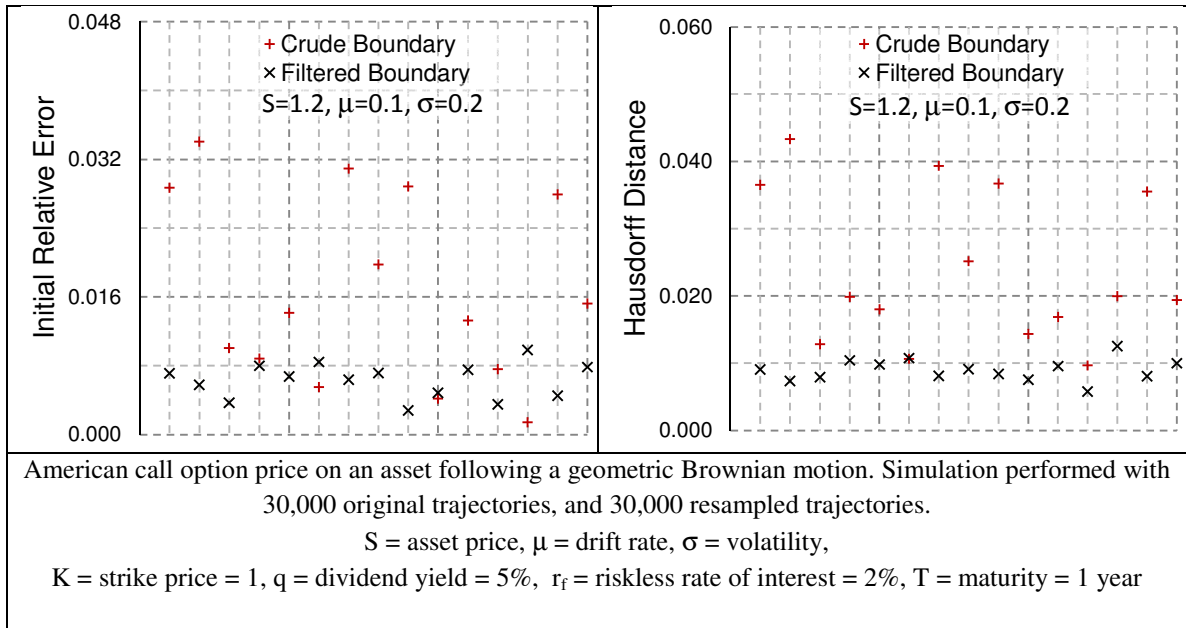
Figure 58: Removing outliers using semi-Studentized residuals

With the outliers filtered out, a new regression of the critical prices is performed and it yields an improved trigger boundary. Comparisons between the crude boundaries and the improved boundaries are provided in Table 35. The same two metrics are retained for these comparisons: the initial relative error and the Hausdorff distance.

Table 35: Comparison of boundary position errors for a crude locus of critical prices (+) and for the filtered and regressed trigger boundary (x) using five call options with each test case repeated fifteen times.







Improvements can be seen in each of the five cases but are particularly obvious when volatility is high. This is not surprising as this corresponds to cases where the experimental critical points are most noisy and therefore to cases where filtering out outliers and where regressing critical prices bring most benefits. Out of interest, the Hausdorff distances for the crude boundaries are all between 0.00 and 0.25 while they are between 0.00 and 0.03 for the filtered boundaries. This highlights the drastic improvements produced by this heuristic approach and this verifies the hypothesis stating that filtering outliers and regressing the locus of filtered critical points yield a better early-exercise boundary.

7.2.9 Matching new research questions and hypotheses

Preliminary testing was useful as it indicated that the methodology, as initially proposed, was not sufficient to achieve the required level of accuracy for the computation of option prices and the generation of early-exercise boundaries. As a result a new research question is formulated as follows:

Research Question 1.2.2 – Reducing variability of results

How can the variance of results obtained from the least-squares Monte Carlo simulation be reduced to yield consistent real options price estimations and consistent early-exercise boundary shapes?

To try to answer this research question, several hypotheses were formulated. The first one deals with the generation of trajectories and propose an improved bootstrap technique to ensure that the distributional properties of returns are preserved and that no artifacts are introduced when using the sampling wheel algorithm.

Hypothesis 1.1.3.1 – Pooling returns to increase size of sample to bootstrap

Pooling samples of returns from different time cross-sections or increasing the relative size from the original sample with respect to the bootstrap sample limits the repetitive sampling of the same highly-weighted return values.

The second hypothesis also deals with the generation of trajectories but takes a different approach: instead of trying to generate perfectly distributed returns, the technique recognizes that sampling errors are unavoidable but that they can be corrected “on the go” by using the same set of sampled returns to compute known quantities and estimate sampling errors.

Hypothesis 1.1.3.2 – Moment matching and control variates

By using the moment matching technique during the generation of trajectories and by sampling control variates at exercise of the option, the variability of the option prices estimate is reduced.

The third hypothesis also deals with the generation of trajectories but takes a new approach: instead of generating pseudo-random and “independent” numbers on the unit segment for the purpose of the inverse transform sampling, it proposes to use sequences of numbers, known as Sobol’s low-discrepancy sequences, which are known to be well (uniformly) distributed on the unit segment.

Hypothesis 1.1.3.3 – Quasi-Monte Carlo simulations

Using Sobol’s low discrepancy sequences in lieu of pseudo-random numbers increases the convergence of the least-squares Monte Carlo method.

The fourth hypothesis aims at improving the generation of the early-exercise boundary by limiting the scope of the conditional expectation of the continuation value regressions. By doing so, it is hypothesized that regression will be improved and that the estimation of critical prices will be more accurate.

Hypothesis 1.1.3.4 – Scoping the regression domain

Reducing the domain over which the continuation value conditional expectation is regressed facilitates the search for the critical price.

The fifth hypothesis aims also at improving the generation of the early-exercise boundary by ensuring that there are enough trajectories and therefore enough data points in the neighborhood of the (still unknown) critical price. It is hypothesized that having many points in the neighborhood of the critical price will improve the quality of the regression on the one hand and prevent the use of extrapolations during the critical price search on the other hand.

Hypothesis 1.1.3.5 – Multi-start Monte Carlo simulations

Using several starting points for the generation of trajectories facilitates the construction of the early-exercise boundary by improving the quality of the least square regression in the neighborhood of the critical price.

Finally, the last hypothesis aims at improving the generation of the early-exercise boundary by first removing critical price outliers and then by smoothing it out using regression. Indeed, once the outliers are removed, it is hypothesized that the filtered set of critical prices lends itself well for regression and that the regressed boundary is a better approximation of the true boundary than the set of *raw* critical prices.

Hypothesis 1.1.3.6 – Regression and filtering of the set of critical prices

Regressing the set of critical prices and removing critical price outliers improves the construction of the early-exercise boundary.

This leads to a new and more extensive mapping between research questions and hypotheses. This updated mapping is described in Figure 59.

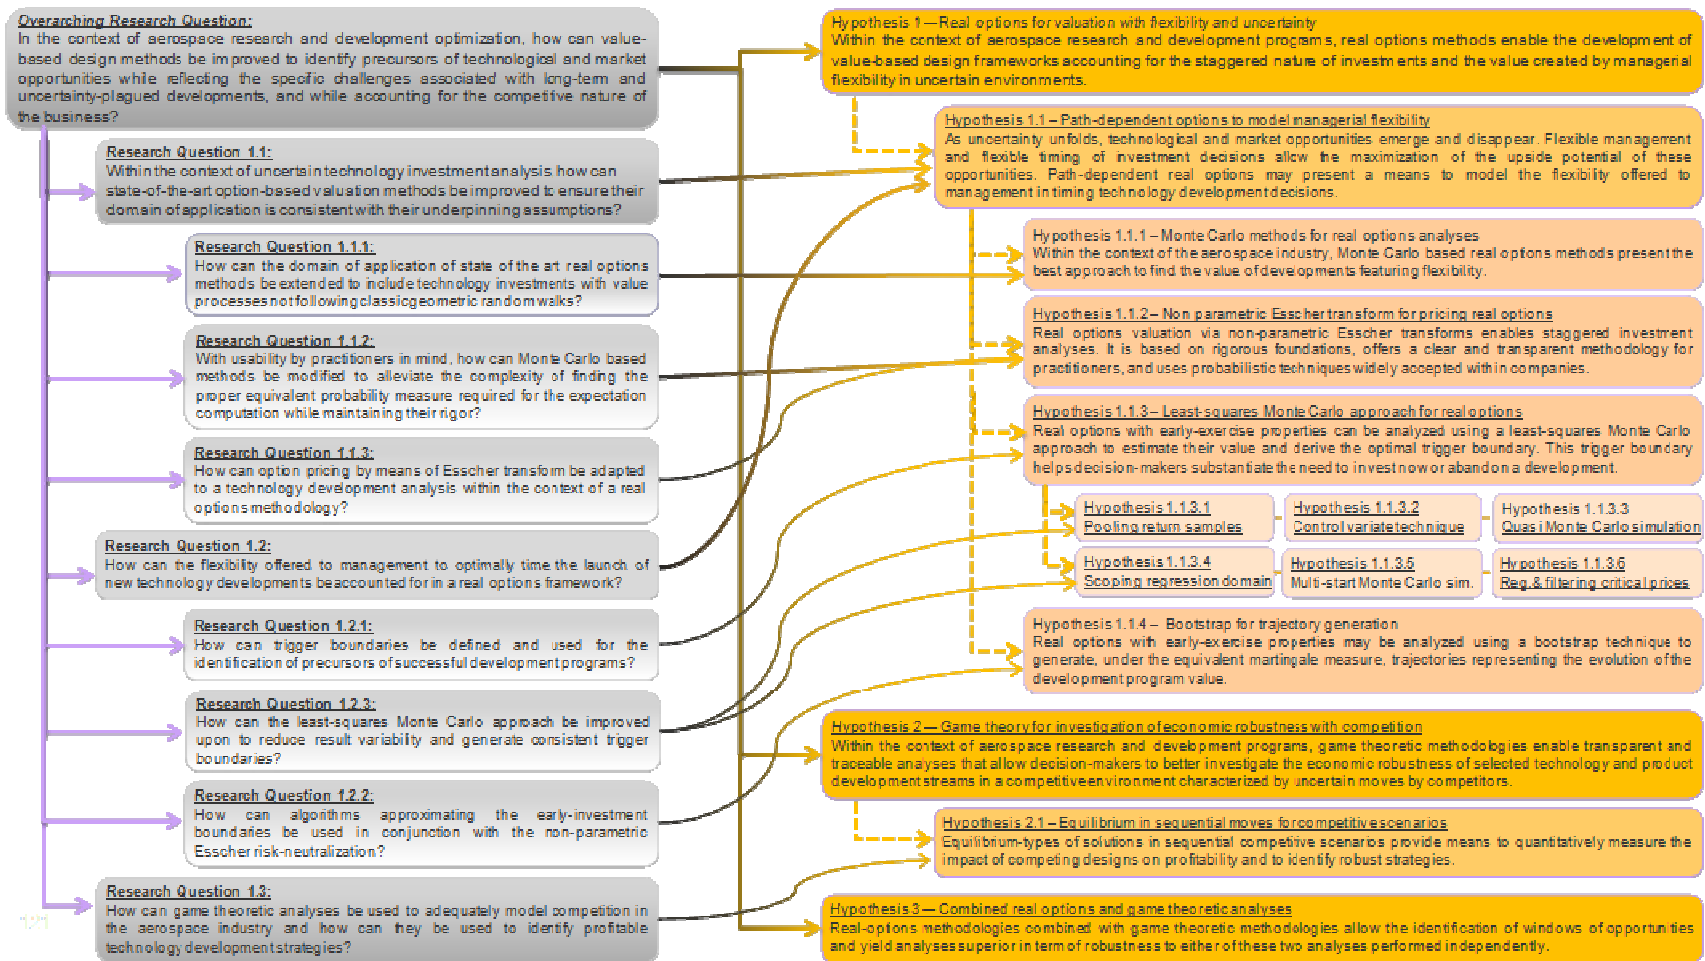


Figure 59: Updated set of research questions and hypotheses

7.3 Verification of technical hypotheses

According to the Project Management Institute [205], the purpose of the verification process is to evaluate “*whether or not a product, service, or system complies with a regulation, requirement, specification, or imposed condition.*” This is achieved by following the bottom-up approach of the definition-decomposition and verification-validation “V-model” previously described. The verification is performed on each technical and modeling hypothesis by studying them independently and testing them on canonical tests. These hypotheses are the foundation of the proposed methodology. Consequently, demonstrating in a crystal-clear manner that these individual pieces work as intended is critical to get the buy-in from practitioners. First, the non-parametric Esscher transformation implementation is checked to verify that the risk-neutralization process is working correctly for commonly used stochastic processes in real options analyses. Next, the implementation of the bootstrap technique is checked to ensure that it can be used for resampling purposes. The implementation of the least-squares Monte Carlo options pricing technique is subsequently checked to verify that real options featuring early-exercise opportunities can be properly evaluated. Finally, once the individual verification of these three implementations is achieved, the verification of the *ensemble* is performed to ensure that the proposed Monte Carlo based method for the evaluation of real options works properly.

7.3.1 Non-parametric Esscher transform technique

The Esscher transform and its non-parametric empirical approximation have been presented as a means to perform the risk-neutralization step required to evaluate real

options. In a real options framework, the practitioner typically starts by calibrating a stochastic process under the historical probability measure and then generates trajectories representing the possible evolutions over time. To price options, practitioners must nevertheless use trajectories under an equivalent risk-neutral measure to estimate the expected option payoff which is then discounted back to the present time at the risk-free rate of return. The Esscher transformation and its non-parametric approximation were introduced as a technical solution to perform this change of probability measure.

Research Question 1.1.2 – Improving Monte Carlo methods for real options

With usability by practitioners in mind, how can Monte Carlo methods be modified to alleviate the complexity of finding the proper equivalent probability measure?

Research Question 1.1.3 – Adaptation of Esscher transform for real options

How can option pricing by means of Esscher transform be adapted to a corporate investment analysis within the context of a real options methodology?

Hypothesis 1.1.2 – Non-parametric Esscher transform for pricing real options

Real options valuation via non-parametric Esscher transforms is a promising framework for staggered investment analyses. It is based on rigorous foundations, offers a clear and transparent methodology for practitioners and uses probabilistic techniques widely accepted within companies.

The purpose of the non-parametric Esscher transformation is to transform an arbitrary distribution such that it exhibits risk-neutral properties. In doing so, the Esscher transform preserves the nature of some classes of stochastic processes which means for instance that a lognormal distribution remains a lognormal distribution. These induced

distributions are therefore investigated to check whether they satisfy the risk-neutrality and the type-preservation properties. The verification starts with a Monte Carlo simulation of the evolutions of primary uncertainties affecting the value of the underlying asset which is then simulated under the equivalent martingale measure using the non-parametric Esscher transform. This yields, at each time step of the simulation, distributions of both underlying asset values and underlying asset returns. The distribution of returns is compared to the known theoretical counterpart. Since one requirement for the proposed real options methodology is the ability to capture a complex reality featuring uncertainties following non-standard stochastic processes, the verification is performed for two completely different processes: a classic geometric Brownian motion (GBM) for which a single equivalent martingale measure exists and the Merton jump-diffusion process (JD) for which the equivalent martingale measure is not unique since the market is incomplete. However, the measure induced by the Esscher transformation leads to one specific combination of jump-diffusion parameters (i.e. new drift, jump arrival rates, and jump amplitudes) which are discussed in Schoutens [206].

Verification process and criteria for success

Several tests are performed to verify the implementation of the change of probability measure using non-parametric Esscher transformations. The first is a qualitative test that visually compares the empirical probability distribution induced by the non-parametric Esscher transformation to the known risk-neutral probability distribution. This test is considered successful if there is no apparent departure from the bisecting line in the Q-Q plot. The second test is quantitative and uses the one-sample Kolmogorov-Smirnov statistical test to confirm whether the empirical probability

distribution and the known risk-neutral probability distribution differ in any way. This test is considered successful if the equality of distributions cannot be rejected at a five percent level of significance. The third test is also quantitative and uses the z -test and the Student's t -test to confirm whether the means of these two distributions differ in any way. This test is considered successful if the equality of means cannot be rejected at a five percent level of significance. All these tests are performed for two popular stochastic processes: the geometric Brownian motion and the Merton jump-diffusion process.

Graphical tests


The first batch of tests deals with the geometric Brownian motion. The stochastic process is sampled to generate 80,000 trajectories representing 80,000 possible evolutions of an underlying asset value over time. At each time step, these values induce a sample of 80,000 returns. One sample is subsequently transformed (weighted) using the non-parametric Esscher transform to yield a new risk-neutral sample which is compared to the known theoretical equivalent risk-neutral distribution using a Q-Q plot. This exercise is repeated for different cases of geometric Brownian motions for which the risk-free rate of return is varied between 2% and 8%, the drift rate is varied between 5% and 20%, the dividend yield is varied between 0% and 15%, the volatility is varied between 20% and 40%, and finally, the time step is varied between four and eight days (90 time steps with a maturity varied between one and two years). The Q-Q plots corresponding to 20 different test cases are presented in Table 37 on pages 264 and 265.

A summary of the results is provided in Table 36. All Q-Q plots exhibit loci of quantiles aligned almost perfectly on the bisecting lines. Minor deviations appear on some Q-Q plots for extreme quantiles (usually the first five and last five quantiles in


several thousand quantiles) but these do not exhibit any specific pattern. Furthermore, upon repetition of the questionable test cases, the minor deviations do not re-appear leading to the conclusion that these are artifacts of simulations using imperfect random number generators.

Table 36: Summary of Q-Q plot tests with GBM


| | | | |
|--|--|--|--|
| GBM $r_f=0.02, \mu=0.05,$ $q=0.00, \sigma=0.20, T=1$ <i>Minor negative kurtosis</i> | GBM $r_f=0.02, \mu=0.05,$ $q=0.00, \sigma=0.40, T=1$ <i>Minor positive kurtosis</i> | GBM $r_f=0.02, \mu=0.20,$ $q=0.00, \sigma=0.20, T=1$ <i>On bisecting line</i> | GBM $r_f=0.02, \mu=0.20,$ $q=0.00, \sigma=0.40, T=1$ <i>On bisecting line</i> |
| GBM $r_f=0.02, \mu=0.05,$ $q=0.00, \sigma=0.20, T=2$ <i>On bisecting line</i> | GBM $r_f=0.02, \mu=0.05,$ $q=0.00, \sigma=0.40, T=2$ <i>On bisecting line</i> | GBM $r_f=0.02, \mu=0.20,$ $q=0.00, \sigma=0.20, T=2$ <i>Two minor bumps</i> | GBM $r_f=0.02, \mu=0.20,$ $q=0.00, \sigma=0.40, T=2$ <i>Minor positive kurtosis</i> |
| GBM $r_f=0.02, \mu=0.05,$ $q=0.05, \sigma=0.20, T=1$ <i>On bisecting line</i> | GBM $r_f=0.02, \mu=0.05,$ $q=0.15, \sigma=0.20, T=1$ <i>On bisecting line</i> | GBM $r_f=0.02, \mu=0.20,$ $q=0.05, \sigma=0.20, T=1$ <i>On bisecting line</i> | GBM $r_f=0.02, \mu=0.20,$ $q=0.15, \sigma=0.20, T=1$ <i>Minor positive kurtosis</i> |
| GBM $r_f=0.08, \mu=0.05,$ $q=0.00, \sigma=0.20, T=1$ <i>Minor positive kurtosis</i> | GBM $r_f=0.08, \mu=0.05,$ $q=0.00, \sigma=0.40, T=1$ <i>Minor positive kurtosis</i> | GBM $r_f=0.08, \mu=0.20,$ $q=0.00, \sigma=0.20, T=1$ <i>On bisecting line</i> | GBM $r_f=0.08, \mu=0.20,$ $q=0.00, \sigma=0.40, T=1$ <i>On bisecting line</i> |
| GBM $r_f=0.08, \mu=0.05,$ $q=0.00, \sigma=0.20, T=2$ <i>On bisecting line</i> | GBM $r_f=0.08, \mu=0.05,$ $q=0.00, \sigma=0.40, T=2$ <i>On bisecting line</i> | GBM $r_f=0.08, \mu=0.20,$ $q=0.00, \sigma=0.20, T=2$ <i>On bisecting line</i> | GBM $r_f=0.08, \mu=0.20,$ $q=0.00, \sigma=0.40, T=2$ <i>On bisecting line</i> |




$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20, T=1$




$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.40, T=1$




$r_f=0.02, \mu=0.20, q=0.00, \sigma=0.20, T=1$




$r_f=0.02, \mu=0.20, q=0.00, \sigma=0.40, T=1$




$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20, T=2$




$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.40, T=2$




$r_f=0.02, \mu=0.20, q=0.00, \sigma=0.20, T=2$




$r_f=0.02, \mu=0.20, q=0.00, \sigma=0.40, T=2$




$r_f=0.02, \mu=0.05, q=0.05, \sigma=0.20, T=1$




$r_f=0.02, \mu=0.05, q=0.15, \sigma=0.20, T=1$




$r_f=0.02, \mu=0.20, q=0.05, \sigma=0.20, T=1$




$r_f=0.02, \mu=0.20, q=0.15, \sigma=0.20, T=1$




$r_f=0.08, \mu=0.05, q=0.00, \sigma=0.20,$
T=1




$r_f=0.08, \mu=0.05, q=0.00, \sigma=0.40,$
T=1




$r_f=0.08, \mu=0.20, q=0.00, \sigma=0.20,$
T=1




$r_f=0.08, \mu=0.20, q=0.00, \sigma=0.40,$
T=1




$r_f=0.08, \mu=0.05, q=0.00, \sigma=0.20,$
T=2



$r_f=0.08, \mu=0.05, q=0.00, \sigma=0.40,$
T=2



$r_f=0.08, \mu=0.20, q=0.00, \sigma=0.20,$
T=2



$r_f=0.08, \mu=0.20, q=0.00, \sigma=0.40,$
T=2

r_f = riskless rate of interest; μ = diffusion statistical drift; σ = diffusion volatility; q = dividend yield; and T = simulation horizon (years)

Experiment parameters: time step number = 90; simulation number = 80,000; return sample pooling = 1

Table 37: Q-Q Plots for verification of the non-parametric Esscher transformation applied to geometric Brownian motions

The second batch of tests deals with the Merton jump-diffusion process. The stochastic process is sampled to generate 50,000 possible trajectories which are used to construct samples of 50,000 returns at each time step. Ten return samples are pooled together to yield a sample of 500,000 returns which are transformed using the non-parametric Esscher procedure. The resulting sample is then compared to the theoretical equivalent risk-neutral distribution.

In these tests, the Q-Q plots are graphed with fewer points than were obtained from the simulations in order to speed-up the verification process. A down-sampling of 50 which consists in ordering returns and then selecting only 10,000 returns (one every fifty returns) is first performed to graph the Q-Q plots. Indeed, the estimation of theoretical quantiles is a time consuming endeavor as the cumulative probability distribution of returns induced by a Merton jump-diffusion process is expressed as an infinite series and it must be inverted to find quantiles.


This exercise is repeated for different parameters with emphasis on varying parameters governing jumps: the jump arrival rate is varied between 200% and 800% while jump size volatility is varied between 20% and 40%. The risk-free rate of return is varied between 2% and 8%, the drift rate is varied between 5% and 20%, the dividend yield is varied between 0% and 5%, and the time step is varied between four and eight days (90 time steps with a maturity between one and two years). The resulting Q-Q plots are presented in Table 39 on pages 268 and 269 with a summary of findings highlighted in Table 38.

Q-Q plots are almost as good as in the previous case and small deviations are still observed. These deviations are mostly restricted to the extreme end of the tails and are


certainly due to the difficulty in capturing rare events (jumps) in simulations. Most of the Q-Q plots exhibit loci of quantiles well aligned on the bisecting lines and the minor deviations appearing in some plots do not exhibit specific patterns and are not repeatable.

Table 38: Summary of Q-Q plot tests with Merton jump diffusion model


| | | | |
|---|---|---|---|
| Merton JD $r_f=0.02, \mu=0.05,$ $q=0.00, \sigma=0.20, T=1,$ $\lambda=4.0, \gamma=-0.02, \delta=0.20$ <i>Minor negative kurtosis</i> | Merton JD $r_f=0.02, \mu=0.05,$ $q=0.00, \sigma=0.40, T=1,$ $\lambda=4.0, \gamma=-0.08, \delta=0.40$ <i>Minor deviation</i> | Merton JD $r_f=0.02, \mu=0.20,$ $q=0.00, \sigma=0.20, T=1,$ $\lambda=4.0, \gamma=-0.02, \delta=0.20$ <i>On bisecting line</i> | Merton JD $r_f=0.02, \mu=0.20,$ $q=0.00, \sigma=0.40, T=1,$ $\lambda=6.0, \gamma=-0.08, \delta=0.40$ <i>On bisecting line</i> |
| Merton JD $r_f=0.02, \mu=0.05,$ $q=0.00, \sigma=0.20, T=2,$ $\lambda=4.0, \gamma=-0.02, \delta=0.20$ <i>On bisecting line</i> | Merton JD $r_f=0.02, \mu=0.05,$ $q=0.00, \sigma=0.40, T=2,$ $\lambda=4.0, \gamma=-0.08, \delta=0.40$ <i>On bisecting line</i> | Merton JD $r_f=0.02, \mu=0.20,$ $q=0.00, \sigma=0.20, T=2,$ $\lambda=4.0, \gamma=-0.02, \delta=0.20$ <i>Two minor bumps</i> | Merton JD $r_f=0.02, \mu=0.20,$ $q=0.00, \sigma=0.40, T=2,$ $\lambda=6.0, \gamma=-0.08, \delta=0.40$ <i>Minor positive kurtosis</i> |
| Merton JD $r_f=0.02, \mu=0.05,$ $q=0.05, \sigma=0.20, T=1,$ $\lambda=4.0, \gamma=-0.02, \delta=0.20$ <i>On bisecting line</i> | Merton JD $r_f=0.02, \mu=0.05,$ $q=0.15, \sigma=0.20, T=1,$ $\lambda=4.0, \gamma=-0.08, \delta=0.40$ <i>On bisecting line</i> | Merton JD $r_f=0.02, \mu=0.20,$ $q=0.05, \sigma=0.20, T=1,$ $\lambda=4.0, \gamma=-0.02, \delta=0.20$ <i>On bisecting line</i> | Merton JD $r_f=0.02, \mu=0.20,$ $q=0.15, \sigma=0.20, T=1,$ $\lambda=6.0, \gamma=-0.08, \delta=0.40$ <i>Minor positive kurtosis</i> |
| Merton JD $r_f=0.08, \mu=0.05,$ $q=0.00, \sigma=0.20, T=1,$ $\lambda=4.0, \gamma=-0.02, \delta=0.20$ <i>On bisecting line</i> | Merton JD $r_f=0.08, \mu=0.05,$ $q=0.00, \sigma=0.40, T=1,$ $\lambda=4.0, \gamma=-0.08, \delta=0.40$ <i>Minor positive kurtosis</i> | Merton JD $r_f=0.08, \mu=0.20,$ $q=0.00, \sigma=0.20, T=1,$ $\lambda=4.0, \gamma=-0.02, \delta=0.20$ <i>On bisecting line</i> | Merton JD $r_f=0.08, \mu=0.20,$ $q=0.00, \sigma=0.40, T=1,$ $\lambda=6.0, \gamma=-0.08, \delta=0.40$ <i>On bisecting line</i> |
| Merton JD $r_f=0.08, \mu=0.05,$ $q=0.00, \sigma=0.20, T=2,$ $\lambda=4.0, \gamma=-0.02, \delta=0.20$ <i>On bisecting line</i> | Merton JD $r_f=0.08, \mu=0.05,$ $q=0.00, \sigma=0.40, T=2,$ $\lambda=4.0, \gamma=-0.08, \delta=0.40$ <i>On bisecting line</i> | Merton JD $r_f=0.08, \mu=0.20,$ $q=0.00, \sigma=0.20, T=2,$ $\lambda=8.0, \gamma=-0.02, \delta=0.20$ <i>On bisecting line</i> | Merton JD $r_f=0.08, \mu=0.20,$ $q=0.00, \sigma=0.40, T=2,$ $\lambda=8.0, \gamma=-0.08, \delta=0.40$ <i>On bisecting line</i> |




$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20,$
 $\lambda=4.00, \gamma=-0.02, \delta=0.20, T=1$




$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20,$
 $\lambda=4.00, \gamma=-0.08, \delta=0.40, T=1$




$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20,$
 $\lambda=6.00, \gamma=-0.02, \delta=0.20, T=1$




$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20,$
 $\lambda=6.00, \gamma=-0.08, \delta=0.40, T=1$




$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20,$
 $\lambda=4.00, \gamma=-0.02, \delta=0.20, T=2$




$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20,$
 $\lambda=4.00, \gamma=-0.08, \delta=0.40, T=2$




$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20,$
 $\lambda=6.00, \gamma=-0.02, \delta=0.20, T=2$




$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20,$
 $\lambda=6.00, \gamma=-0.08, \delta=0.40, T=2$




$r_f=0.02, \mu=0.15, q=0.00, \sigma=0.20,$
 $\lambda=4.00, \gamma=-0.02, \delta=0.20, T=1$




$r_f=0.02, \mu=0.15, q=0.00, \sigma=0.20,$
 $\lambda=4.00, \gamma=-0.08, \delta=0.40, T=1$




$r_f=0.02, \mu=0.15, q=0.00, \sigma=0.20,$
 $\lambda=6.00, \gamma=-0.02, \delta=0.20, T=1$




$r_f=0.02, \mu=0.15, q=0.00, \sigma=0.20,$
 $\lambda=6.00, \gamma=-0.08, \delta=0.40, T=1$




$r_f=0.02, \mu=0.15, q=0.00, \sigma=0.20,$
 $\lambda=4.00, \gamma=-0.02, \delta=0.20, T=2$




$r_f=0.02, \mu=0.15, q=0.00, \sigma=0.20,$
 $\lambda=4.00, \gamma=-0.08, \delta=0.40, T=2$




$r_f=0.02, \mu=0.15, q=0.00, \sigma=0.20,$
 $\lambda=6.00, \gamma=-0.02, \delta=0.20, T=2$




$r_f=0.02, \mu=0.15, q=0.00, \sigma=0.20,$
 $\lambda=6.00, \gamma=-0.08, \delta=0.40, T=2$




$r_f=0.08, \mu=0.20, q=0.05, \sigma=0.40,$
 $\lambda=4.00, \gamma=-0.02, \delta=0.20, T=1$



$r_f=0.08, \mu=0.20, q=0.05, \sigma=0.40,$
 $\lambda=4.00, \gamma=-0.08, \delta=0.40, T=1$



$r_f=0.08, \mu=0.20, q=0.05, \sigma=0.40,$
 $\lambda=8.00, \gamma=-0.02, \delta=0.20, T=1$



$r_f=0.08, \mu=0.20, q=0.05, \sigma=0.40,$
 $\lambda=8.00, \gamma=-0.08, \delta=0.40, T=1$

r_f = riskless rate of interest; μ = diffusion statistical drift; σ = diffusion volatility; q = dividend yield; λ = arrival rate of jumps (per year);
 γ = mean amplitude of jumps; δ = volatility of jump amplitude; and T = simulation horizon (years)

Experiment parameters: time step number = 90; simulation number = 50,000; return sample pooling = 10; down-sampling factor = 50

Table 39: Q-Q Plots for verification of the non-parametric Esscher transformation applied to Merton jump diffusion processes

Statistical tests – Kolmogorov-Smirnov test

If Q-Q plots are helpful to qualitatively compare experimental samples and reference distributions, they do not quantify whether the observed deviations are statistically significant. Thus, more testing is warranted to quantify the observed departures from the bisecting lines and to perform statistical testing. For each test case, the observed samples are compared to the corresponding theoretical risk-neutral distribution using the Kolmogorov-Smirnov test. The null hypothesis for the Kolmogorov-Smirnov test is that the observed samples induced by simulations and non-parametric Esscher transformations are drawn from the known theoretical distributions. The Kolmogorov-Smirnov test yields test statistics (and p -values) that can be compared to critical values in order to assess the likelihood of observing such difference between the sample and the reference distribution.

The first batch of tests deals with geometric Brownian motions using different combinations of parameters (risk free rate, drift, volatility, dividend yield and maturities) representative of what would be used for the pricing of real options. The risk-free rate of return is varied between 2% and 8%, the drift rate is varied between 5% and 20%, the dividend yield is varied between 0% and 15%, the volatility is varied between 20% and 40%, and finally, the time step is varied between four and eight days (90 time steps with a maturity varied between one and two years). The simulations are run for 80,000 different paths leading to a sample of 80,000 returns at each time step. For twenty different test cases reported in Table 40, the one-sample Kolmogorov-Smirnov test statistic is computed as well as the corresponding p -value under the null hypothesis (p -value represents the likelihood of having a result equal or more extreme than what is

currently observed given that the null hypothesis is correct). The p -values for the test cases are very large and all of them are much greater than the 5% level of significance chosen for this test. As a result, it is not possible to reject the null hypothesis at the 5% level of significance.

Table 40: Kolmogorov-Smirnov test for twenty combinations of parameters

| r_f | μ | q | σ | T | Kolmogorov-Smirnov test statistic | Kolmogorov-Smirnov test critical value | Kolmogorov-Smirnov test p -value |
|-------|-------|-------|----------|------|-----------------------------------|--|------------------------------------|
| 2.0% | 5.0% | 0.0% | 20.0% | 1.00 | 0.532 | 1.358 | 94.0% |
| 2.0% | 20.0% | 0.0% | 20.0% | 1.00 | 0.498 | 1.358 | 96.5% |
| 2.0% | 5.0% | 0.0% | 40.0% | 1.00 | 0.469 | 1.358 | 98.0% |
| 2.0% | 20.0% | 0.0% | 40.0% | 1.00 | 0.662 | 1.358 | 77.4% |
| 8.0% | 5.0% | 0.0% | 20.0% | 1.00 | 0.494 | 1.358 | 96.8% |
| 8.0% | 20.0% | 0.0% | 20.0% | 1.00 | 0.749 | 1.358 | 62.9% |
| 8.0% | 5.0% | 0.0% | 40.0% | 1.00 | 0.759 | 1.358 | 61.3% |
| 8.0% | 20.0% | 0.0% | 40.0% | 1.00 | 0.438 | 1.358 | 99.1% |
| 2.0% | 5.0% | 0.0% | 20.0% | 2.00 | 0.586 | 1.358 | 88.2% |
| 2.0% | 20.0% | 0.0% | 20.0% | 2.00 | 0.616 | 1.358 | 84.3% |
| 2.0% | 5.0% | 0.0% | 40.0% | 2.00 | 0.499 | 1.358 | 96.4% |
| 2.0% | 20.0% | 0.0% | 40.0% | 2.00 | 0.465 | 1.358 | 98.2% |
| 8.0% | 5.0% | 0.0% | 20.0% | 2.00 | 0.588 | 1.358 | 88.0% |
| 8.0% | 20.0% | 0.0% | 20.0% | 2.00 | 0.560 | 1.358 | 91.3% |
| 8.0% | 5.0% | 0.0% | 40.0% | 2.00 | 0.419 | 1.358 | 99.5% |
| 8.0% | 20.0% | 0.0% | 40.0% | 2.00 | 0.597 | 1.358 | 86.8% |
| 2.0% | 5.0% | 5.0% | 20.0% | 1.00 | 0.567 | 1.358 | 90.5% |
| 2.0% | 20.0% | 5.0% | 20.0% | 1.00 | 0.597 | 1.358 | 86.8% |
| 2.0% | 5.0% | 15.0% | 20.0% | 1.00 | 0.538 | 1.358 | 93.4% |
| 2.0% | 20.0% | 15.0% | 20.0% | 1.00 | 0.504 | 1.358 | 96.2% |

r_f = risk-free rate; μ = drift of GBM; σ = volatility of GBM; q = dividend yield; T = simulation horizon
Experiment parameters: time step number = 90; simulation number = 80,000; return sample pooling = 1

Monte Carlo simulations introduce some variability as new pseudo-random number sequences are used for the generation of trajectories. If the tests were repeated, this could possibly lead to contradicting conclusions to the Kolmogorov-Smirnov tests. To check the robustness of these conclusions, each of the twenty test cases is repeated

thirty times to assess how robust the outcomes of the hypothesis testing are. The p -values are computed for each trial and a synthetic view representing the experimental distribution of p -values is reported in Figure 60. Interestingly, the experimental distribution is not uniform and is in fact skewed towards the larger p -values (right skewed as shown). Of those six hundred trials, a single one has a p -value below the 5% level of significance which represents less than 0.2% of all trials. Furthermore, 95% of trials have p -values above 30% and therefore are far-away from the critical area. At the 5% level of significance, there is thus no reason to reject the null hypothesis that the experimental return samples are drawn from the theoretical risk-neutral distribution.




Figure 60: Distribution of p -values for 600 Kolmogorov-Smirnov tests (geometric Brownian motions)

The second batch of tests deals with the return distributions induced by Merton jump diffusion processes. Table 41 corresponds to twenty test cases with emphasis put on varying parameters governing jumps: jump arrival rate is varied between 200% and 800% while jump size volatility is varied between 20% and 40%. The risk-free rate of return is varied between 2% and 8%, the drift rate is varied between 5% and 20%, the dividend yield is varied between 0% and 5%, and the time step is varied between four and eight

days. The simulations are run for 50,000 different paths leading to samples of 50,000 returns at each time step. Ten different samples of returns are pooled together and weighted using the non-parametric Esscher transform so as to construct a new sample of 500,000 returns. This sample is reduced in size by ordering and then selecting one return every 50 returns to speed-up the verification process. A Kolmogorov-Smirnov test is then carried out with this subsample of the risk-neutral experimental sample. The null hypothesis for the Kolmogorov-Smirnov test is that the observed sample induced by simulations and non-parametric Esscher transformations is drawn from the known theoretical distributions. All p -values are substantially above the 5% significance level retained for the test and consequently the null hypothesis cannot be rejected.

Table 41: Kolmogorov-Smirnov statistical tests for twenty cases of Merton Jump diffusion process

| r_f | μ | q | σ | λ | δ | T | Kolmogorov Smirnov statistic | p -value |
|-------|-------|------|----------|-----------|----------|-----|------------------------------|------------|
| 2.0% | 5% | 0.0% | 20% | 400% | 20% | 1.0 | 0.717 | 68% |
| 2.0% | 5% | 0.0% | 20% | 600% | 20% | 1.0 | 0.619 | 84% |
| 2.0% | 5% | 0.0% | 20% | 400% | 40% | 1.0 | 0.560 | 91% |
| 2.0% | 5% | 0.0% | 20% | 600% | 40% | 1.0 | 0.610 | 85% |
| 2.0% | 15% | 0.0% | 20% | 400% | 20% | 1.0 | 0.866 | 44% |
| 2.0% | 15% | 0.0% | 20% | 600% | 20% | 1.0 | 0.687 | 73% |
| 2.0% | 15% | 0.0% | 20% | 400% | 40% | 1.0 | 0.882 | 42% |
| 2.0% | 15% | 0.0% | 20% | 600% | 40% | 1.0 | 0.647 | 80% |
| 2.0% | 5% | 0.0% | 20% | 400% | 20% | 2.0 | 0.916 | 37% |
| 2.0% | 5% | 0.0% | 20% | 600% | 20% | 2.0 | 0.848 | 47% |
| 2.0% | 5% | 0.0% | 20% | 400% | 40% | 2.0 | 0.716 | 68% |
| 2.0% | 5% | 0.0% | 20% | 600% | 40% | 2.0 | 0.600 | 86% |
| 2.0% | 15% | 0.0% | 20% | 400% | 20% | 2.0 | 0.860 | 45% |
| 2.0% | 15% | 0.0% | 20% | 600% | 20% | 2.0 | 0.551 | 92% |
| 2.0% | 15% | 0.0% | 20% | 400% | 40% | 2.0 | 1.101 | 18% |
| 2.0% | 15% | 0.0% | 20% | 600% | 40% | 2.0 | 0.567 | 90% |
| 8.0% | 20% | 5.0% | 40% | 400% | 20% | 1.0 | 1.006 | 26% |
| 8.0% | 20% | 5.0% | 40% | 800% | 20% | 1.0 | 0.837 | 49% |

Table 41 Continued

| | | | | | | | | |
|------|-----|------|-----|------|-----|-----|-------|-----|
| 8.0% | 20% | 5.0% | 40% | 400% | 40% | 1.0 | 0.833 | 49% |
| 8.0% | 20% | 5.0% | 40% | 800% | 40% | 1.0 | 0.638 | 81% |

r_f = riskless rate of interest; μ = diffusion statistical drift; σ = diffusion volatility; q = dividend yield;
 λ = arrival rate of jumps (per year); $\gamma = -\sigma^2/2$ = jump amplitude; δ = volatility of jump amplitude;
 T = simulation horizon (years)
 Experiment parameters: time step number = 90; simulation number = 50,000; return sample pooling = 10;
 down-sampling factor = 50

As previously done, each of these twenty test cases is now repeated thirty times in order to assess the sensitivity and robustness of the conclusions to the Kolmogorov-Smirnov tests. This yields an experiment with a grand total of six hundred trials for which the experimental distribution of p -values is reported in Figure 61. It is again skewed towards the larger p -values (right skewed as shown) and a single case has a p -value below the 5% level of significance. This represents less than 0.2% of all trials and as a consequence, the null hypothesis cannot be rejected at the 5% level of significance.




Figure 61: Distribution of p -values for 600 Kolmogorov-Smirnov tests (Merton jump diffusion processes)

Statistical tests – Testing the mean with z -tests and t -tests

Switching to the first moment of the risk-neutral return distributions, the mean values of samples of returns induced by simulation and subsequent change of measure with the non-parametric Esscher transform are compared to the expected values of the corresponding known risk-neutral distributions. Statistical testing is performed by repeating each test case thirty times yielding an experimental average mean and an experimental standard error which are then used to carry out two-tailed z -tests and t -tests. The null hypothesis for the z -tests and t -tests is the equality of the experimental average mean and the expected value of the known risk neutral distribution.

The first batch of experiments deals with geometric Brownian motions using different combinations of parameters (risk free rate, drift, volatility, dividend yield and maturities) representative of what would be used for the pricing of real options. The risk-free rate of return is varied between 2% and 8%, the drift rate is varied between 5% and 20%, the dividend yield is varied between 0% and 15%, the volatility is varied between 20% and 40%, and finally, the time step is varied between four and eight days (90 time steps with a maturity varied between one and two years). The simulations are run for 80,000 different paths leading to a sample of 80,000 returns at each time step. For twenty test cases of geometric Brownian motions, Table 42 reports the outcomes of the two-tailed z -tests and t -tests. Approximate z -statistics and t -statistics are computed, and corresponding p -values are derived. All p -value estimates are above the 5% significance level retained for the test. Consequently, the null hypothesis cannot be rejected at the 5% level of significance. As expected, the two-tailed t -tests and the z -tests yield similar

approximations of the p -values since both tests are applicable “in the limit” (due to the large sample size).

Table 42: z -tests and t -tests for the mean returns of twenty cases of geometric Brownian motions

| r_f | μ | q | σ | T | Experimental Sample Mean Return | Experimental Sample Standard Error | Theoretical Mean Return | z -test and t -test statistic | z -test p -value | t -test p -value |
|-------|-------|-----|----------|---|---------------------------------|------------------------------------|-------------------------|-----------------------------------|----------------------|----------------------|
| 2% | 5% | 0% | 20% | 1 | -2.02E-07 | 1.45E-07 | -3.85E-20 | 1.393 | 16% | 17% |
| 2% | 20% | 0% | 20% | 1 | 1.57E-07 | 1.10E-07 | -3.85E-20 | 1.436 | 15% | 16% |
| 2% | 5% | 0% | 40% | 1 | -6.67E-04 | 5.73E-07 | -6.67E-04 | 0.163 | 87% | 87% |
| 2% | 20% | 0% | 40% | 1 | -6.67E-04 | 6.41E-07 | -6.67E-04 | 0.252 | 80% | 80% |
| 8% | 5% | 0% | 20% | 1 | 6.67E-04 | 1.30E-07 | 6.67E-04 | 0.076 | 94% | 94% |
| 8% | 20% | 0% | 20% | 1 | 6.67E-04 | 1.67E-07 | 6.67E-04 | 0.873 | 38% | 39% |
| 8% | 5% | 0% | 40% | 1 | 2.84E-07 | 5.88E-07 | -1.54E-19 | 0.484 | 63% | 63% |
| 8% | 20% | 0% | 40% | 1 | -6.27E-07 | 6.12E-07 | -1.54E-19 | 1.025 | 31% | 31% |
| 2% | 5% | 0% | 20% | 2 | 1.54E-07 | 1.93E-07 | -7.71E-20 | 0.801 | 42% | 43% |
| 2% | 20% | 0% | 20% | 2 | 1.24E-07 | 2.58E-07 | -7.71E-20 | 0.482 | 63% | 63% |
| 2% | 5% | 0% | 40% | 2 | -1.33E-03 | 1.37E-06 | -1.33E-03 | 0.512 | 61% | 61% |
| 2% | 20% | 0% | 40% | 2 | -1.33E-03 | 1.03E-06 | -1.33E-03 | 0.615 | 54% | 54% |
| 8% | 5% | 0% | 20% | 2 | 1.33E-03 | 3.20E-07 | 1.33E-03 | 1.156 | 25% | 26% |
| 8% | 20% | 0% | 20% | 2 | 1.33E-03 | 2.64E-07 | 1.33E-03 | 0.668 | 50% | 51% |
| 8% | 5% | 0% | 40% | 2 | -1.12E-06 | 1.05E-06 | -3.08E-19 | 1.063 | 29% | 30% |
| 8% | 20% | 0% | 40% | 2 | 5.36E-07 | 9.04E-07 | -3.08E-19 | 0.593 | 55% | 56% |
| 2% | 5% | 5% | 20% | 1 | -5.55E-04 | 1.43E-07 | -5.56E-04 | 0.659 | 51% | 52% |
| 2% | 20% | 5% | 20% | 1 | -5.56E-04 | 1.52E-07 | -5.56E-04 | 1.447 | 15% | 16% |
| 2% | 5% | 15% | 20% | 1 | -1.67E-03 | 1.27E-07 | -1.67E-03 | 0.670 | 50% | 51% |
| 2% | 20% | 15% | 20% | 1 | -1.67E-03 | 1.50E-07 | -1.67E-03 | 0.563 | 57% | 58% |

r_f = riskless rate of interest; μ = diffusion statistical drift; σ = diffusion volatility;
 q = dividend yield; T = simulation horizon (years)
 Experiment parameters: time step number = 90; simulation number = 80,000; return sample pooling = 1;
 test case repetition = 30

The second batch of experiments deals with the return distributions induced by Merton jump diffusion processes. Table 43 corresponds to twenty test cases with emphasis put on varying parameters governing jumps: jump arrival rate is varied between 200% and 800% while jump size volatility is varied between 20% and 40%. The risk-free rate of return is varied between 2% and 8%, the drift rate is varied between 5% and 20%,

the dividend yield is varied between 0% and 5%, and the time step is varied between four and eight days. The simulations are run for 50,000 different paths leading to a sample of 50,000 returns at each time step. Ten different samples of returns are pooled together and weighted using the non-parametric Esscher transform to construct a new sample of 500,000 returns. Approximate z -statistics and t -statistics are computed, and corresponding p -values are derived.

Table 43: z-tests and t-tests for the mean returns of twenty cases of Merton jump diffusion processes

| r_f | μ | q | σ | λ | δ | T | Exp. Sample Mean Return | Exp. Sample Standard Error | Theo. Mean Return | z -test and t -test statistics | z -test p -value | t -test p -value |
|-------|-------|-----|----------|-----------|----------|---|-------------------------|----------------------------|-------------------|------------------------------------|----------------------|----------------------|
| 2% | 5% | 0% | 20% | 400% | 20% | 1 | -8.92E-04 | 2.05E-06 | -8.93E-04 | 0.304 | 76% | 76% |
| 2% | 5% | 0% | 20% | 600% | 20% | 1 | -1.34E-03 | 2.69E-06 | -1.34E-03 | 0.551 | 58% | 59% |
| 2% | 5% | 0% | 20% | 400% | 40% | 1 | -3.57E-03 | 7.52E-06 | -3.57E-03 | 0.528 | 60% | 60% |
| 2% | 5% | 0% | 20% | 600% | 40% | 1 | -5.35E-03 | 9.79E-06 | -5.35E-03 | 0.066 | 95% | 95% |
| 2% | 15% | 0% | 20% | 400% | 20% | 1 | -9.22E-04 | 1.65E-06 | -9.23E-04 | 0.750 | 45% | 46% |
| 2% | 15% | 0% | 20% | 600% | 20% | 1 | -1.36E-03 | 2.63E-06 | -1.36E-03 | 0.290 | 77% | 77% |
| 2% | 15% | 0% | 20% | 400% | 40% | 1 | -3.63E-03 | 5.83E-06 | -3.64E-03 | 1.470 | 14% | 15% |
| 2% | 15% | 0% | 20% | 600% | 40% | 1 | -5.42E-03 | 9.51E-06 | -5.41E-03 | 0.772 | 44% | 45% |
| 2% | 5% | 0% | 20% | 400% | 20% | 2 | -1.78E-03 | 2.54E-06 | -1.79E-03 | 0.965 | 33% | 34% |
| 2% | 5% | 0% | 20% | 600% | 20% | 2 | -2.67E-03 | 3.17E-06 | -2.67E-03 | 0.068 | 95% | 95% |
| 2% | 5% | 0% | 20% | 400% | 40% | 2 | -7.15E-03 | 9.05E-06 | -7.14E-03 | 0.970 | 33% | 34% |
| 2% | 5% | 0% | 20% | 600% | 40% | 2 | -1.07E-02 | 1.13E-05 | -1.07E-02 | 0.315 | 75% | 76% |
| 2% | 15% | 0% | 20% | 400% | 20% | 2 | -1.84E-03 | 3.02E-06 | -1.85E-03 | 0.359 | 72% | 72% |
| 2% | 15% | 0% | 20% | 600% | 20% | 2 | -2.73E-03 | 2.77E-06 | -2.73E-03 | 0.097 | 92% | 92% |
| 2% | 15% | 0% | 20% | 400% | 40% | 2 | -7.28E-03 | 1.27E-05 | -7.28E-03 | 0.096 | 92% | 92% |
| 2% | 15% | 0% | 20% | 600% | 40% | 2 | -1.08E-02 | 1.17E-05 | -1.08E-02 | 0.092 | 93% | 93% |
| 8% | 20% | 5% | 40% | 400% | 20% | 1 | -1.46E-03 | 2.06E-06 | -1.46E-03 | 0.508 | 61% | 62% |
| 8% | 20% | 5% | 40% | 800% | 20% | 1 | -2.35E-03 | 2.42E-06 | -2.35E-03 | 1.312 | 19% | 20% |
| 8% | 20% | 5% | 40% | 400% | 40% | 1 | -4.17E-03 | 1.00E-05 | -4.17E-03 | 0.542 | 59% | 59% |
| 8% | 20% | 5% | 40% | 800% | 40% | 1 | -7.73E-03 | 1.07E-05 | -7.73E-03 | 0.098 | 92% | 92% |

r_f = riskless rate of interest; μ = diffusion statistical drift; σ = diffusion volatility; q = dividend yield;
 λ = arrival rate of jumps (per year); $\gamma = -\sigma^2/2$ = jump amplitude; δ = volatility of jump amplitude;
T = simulation horizon (years)
Experiment parameters: time step number = 90; simulation number = 50,000; return sample pooling = 10;
test case repetition = 30

All p -value estimates are above the 5% significance level retained for the test. Consequently, the null hypothesis cannot be rejected at the 5% level of significance. Again, the t -tests and the z -tests yield similar approximations of the p -values.

7.3.2 Bootstrapping technique for resampling

The bootstrap technique has been presented as a means to resample a weighted sample leading to a non-weighted sample. In a real options framework, the practitioner typically starts by simulating trajectories for the investment value representing its possible evolutions over time. Risk-neutralization of these trajectories using the non-parametric Esscher transform introduces weighting vectors at each time cross-section of the trajectories. This results in locally-weighted trajectories which prevents the computation of conditional expectation regressions along these trajectories. In turn, this prevents the use of the least-squares Monte Carlo algorithm for option pricing purposes. The bootstrap technique was introduced as an enabling technique to use the output of the non-parametric Esscher transform in a least-squares Monte Carlo option valuation algorithm.

Research Question 1.2.2 – Generating trajectories for finding investment boundaries

How can algorithms approximating the early-investment boundaries be used in conjunction with the non-parametric Esscher risk-neutralization?

Hypothesis 1.1.4 – Bootstrapping for trajectory generation

Real options with early-exercise properties may be analyzed using a bootstrapping technique to generate risk-neutral trajectories for the evolution of the research and development program values via simulations and resampling

The purpose of bootstrapping is to construct a new sample from an already existing sample featuring desirable properties. In doing so, the bootstrap preserves the characteristics as well as the nature of the original distributions. The induced distributions are therefore investigated to check whether they also satisfy the risk-neutrality and the type-preservation properties.

Verification process and criteria for success

Similarly to what was done for the verification of the non-parametric Esscher transformation, several tests are performed to verify the implementation of the bootstrapping. The first is a qualitative test that visually compares the empirical probability distribution induced by the resampling to the empirical probability distribution used for the resampling: in other words, the output of the procedure which is a non-weighted sample of returns is compared to the input which is a weighted sample of returns. This test is considered successful if there is no apparent departure from the bisecting line in the corresponding Q-Q plot. The second test is quantitative and uses the two-sample Kolmogorov-Smirnov statistical test to confirm whether the input (weighted) and output (non-weighted) empirical distributions differ. This test is considered successful if the equality of distributions cannot be rejected at a five percent level of significance. All these tests are performed for two popular stochastic processes: the geometric Brownian motion and the Merton jump-diffusion process.

Graphical tests


The first batch of tests deals with the geometric Brownian motion. The stochastic process is sampled to generate 80,000 trajectories representing the possible evolution of a

development program over time. These trajectories induce samples of 80,000 returns at each time step. Four of these samples are pooled together to yield a larger sample of 320,000 returns and a change of probability measure is performed using the non-parametric Esscher transform. This new sample of 320,000 weighted returns is bootstrapped using the sampling wheel algorithm to yield a new experimental sample of 20,000 non-weighted returns. The resulting empirical distribution made of 20,000 returns is compared to the non-resampled empirical distribution made of 320,000 weighted returns and the exercise is repeated for different test cases. Table 45 on page 282 and page 283 displays Q-Q plots corresponding to these test cases with the risk-free rate of return varied between 2% and 8%, the drift rate varied between 5% and 20%, the dividend yield varied between 0% and 15%, the volatility varied between 20% and 40%, and finally the time step varied between four and eight days.


A summary of the results is presented in Table 44 and all but one Q-Q plot exhibit loci of quantiles aligned almost perfectly on the bisecting lines. While minor deviations appear on some Q-Q plots for extreme quantiles, one of them exhibits some *notable* excess kurtosis ($r_f=0.08$, $\mu=0.05$, $q=0.00$, $\sigma=0.40$, $T=1$). Performing the test again, this excess kurtosis disappears leading to the conclusion that this is probably an artifact of the simulation using an imperfect pseudo-random number generator.

Table 44: Summary of the bootstrap resampling tests using Q-Q plots with GBM


| | | | |
|---|---|---|---|
| <p>GBM $r_f=0.02, \mu=0.05,$ $q=0.00, \sigma=0.20, T=1$ <i>On bisecting line</i></p> | <p>GBM $r_f=0.02, \mu=0.05,$ $q=0.00, \sigma=0.40, T=1$ <i>On bisecting line</i></p> | <p>GBM $r_f=0.02, \mu=0.20,$ $q=0.00, \sigma=0.20, T=1$ <i>Minor deviation</i></p> | <p>GBM $r_f=0.02, \mu=0.20,$ $q=0.00, \sigma=0.40, T=1$ <i>On bisecting line</i></p> |
| <p>GBM $r_f=0.02, \mu=0.05,$ $q=0.00, \sigma=0.20, T=2$ <i>On bisecting line</i></p> | <p>GBM $r_f=0.02, \mu=0.05,$ $q=0.00, \sigma=0.40, T=2$ <i>On bisecting line</i></p> | <p>GBM $r_f=0.02, \mu=0.20,$ $q=0.00, \sigma=0.20, T=2$ <i>On bisecting line</i></p> | <p>GBM $r_f=0.02, \mu=0.20,$ $q=0.00, \sigma=0.40, T=2$ <i>On bisecting line</i></p> |
| <p>GBM $r_f=0.02, \mu=0.05,$ $q=0.05, \sigma=0.20, T=1$ <i>On bisecting line</i></p> | <p>GBM $r_f=0.02, \mu=0.05,$ $q=0.15, \sigma=0.20, T=1$ <i>Minor deviation</i></p> | <p>GBM $r_f=0.02, \mu=0.20,$ $q=0.05, \sigma=0.20, T=1$ <i>On bisecting line</i></p> | <p>GBM $r_f=0.02, \mu=0.20,$ $q=0.15, \sigma=0.20, T=1$ <i>Minor deviation</i></p> |
| <p>GBM $r_f=0.08, \mu=0.05,$ $q=0.00, \sigma=0.20, T=1$ <i>On bisecting line</i></p> | <p>GBM $r_f=0.08, \mu=0.05,$ $q=0.00, \sigma=0.40, T=1$ <i>Notable excess kurtosis</i></p> | <p>GBM $r_f=0.08, \mu=0.20,$ $q=0.00, \sigma=0.20, T=1$ <i>On bisecting line</i></p> | <p>GBM $r_f=0.08, \mu=0.20,$ $q=0.00, \sigma=0.40, T=1$ <i>On bisecting line</i></p> |
| <p>GBM $r_f=0.08, \mu=0.05,$ $q=0.00, \sigma=0.20, T=2$ <i>On bisecting line</i></p> | <p>GBM $r_f=0.08, \mu=0.05,$ $q=0.00, \sigma=0.40, T=2$ <i>On bisecting line</i></p> | <p>GBM $r_f=0.08, \mu=0.20,$ $q=0.00, \sigma=0.20, T=2$ <i>On bisecting line</i></p> | <p>GBM $r_f=0.08, \mu=0.20,$ $q=0.00, \sigma=0.40, T=2$ <i>On bisecting line</i></p> |




$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20, T=1$




$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.40, T=1$




$r_f=0.02, \mu=0.20, q=0.00, \sigma=0.20, T=1$




$r_f=0.02, \mu=0.20, q=0.00, \sigma=0.40, T=1$




$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20, T=2$




$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.40, T=2$




$r_f=0.02, \mu=0.20, q=0.00, \sigma=0.20, T=2$




$r_f=0.02, \mu=0.20, q=0.00, \sigma=0.40, T=2$




$r_f=0.02, \mu=0.05, q=0.05, \sigma=0.20, T=1$




$r_f=0.02, \mu=0.05, q=0.15, \sigma=0.20, T=1$




$r_f=0.02, \mu=0.20, q=0.05, \sigma=0.20, T=1$




$r_f=0.02, \mu=0.20, q=0.15, \sigma=0.20, T=1$




$r_f=0.08, \mu=0.05, q=0.00, \sigma=0.20,$
T=1




$r_f=0.08, \mu=0.05, q=0.00, \sigma=0.40,$
T=1




$r_f=0.08, \mu=0.20, q=0.00, \sigma=0.20,$
T=1




$r_f=0.08, \mu=0.20, q=0.00, \sigma=0.40,$
T=1




$r_f=0.08, \mu=0.05, q=0.00, \sigma=0.20,$
T=2



$r_f=0.08, \mu=0.05, q=0.00, \sigma=0.40,$
T=2



$r_f=0.08, \mu=0.20, q=0.00, \sigma=0.20,$
T=2



$r_f=0.08, \mu=0.20, q=0.00, \sigma=0.40,$
T=2

r_f = riskless rate of interest; μ = diffusion statistical drift; σ = diffusion volatility; q = dividend yield; and T = simulation horizon (years)


Experiment parameters: time step number = 90; simulation number = 80,000; resampling pool number = 4; resampling simulation number = 20,000

Table 45: Q-Q Plots for verification of the bootstrap technique applied to Esscher transformed geometric Brownian motions


The Q-Q plots on page 286 and 287 correspond to experiments for the return distributions induced by Merton jump diffusion processes. For these experiments, emphasis is put on varying parameters governing jumps: jump arrival rate is varied between 200% and 800% while jump size volatility is varied between 20% and 40%. The risk-free rate of return is still varied between 2% and 8%, the drift rate is still varied between 5% and 20%, the dividend yield is varied between 0% and 5%, and the time step is varied between four and eight days. In each experiment, the original sample is made of 500,000 weighted returns resulting from the pooling of 50,000 simulated returns from ten time steps and the subsequent weighting via the non-parametric Esscher transformation. This weighted sample is resampled with replacement using the bootstrap procedure to generate 100,000 trajectories. At each time cross-section, these trajectories yield a sample of 100,000 returns but a subsample of only 10,000 returns is used to construct the empirical distribution retained for the test in order to speed-up the verification process. The aforementioned empirical distribution is then compared to the empirical distribution made of all 500,000 weighted returns. The results are summarized in Table 46. It is observed that the Q-Q plots do not “fit” the bisecting line as well as in the previous case and small deviations occur. These deviations are nevertheless restricted to the end of the distribution tails and are due to the difficulty in capturing rare events (jumps) in simulations. Still, most Q-Q plots exhibit loci of quantiles reasonably well aligned on the bisecting lines. Besides, the minor deviations appearing in some plots do not exhibit consistent behavior and are not repeatable.

Table 46: Summary of the bootstrap resampling tests with Q-Q plots for Merton jump diffusion


| | | | |
|---|---|---|---|
| Merton JD $r_f=0.02, \mu=0.05,$ $q=0.00, \sigma=0.20, T=1,$ $\lambda=4.0, \gamma=-0.02, \delta=0.20$ <i>On bisecting line</i> | Merton JD $r_f=0.02, \mu=0.05,$ $q=0.00, \sigma=0.40, T=1,$ $\lambda=4.0, \gamma=-0.08, \delta=0.40$ <i>Minor deviation</i> | Merton JD $r_f=0.02, \mu=0.20,$ $q=0.00, \sigma=0.20, T=1,$ $\lambda=6.0, \gamma=-0.02, \delta=0.20$ <i>On bisecting line</i> | Merton JD $r_f=0.02, \mu=0.20,$ $q=0.00, \sigma=0.40, T=1,$ $\lambda=6.0, \gamma=-0.08, \delta=0.40$ <i>On bisecting line</i> |
| Merton JD $r_f=0.02, \mu=0.05,$ $q=0.00, \sigma=0.20, T=2,$ $\lambda=4.0, \gamma=-0.02, \delta=0.20$ <i>Negative skewness</i> | Merton JD $r_f=0.02, \mu=0.05,$ $q=0.00, \sigma=0.40, T=2,$ $\lambda=4.0, \gamma=-0.08, \delta=0.40$ <i>Positive skewness</i> | Merton JD $r_f=0.02, \mu=0.20,$ $q=0.00, \sigma=0.20, T=2,$ $\lambda=6.0, \gamma=-0.02, \delta=0.20$ <i>Minor deviation</i> | Merton JD $r_f=0.02, \mu=0.20,$ $q=0.00, \sigma=0.40, T=2,$ $\lambda=6.0, \gamma=-0.08, \delta=0.40$ <i>Minor deviation</i> |
| Merton JD $r_f=0.02, \mu=0.05,$ $q=0.05, \sigma=0.20, T=1,$ $\lambda=4.0, \gamma=-0.02, \delta=0.20$ <i>Minor deviation</i> | Merton JD $r_f=0.02, \mu=0.05,$ $q=0.15, \sigma=0.20, T=1,$ $\lambda=4.0, \gamma=-0.08, \delta=0.40$ <i>Minor deviation</i> | Merton JD $r_f=0.02, \mu=0.20,$ $q=0.05, \sigma=0.20, T=1,$ $\lambda=6.0, \gamma=-0.02, \delta=0.20$ <i>On bisecting line</i> | Merton JD $r_f=0.02, \mu=0.20,$ $q=0.15, \sigma=0.20, T=1,$ $\lambda=6.0, \gamma=-0.08, \delta=0.40$ <i>On bisecting line</i> |
| Merton JD $r_f=0.08, \mu=0.05,$ $q=0.00, \sigma=0.20, T=1,$ $\lambda=4.0, \gamma=-0.02, \delta=0.20$ <i>Minor deviation</i> | Merton JD $r_f=0.08, \mu=0.05,$ $q=0.00, \sigma=0.40, T=1,$ $\lambda=4.0, \gamma=-0.08, \delta=0.40$ <i>On bisecting line</i> | Merton JD $r_f=0.08, \mu=0.20,$ $q=0.00, \sigma=0.20, T=1,$ $\lambda=6.0, \gamma=-0.02, \delta=0.20$ <i>On bisecting line</i> | Merton JD $r_f=0.08, \mu=0.20,$ $q=0.00, \sigma=0.40, T=1,$ $\lambda=6.0, \gamma=-0.08, \delta=0.40$ <i>Minor deviation</i> |
| Merton JD $r_f=0.08, \mu=0.05,$ $q=0.00, \sigma=0.20, T=2,$ $\lambda=4.0, \gamma=-0.02, \delta=0.20$ <i>On bisecting line</i> | Merton JD $r_f=0.08, \mu=0.05,$ $q=0.00, \sigma=0.40, T=2,$ $\lambda=4.0, \gamma=-0.08, \delta=0.40$ <i>On bisecting line</i> | Merton JD $r_f=0.08, \mu=0.20,$ $q=0.00, \sigma=0.20, T=2,$ $\lambda=8.0, \gamma=-0.02, \delta=0.20$ <i>On bisecting line</i> | Merton JD $r_f=0.08, \mu=0.20,$ $q=0.00, \sigma=0.40, T=2,$ $\lambda=8.0, \gamma=-0.08, \delta=0.40$ <i>Positive skewness</i> |




$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20,$
 $\lambda=4.00, \gamma=-0.02, \delta=0.20, T=1$




$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20,$
 $\lambda=4.00, \gamma=-0.08, \delta=0.40, T=1$




$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20,$
 $\lambda=6.00, \gamma=-0.02, \delta=0.20, T=1$




$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20,$
 $\lambda=6.00, \gamma=-0.08, \delta=0.40, T=1$




$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20,$
 $\lambda=4.00, \gamma=-0.02, \delta=0.20, T=2$







$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20,$
 $\lambda=4.00, \gamma=-0.08, \delta=0.40, T=2$

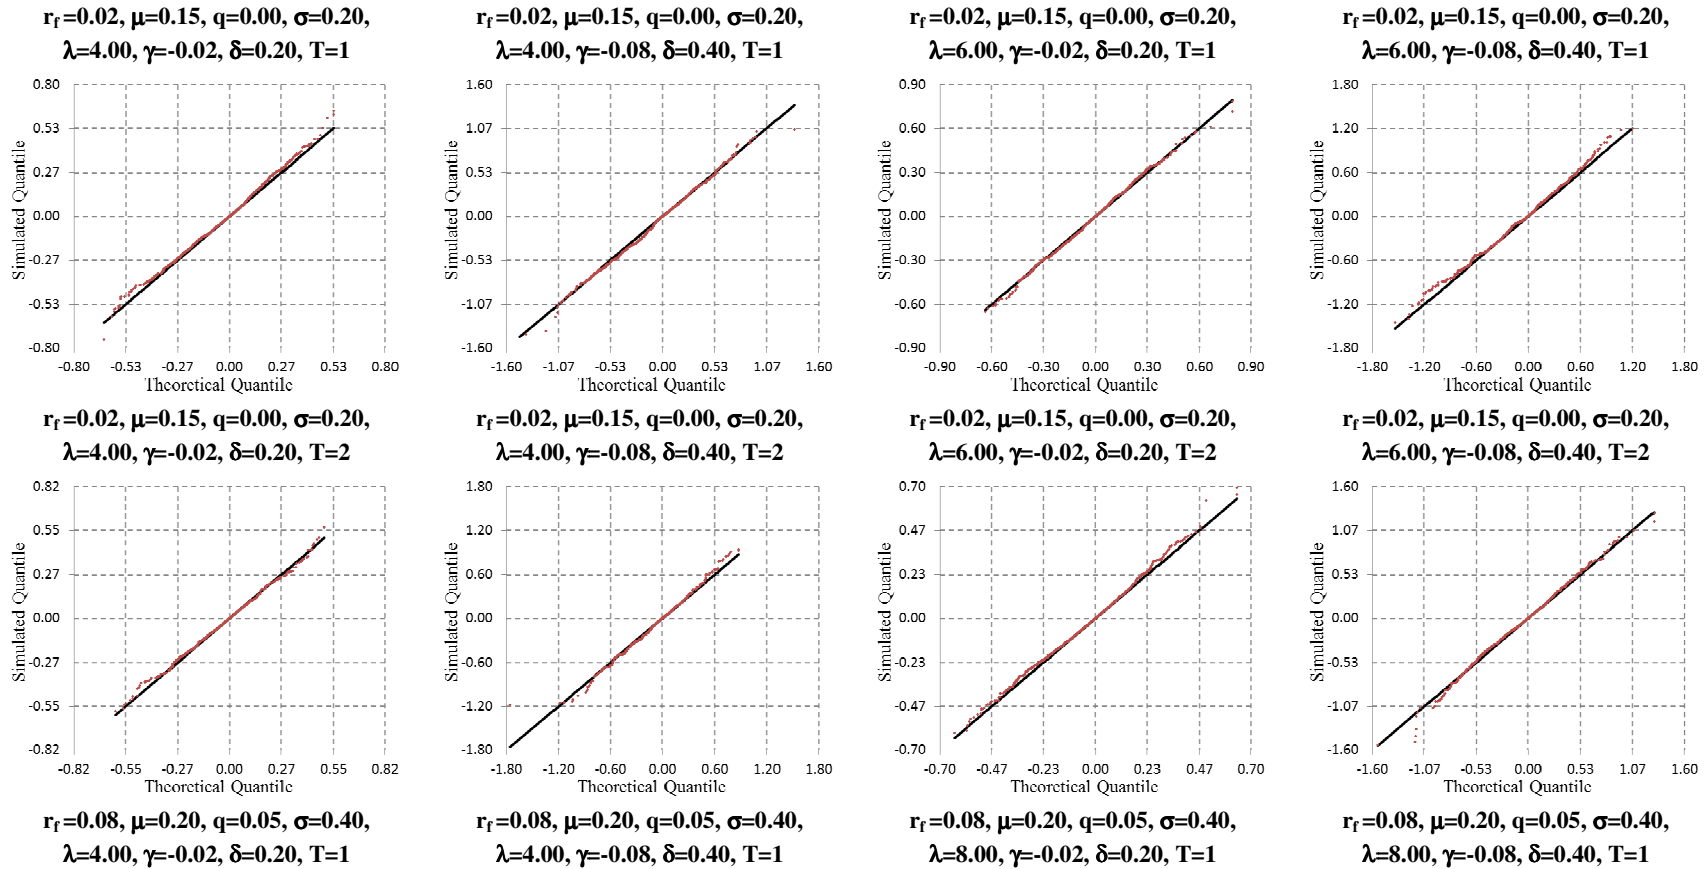


$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20,$
 $\lambda=6.00, \gamma=-0.02, \delta=0.20, T=2$



$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20,$
 $\lambda=6.00, \gamma=-0.08, \delta=0.40, T=2$





r_f = riskless rate of interest; μ = diffusion statistical drift; σ = diffusion volatility; q = dividend yield; λ = arrival rate of jumps (per year);
 γ = mean amplitude of jumps; δ = volatility of jump amplitude; and T = simulation horizon (years)

Experiment parameters: time step number = 90; simulation number = 50,000; resampling pool number= 10; resampling simulation number = 100,000;
 down-sampling factor = 10

Table 47: Q-Q Plots for verification of the bootstrap technique applied to Esscher transformed Jump Diffusion processes

Statistical tests – Kolmogorov-Smirnov test

If Q-Q plots are helpful to qualitatively compare experimental samples and reference distributions, they do not quantify whether the observed deviations are statistically significant. Thus, more testing is performed and for each test case, the non-weighted bootstrapped sample is compared to the original weighted empirical distribution using the two-sample Kolmogorov-Smirnov test. The null hypothesis for the Kolmogorov-Smirnov test is that the experimental samples induced by simulations and non-parametric Esscher transformations are drawn from the original weighted empirical distributions. The Kolmogorov-Smirnov test yields test statistics (and p -values) that can be compared to critical values in order to assess the likelihood of observing such difference between the samples and the original empirical distributions.

Table 48 corresponds to twenty experiments for the return distributions induced by geometric Brownian motions. For these experiments, the risk-free rate of return is varied between 2% and 8%, the drift rate is varied between 5% and 20%, the dividend yield is varied between 0% and 5%, the volatility is varied between 20% and 40%, and the time step is varied between four and eight days. For each experiment, the original sample is made of 320,000 weighted returns resulting from the pooling of 80,000 simulated returns from four time steps and subsequent weighting via the non-parametric Esscher transformation. This weighted distribution is bootstrapped to generate 20,000 trajectories. At each time cross-section, these 20,000 trajectories yield a sample of 20,000 returns. One such sample of 20,000 returns is used to perform the Kolmogorov-Smirnov test. The p -values are computed and all of them are substantially above the 5%

significance level retained for the test. Consequently, the null hypothesis cannot be rejected at the 5% level of significance.

Table 48: Kolmogorov-Smirnov statistical tests for twenty cases of geometric Brownian motions

| r_f | μ | q | σ | T | Kolmogorov Smirnov statistic | p -value |
|---|-------|-----|----------|-----|------------------------------|------------|
| 2.0% | 5% | 0% | 20% | 1.0 | 0.820 | 51% |
| 2.0% | 20% | 0% | 20% | 1.0 | 0.735 | 65% |
| 2.0% | 5% | 0% | 40% | 1.0 | 0.580 | 89% |
| 2.0% | 20% | 0% | 40% | 1.0 | 0.755 | 62% |
| 8.0% | 5% | 0% | 20% | 1.0 | 0.955 | 32% |
| 8.0% | 20% | 0% | 20% | 1.0 | 0.590 | 88% |
| 8.0% | 5% | 0% | 40% | 1.0 | 0.790 | 56% |
| 8.0% | 20% | 0% | 40% | 1.0 | 0.760 | 61% |
| 2.0% | 5% | 0% | 20% | 2.0 | 0.720 | 68% |
| 2.0% | 20% | 0% | 20% | 2.0 | 1.080 | 19% |
| 2.0% | 5% | 0% | 40% | 2.0 | 1.065 | 21% |
| 2.0% | 20% | 0% | 40% | 2.0 | 0.810 | 53% |
| 8.0% | 5% | 0% | 20% | 2.0 | 0.980 | 29% |
| 8.0% | 20% | 0% | 20% | 2.0 | 1.005 | 26% |
| 8.0% | 5% | 0% | 40% | 2.0 | 0.725 | 67% |
| 8.0% | 20% | 0% | 40% | 2.0 | 0.810 | 53% |
| 2.0% | 5% | 5% | 20% | 1.0 | 1.100 | 18% |
| 2.0% | 20% | 5% | 20% | 1.0 | 0.850 | 47% |
| 2.0% | 5% | 15% | 20% | 1.0 | 1.025 | 24% |
| 2.0% | 20% | 15% | 20% | 1.0 | 0.565 | 91% |
| r_f = riskless rate of interest; μ = diffusion statistical drift; σ = diffusion volatility; q = dividend yield; T = simulation horizon (years) Experiment parameters: time step number = 90; simulation number = 80,000; resampling pool number= 4; resampling simulation number = 20,000 | | | | | | |

Monte Carlo simulations can introduce some variability since new pseudo-random number sequences are used each time a test is carried out. If the tests were repeated, this could possibly lead to contradicting conclusions to the two-sample Kolmogorov-Smirnov test. To check the robustness of these conclusions, each of the twenty test cases is repeated thirty times to help gauge how robust the outcomes of the

hypothesis testing are. This yields twenty experiments, each made of thirty repeated tests, for a grand total of six hundred trials. The p -values are computed for each trial and the aggregate experimental distribution of p -values is reported in Figure 62. The distribution of p -values looks a bit more evenly distributed than in the previous verification but it is still right-skewed as shown. Of interest, only three tests in the 600 trials have p -values below the 5% level of significance. This represents less than 0.5% of all tests and consequently, the null hypothesis cannot be rejected at the 5% level of significance.




Figure 62: Distribution of p -values for 600 Kolmogorov-Smirnov tests (geometric Brownian motions)

The second batch of tests deals with the return distributions induced by Merton jump diffusion processes. Table 49 corresponds to twenty experiments with emphasis put on varying parameters governing jumps: jump arrival rate is varied between 200% and 800% while jump size volatility is varied between 20% and 40%. The risk-free rate of return is varied between 2% and 8%, the drift rate is varied between 5% and 20%, the dividend yield is varied between 0% and 5%, and the time step is varied between four and eight days. The simulations are run for 50,000 different paths leading to samples of 50,000 returns at each time step. Ten different samples of returns are pooled together to form a pool of 500,000 returns which are weighted using the non-parametric Esscher

transform. This weighted sample is bootstrapped to generate 100,000 trajectories. At each time cross-section, these 100,000 trajectories yield a sample of 100,000 returns. One such sample is selected and reduced in size (down-sampling achieved by ordering the returns and selecting one return every ten returns) to speed-up the verification process. This new subsample of 10,000 non-weighted returns is compared to the original empirical distribution of 500,000 weighted returns using the two-sample Kolmogorov-Smirnov test.

Table 49: Kolmogorov-Smirnov statistical tests for twenty cases of Merton Jump Diffusion processes

| r_f | μ | q | σ | λ | δ | T | Kolmogorov Smirnov statistic | p-value |
|-------|-------|------|----------|-----------|----------|-----|------------------------------|---------|
| 2.0% | 5% | 0.0% | 20% | 400% | 20% | 1.0 | 0.792 | 56% |
| 2.0% | 5% | 0.0% | 20% | 600% | 20% | 1.0 | 0.516 | 95% |
| 2.0% | 5% | 0.0% | 20% | 400% | 40% | 1.0 | 0.608 | 85% |
| 2.0% | 5% | 0.0% | 20% | 600% | 40% | 1.0 | 0.594 | 87% |
| 2.0% | 15% | 0.0% | 20% | 400% | 20% | 1.0 | 0.820 | 51% |
| 2.0% | 15% | 0.0% | 20% | 600% | 20% | 1.0 | 0.870 | 44% |
| 2.0% | 15% | 0.0% | 20% | 400% | 40% | 1.0 | 0.700 | 71% |
| 2.0% | 15% | 0.0% | 20% | 600% | 40% | 1.0 | 0.940 | 34% |
| 2.0% | 5% | 0.0% | 20% | 400% | 20% | 2.0 | 0.813 | 52% |
| 2.0% | 5% | 0.0% | 20% | 600% | 20% | 2.0 | 0.940 | 34% |
| 2.0% | 5% | 0.0% | 20% | 400% | 40% | 2.0 | 0.467 | 98% |
| 2.0% | 5% | 0.0% | 20% | 600% | 40% | 2.0 | 0.714 | 69% |
| 2.0% | 15% | 0.0% | 20% | 400% | 20% | 2.0 | 0.566 | 91% |
| 2.0% | 15% | 0.0% | 20% | 600% | 20% | 2.0 | 0.940 | 34% |
| 2.0% | 15% | 0.0% | 20% | 400% | 40% | 2.0 | 0.898 | 40% |
| 2.0% | 15% | 0.0% | 20% | 600% | 40% | 2.0 | 0.636 | 81% |
| 8.0% | 20% | 5.0% | 40% | 400% | 20% | 1.0 | 0.601 | 86% |
| 8.0% | 20% | 5.0% | 40% | 800% | 20% | 1.0 | 0.509 | 96% |
| 8.0% | 20% | 5.0% | 40% | 400% | 40% | 1.0 | 0.955 | 32% |
| 8.0% | 20% | 5.0% | 40% | 800% | 40% | 1.0 | 0.714 | 69% |

r_f = riskless rate of interest; μ = diffusion statistical drift; σ = diffusion volatility; q = dividend yield;
 λ = arrival rate of jumps (per year); $\gamma = -\sigma^2/2$ = jump amplitude; δ = volatility of jump amplitude;
T = simulation horizon (years)
Experiment parameters: time step number = 90; simulation number = 50,000; resampling pool number = 10;
resampling simulation number = 100,000; down-sampling ratio = 10

The p -values are computed and all of them are substantially above the 5% significance level retained for the test. Consequently, the null hypothesis cannot be rejected at the 5% level of significance.

As previously done, the robustness of these conclusions is checked by repeating tests: each of the twenty test cases is now repeated thirty times. This yields twenty experiments, each made of thirty repeated tests, for a grand total of six hundred trials. The p -values are computed for each trial and the aggregate experimental distribution of p -values is reported in Figure 63. The distribution of p -values looks more evenly distributed than in the previous verification despite still being slightly right-skewed as shown. Of interest, only five tests out of the 600 trials have p -values below the 5% level of significance. This represents less than 0.9% of all tests and consequently, the null hypothesis cannot be rejected at the 5% level of significance.




Figure 63: Distribution of p -values for 600 Kolmogorov-Smirnov tests (Jump Diffusion processes)

7.3.3 Combined non-parametric Esscher transform and resampling

The simulation, non-parametric Esscher transform, and resampling have been introduced as a way to change the probability measure, and then simulate the evolution of complex stochastic processes under the equivalent martingale measure. Thus, the next step in the verification process consists in combining the simulation, non-parametric Esscher transformation, and resampling using the bootstrap procedure and checking whether the combined approach performs adequately. Only then can conclusions be drawn about hypothesis 1.1.2 and 1.1.4.

Verification process and criteria for success

Similarly to what was done previously, several tests are performed to verify the implementation of the combined approach. The first is a qualitative test that visually compares the empirical probability distribution induced by the simulation, non-parametric Esscher transform, and resampling, to the known probability distribution under the equivalent martingale measure. This test is considered successful if there is no apparent departure from the bisecting line in the Q-Q plot. The second test is quantitative and uses the one-sample Kolmogorov-Smirnov statistical test to confirm whether the empirical probability distribution and the corresponding theoretical probability distribution differ. This test is considered successful if the equality of distributions cannot be rejected at a five percent level of significance. The third test is also quantitative and uses the z -test and the Student's t -test to confirm whether the means of these two distributions differ in any way. This test is considered successful if the equality of means cannot be rejected at a five percent level of significance. All these tests are performed for

two popular stochastic processes: the geometric Brownian motion and the Merton jump-diffusion process.

Graphical tests





The experimental distributions obtained via combined simulation, non-parametric Esscher transformation, and bootstrapping are compared to the corresponding known theoretical distribution using Q-Q plots. The first batch of tests deals with the geometric Brownian motion. The stochastic process is sampled to generate 80,000 trajectories representing the possible evolution of a development program over time. These trajectories induce samples of 80,000 returns at each time step. Four of these samples are pooled together to yield a larger sample of 320,000 returns and a change of probability measure is performed using the non-parametric Esscher transform. This sample of 320,000 weighted returns is bootstrapped using the sampling wheel algorithm to yield 20,000 new trajectories. These trajectories induce a sample of 20,000 non-weighted returns at each time step. One such sample is selected and compared to the theoretical return distribution under the equivalent martingale measure. Table 51 on page 296 and 297 displays several plots corresponding to twenty test cases with the risk free rate of return varied between 2% and 8%, the drift rate varied between 5% and 20%, the dividend yield varied between 0% and 15%, the volatility varied between 20% and 40% and finally the time step varied between four and eight days.

A summary of the results is provided in Table 50. All Q-Q plots exhibit loci of quantiles aligned almost perfectly on the bisecting lines. While minor deviations appear on some Q-Q plots, these deviations are restricted to extreme quantiles at the far-end of the distribution tails (usually first and last three quantiles in several thousands). Besides,

when the tests are repeated and new Q-Q plots are graphed, these minor deviations often disappear and reappear randomly. These deviations are thus not repeatable and, as such, are just artifacts of the simulations.

Table 50: Summary of the combined approach tests using Q-Q plots with GBM

| | | | |
|--|--|--|--|
| GBM $r_f=0.02, \mu=0.05,$ $q=0.00, \sigma=0.20, T=1$ <i>On bisecting line</i> | GBM $r_f=0.02, \mu=0.05,$ $q=0.00, \sigma=0.40, T=1$ <i>On bisecting line</i> | GBM $r_f=0.02, \mu=0.20,$ $q=0.00, \sigma=0.20, T=1$ <i>Minor deviation</i> | GBM $r_f=0.02, \mu=0.20,$ $q=0.00, \sigma=0.40, T=1$ <i>On bisecting line</i> |
| GBM $r_f=0.02, \mu=0.05,$ $q=0.00, \sigma=0.20, T=2$ <i>On bisecting line</i> | GBM $r_f=0.02, \mu=0.05,$ $q=0.00, \sigma=0.40, T=2$ <i>Minor deviation</i> | GBM $r_f=0.02, \mu=0.20,$ $q=0.00, \sigma=0.20, T=2$ <i>On bisecting line</i> | GBM $r_f=0.02, \mu=0.20,$ $q=0.00, \sigma=0.40, T=2$ <i>Minor deviation</i> |
| GBM $r_f=0.02, \mu=0.05,$ $q=0.05, \sigma=0.20, T=1$ <i>On bisecting line</i> | GBM $r_f=0.02, \mu=0.05,$ $q=0.15, \sigma=0.20, T=1$ <i>Minor deviation</i> | GBM $r_f=0.02, \mu=0.20,$ $q=0.05, \sigma=0.20, T=1$ <i>On bisecting line</i> | GBM $r_f=0.02, \mu=0.20,$ $q=0.15, \sigma=0.20, T=1$ <i>Minor deviation</i> |
| GBM $r_f=0.08, \mu=0.05,$ $q=0.00, \sigma=0.20, T=1$ <i>On bisecting line</i> | GBM $r_f=0.08, \mu=0.05,$ $q=0.00, \sigma=0.40, T=1$ <i>Minor deviation</i> | GBM $r_f=0.08, \mu=0.20,$ $q=0.00, \sigma=0.20, T=1$ <i>On bisecting line</i> | GBM $r_f=0.08, \mu=0.20,$ $q=0.00, \sigma=0.40, T=1$ <i>On bisecting line</i> |
| GBM $r_f=0.08, \mu=0.05,$ $q=0.00, \sigma=0.20, T=2$ <i>On bisecting line</i> | GBM $r_f=0.08, \mu=0.05,$ $q=0.00, \sigma=0.40, T=2$ <i>On bisecting line</i> | GBM $r_f=0.08, \mu=0.20,$ $q=0.00, \sigma=0.20, T=2$ <i>On bisecting line</i> | GBM $r_f=0.08, \mu=0.20,$ $q=0.00, \sigma=0.40, T=2$ <i>On bisecting line</i> |







$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20, T=1$

$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.40, T=1$

$r_f=0.02, \mu=0.20, q=0.00, \sigma=0.20, T=1$

$r_f=0.02, \mu=0.20, q=0.00, \sigma=0.40, T=1$







$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20, T=2$

$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.40, T=2$

$r_f=0.02, \mu=0.20, q=0.00, \sigma=0.20, T=2$

$r_f=0.02, \mu=0.20, q=0.00, \sigma=0.40, T=2$




$r_f=0.02, \mu=0.05, q=0.05, \sigma=0.20, T=1$


$r_f=0.02, \mu=0.05, q=0.15, \sigma=0.20, T=1$

$r_f=0.02, \mu=0.20, q=0.05, \sigma=0.20, T=1$


$r_f=0.02, \mu=0.20, q=0.15, \sigma=0.20, T=1$




$r_f=0.08, \mu=0.05, q=0.00, \sigma=0.20,$
T=1




$r_f=0.08, \mu=0.05, q=0.00, \sigma=0.40,$
T=1




$r_f=0.08, \mu=0.20, q=0.00, \sigma=0.20,$
T=1




$r_f=0.08, \mu=0.20, q=0.00, \sigma=0.40,$
T=1




$r_f=0.08, \mu=0.05, q=0.00, \sigma=0.20,$
T=2



$r_f=0.08, \mu=0.05, q=0.00, \sigma=0.40,$
T=2



$r_f=0.08, \mu=0.20, q=0.00, \sigma=0.20,$
T=2



$r_f=0.08, \mu=0.20, q=0.00, \sigma=0.40,$
T=2

r_f = riskless rate of interest; μ = diffusion statistical drift; σ = diffusion volatility; q = dividend yield; and T = simulation horizon (years)

Experiment parameters: time step number = 90; simulation number = 80,000; resampling pool number = 4; resampling simulation number = 20,000

Table 51: Q-Q Plots for verification of the combined approach for geometric Brownian motions


The Q-Q plots in Table 53 pages 300 and 301 correspond to experiments with return distributions induced by Merton jump diffusion processes. For these experiments, emphasis is put on varying parameters governing jumps: the jump arrival rate is varied between 200% and 800% while the jump size volatility is varied between 20% and 40%. The risk-free rate of return is still varied between 2% and 8%, the drift rate is still varied between 5% and 20%, the dividend yield is varied between 0% and 5%, and the time step is varied between four and eight days. In each experiment, Monte Carlo simulations are used to generate 50,000 trajectories. A sample of 500,000 weighted returns is constructed from the pooling of 50,000 simulated returns from ten different time steps and weighting is performed with the non-parametric Esscher transform. This sample of weighted returns is bootstrapped using the sampling wheel algorithm to generate 100,000 trajectories. At each time cross-section, these 100,000 trajectories yield a sample of 100,000 returns. To speed-up the verification, only one return every ten returns is selected for the construction of Q-Q plots.

A summary of the results is provided in Table 52. It is observed that the Q-Q plots do not perfectly “fit” the bisecting line and small deviations occur. For the most part, these deviations are restricted to the extreme end of the distribution tails and are due to the difficulty in capturing rare events in the simulations (jumps). In addition, the minor deviations appearing in some of the plots do not have consistent patterns and cannot be reproduced during repeated experiments. Moreover, in two experiments, there is a *notable* departure from the bisecting line. These two experiments are investigated further with the Kolmogorov-Smirnov test to check the statistical significance of the departure from the bisecting line. It is worth noting that Q-Q plots seem to improve as the arrival

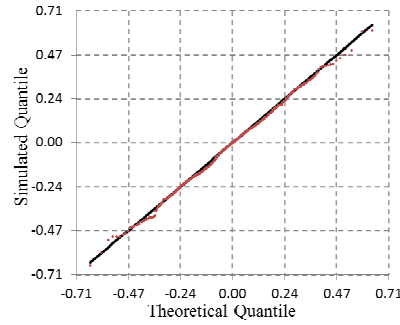
rate of jumps increases: this is due to more frequent jumps that are therefore easier to capture in the simulations. Despite all this, most Q-Q plots exhibit loci of quantiles reasonably well aligned on the bisecting lines.

Table 52: Summary of the combined approach tests with Q-Q plots for Merton jump diffusion


| | | | |
|--|--|--|--|
| Merton JD $r_f=0.02, \mu=0.05,$ $q=0.00, \sigma=0.20, T=1$ $\lambda=4.0, \gamma=-0.02, \delta=0.20$ <i>Minor deviation</i> | Merton JD $r_f=0.02, \mu=0.05,$ $q=0.00, \sigma=0.40, T=1$ $\lambda=4.0, \gamma=-0.08, \delta=0.40$ <i>On bisecting line</i> | Merton JD $r_f=0.02, \mu=0.20,$ $q=0.00, \sigma=0.20, T=1$ $\lambda=6.0, \gamma=-0.02, \delta=0.20$ <i>On bisecting line</i> | Merton JD $r_f=0.02, \mu=0.20,$ $q=0.00, \sigma=0.40, T=1$ $\lambda=6.0, \gamma=-0.08, \delta=0.40$ <i>Notable deviation</i> |
| Merton JD $r_f=0.02, \mu=0.05,$ $q=0.00, \sigma=0.20, T=2$ $\lambda=4.0, \gamma=-0.02, \delta=0.20$ <i>Minor deviation</i> | Merton JD $r_f=0.02, \mu=0.05,$ $q=0.00, \sigma=0.40, T=2$ $\lambda=4.0, \gamma=-0.08, \delta=0.40$ <i>Notable deviation</i> | Merton JD $r_f=0.02, \mu=0.20,$ $q=0.00, \sigma=0.20, T=2$ $\lambda=6.0, \gamma=-0.02, \delta=0.20$ <i>On bisecting line</i> | Merton JD $r_f=0.02, \mu=0.20,$ $q=0.00, \sigma=0.40, T=2$ $\lambda=6.0, \gamma=-0.08, \delta=0.40$ <i>On bisecting line</i> |
| Merton JD $r_f=0.02, \mu=0.05,$ $q=0.05, \sigma=0.20, T=1$ $\lambda=4.0, \gamma=-0.02, \delta=0.20$ <i>On bisecting line</i> | Merton JD $r_f=0.02, \mu=0.05,$ $q=0.15, \sigma=0.20, T=1$ $\lambda=4.0, \gamma=-0.08, \delta=0.40$ <i>Minor deviation</i> | Merton JD $r_f=0.02, \mu=0.20,$ $q=0.05, \sigma=0.20, T=1$ $\lambda=6.0, \gamma=-0.02, \delta=0.20$ <i>On bisecting line</i> | Merton JD $r_f=0.02, \mu=0.20,$ $q=0.15, \sigma=0.20, T=1$ $\lambda=6.0, \gamma=-0.08, \delta=0.40$ <i>Minor deviation</i> |
| Merton JD $r_f=0.08, \mu=0.05,$ $q=0.00, \sigma=0.20, T=1$ $\lambda=4.0, \gamma=-0.02, \delta=0.20$ <i>On bisecting line</i> | Merton JD $r_f=0.08, \mu=0.05,$ $q=0.00, \sigma=0.40, T=1$ $\lambda=4.0, \gamma=-0.08, \delta=0.40$ <i>On bisecting line</i> | Merton JD $r_f=0.08, \mu=0.20,$ $q=0.00, \sigma=0.20, T=1$ $\lambda=6.0, \gamma=-0.02, \delta=0.20$ <i>On bisecting line</i> | Merton JD $r_f=0.08, \mu=0.20,$ $q=0.00, \sigma=0.40, T=1$ $\lambda=6.0, \gamma=-0.08, \delta=0.40$ <i>On bisecting line</i> |
| Merton JD $r_f=0.08, \mu=0.05,$ $q=0.00, \sigma=0.20, T=2$ $\lambda=4.0, \gamma=-0.02, \delta=0.20$ <i>On bisecting line</i> | Merton JD $r_f=0.08, \mu=0.05,$ $q=0.00, \sigma=0.40, T=2$ $\lambda=4.0, \gamma=-0.08, \delta=0.40$ <i>On bisecting line</i> | Merton JD $r_f=0.08, \mu=0.20,$ $q=0.00, \sigma=0.20, T=2$ $\lambda=8.0, \gamma=-0.02, \delta=0.20$ <i>Minor deviation</i> | Merton JD $r_f=0.08, \mu=0.20,$ $q=0.00, \sigma=0.40, T=2$ $\lambda=8.0, \gamma=-0.08, \delta=0.40$ <i>On bisecting line</i> |




$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20,$
 $\lambda=4.00, \gamma=-0.02, \delta=0.20, T=1$




$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20,$
 $\lambda=4.00, \gamma=-0.08, \delta=0.40, T=1$




$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20,$
 $\lambda=6.00, \gamma=-0.02, \delta=0.20, T=1$




$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20,$
 $\lambda=6.00, \gamma=-0.08, \delta=0.40, T=1$




$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20,$
 $\lambda=4.00, \gamma=-0.02, \delta=0.20, T=2$







$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20,$
 $\lambda=4.00, \gamma=-0.08, \delta=0.40, T=2$







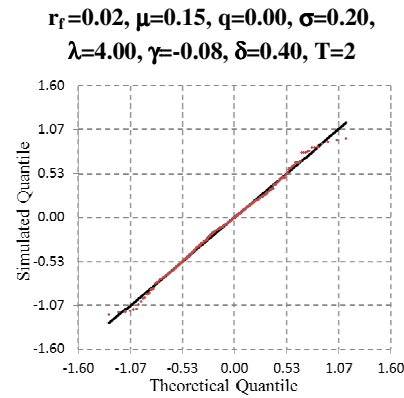







$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20,$
 $\lambda=6.00, \gamma=-0.02, \delta=0.20, T=2$



$r_f=0.02, \mu=0.05, q=0.00, \sigma=0.20,$
 $\lambda=6.00, \gamma=-0.08, \delta=0.40, T=2$





r_f = riskless rate of interest; μ = diffusion statistical drift; σ = diffusion volatility; q = dividend yield; λ = arrival rate of jumps (per year);
 γ = mean amplitude of jumps; δ = volatility of jump amplitude; and T = simulation horizon (years)

Experiment parameters: time step number = 90; simulation number = 50,000; resampling pool number = 10; resampling draws = 100,000;
down-sampling factor = 10

Table 53: Q-Q Plots for verification of the combined approach for Merton jump diffusion processes

Statistical tests – Kolmogorov-Smirnov test

The experimental distributions obtained via combined simulation, non-parametric Esscher transformation, and bootstrapping are compared to the corresponding theoretical distribution using Kolmogorov-Smirnov tests.

Table 54 corresponds to twenty experiments with the return distributions induced by geometric Brownian motions. For these experiments, the risk-free rate of return is varied between 2% and 8%, the drift rate is varied between 5% and 20%, the dividend yield is varied between 0% and 5%, the volatility is varied between 20% and 40%, and the time step is varied between four and eight days. In each experiment, Monte Carlo simulations are used to generate 80,000 trajectories. A sample of 320,000 weighted returns is constructed from the pooling of 80,000 simulated returns from four different time steps. It is subsequently weighted with the non-parametric Esscher transformation and then bootstrapped to generate 20,000 trajectories. At each time cross-section, these 20,000 trajectories yield a sample of 20,000 returns. One such sample is used to carry out the one-sample Kolmogorov-Smirnov test. The null hypothesis for the test is that the experimental sample is drawn from the known theoretical distribution. The p -values are computed and all of them are above the 5% significance level retained for the test. Thus, the null hypothesis cannot be rejected at the 5% level of significance. Interestingly, p -values are more evenly distributed than in previous tests and some of them are actually close to the critical region.

Table 54: Kolmogorov-Smirnov statistical tests for twenty cases of geometric Brownian motions

| r_f | μ | q | σ | T | Kolmogorov Smirnov statistic | p -value |
|--|-------|-----|----------|-----|------------------------------|------------|
| 2.0% | 5% | 0% | 20% | 1.0 | 0.813 | 52% |
| 2.0% | 20% | 0% | 20% | 1.0 | 0.877 | 43% |
| 2.0% | 5% | 0% | 40% | 1.0 | 0.594 | 87% |
| 2.0% | 20% | 0% | 40% | 1.0 | 0.827 | 50% |
| 8.0% | 5% | 0% | 20% | 1.0 | 0.700 | 71% |
| 8.0% | 20% | 0% | 20% | 1.0 | 0.580 | 89% |
| 8.0% | 5% | 0% | 40% | 1.0 | 0.700 | 71% |
| 8.0% | 20% | 0% | 40% | 1.0 | 0.453 | 99% |
| 2.0% | 5% | 0% | 20% | 2.0 | 0.870 | 44% |
| 2.0% | 20% | 0% | 20% | 2.0 | 0.919 | 37% |
| 2.0% | 5% | 0% | 40% | 2.0 | 1.223 | 10% |
| 2.0% | 20% | 0% | 40% | 2.0 | 1.216 | 10% |
| 8.0% | 5% | 0% | 20% | 2.0 | 0.735 | 65% |
| 8.0% | 20% | 0% | 20% | 2.0 | 1.181 | 12% |
| 8.0% | 5% | 0% | 40% | 2.0 | 0.679 | 75% |
| 8.0% | 20% | 0% | 40% | 2.0 | 1.089 | 19% |
| 2.0% | 5% | 5% | 20% | 1.0 | 1.110 | 17% |
| 2.0% | 20% | 5% | 20% | 1.0 | 0.976 | 30% |
| 2.0% | 5% | 15% | 20% | 1.0 | 1.131 | 15% |
| 2.0% | 20% | 15% | 20% | 1.0 | 0.566 | 91% |
| r_f = riskless rate of interest; μ = diffusion statistical drift; σ = diffusion volatility; q = dividend yield; T = simulation horizon (years) Experiment parameters: time step number = 90; simulation number = 80,000; resampling pool number = 4; resampling simulation number = 20,000 | | | | | | |

Monte Carlo simulations introduce some variability as new pseudo-random number sequences are used and new distributions are generated each time a test is carried out. To check the robustness of these conclusions, each of the twenty test cases is repeated thirty times to help gauge how robust the outcomes of the hypothesis testing are. This yields twenty experiments, each made of thirty repeated tests, for a grand total of six hundred trials. The p -values are computed for each trial and the aggregate experimental distribution of p -values is reported in Figure 64. The distribution of p -values looks more

evenly distributed than in previous verifications. Of interest, only nineteen trials out of the 600 trials have p -values below the 5% level of significance. This represents less than 3.2% of all tests and consequently, the null hypothesis cannot be rejected at the 5% level of significance.




Figure 64: Distribution of p -values for 600 Kolmogorov-Smirnov tests (geometric Brownian motions)

The second batch of tests deals with the return distributions induced by Merton jump diffusion processes. Table 55 corresponds to twenty experiments with emphasis put on varying parameters governing jumps: the jump arrival rate is varied between 200% and 800% and the jump size volatility is varied between 20% and 40%. The risk-free rate of return is still varied between 2% and 8%, the drift rate is still varied between 5% and 20%, the dividend yield is varied between 0% and 5%, and the time step is varied between four and eight days. In each test case, Monte Carlo simulations are used to generate 50,000 trajectories. A sample of 500,000 weighted returns is constructed from the pooling of 50,000 simulated returns from ten different time steps. A change of probability measure is performed with the non-parametric Esscher transform. This yields a weighted sample which is bootstrapped to generate 100,000 trajectories. These trajectories yield samples of 100,000 returns at each time step. In order to speed-up the

verification, down-sampling is performed by ordering returns and selecting one every ten returns. This yields a subsample of 10,000 returns which are used to perform the Kolmogorov-Smirnov test. The null hypothesis for the test is that the experimental sample is drawn from the known theoretical distribution. The p -values are computed and all of them are above the 5% significance level retained for the test. Thus, the null hypothesis cannot be rejected at the 5% level of significance.

Table 55: Kolmogorov-Smirnov statistical tests for twenty cases of Merton Jump Diffusion processes

| r_f | μ | q | σ | λ | δ | T | Kolmogorov Smirnov statistic | p -value |
|-------|-------|------|----------|-----------|----------|-----|------------------------------|------------|
| 2.0% | 5% | 0.0% | 20% | 400% | 20% | 1.0 | 1.010 | 26% |
| 2.0% | 5% | 0.0% | 20% | 600% | 20% | 1.0 | 0.790 | 56% |
| 2.0% | 5% | 0.0% | 20% | 400% | 40% | 1.0 | 0.780 | 58% |
| 2.0% | 5% | 0.0% | 20% | 600% | 40% | 1.0 | 0.660 | 78% |
| 2.0% | 15% | 0.0% | 20% | 400% | 20% | 1.0 | 0.630 | 82% |
| 2.0% | 15% | 0.0% | 20% | 600% | 20% | 1.0 | 0.850 | 47% |
| 2.0% | 15% | 0.0% | 20% | 400% | 40% | 1.0 | 0.650 | 79% |
| 2.0% | 15% | 0.0% | 20% | 600% | 40% | 1.0 | 1.060 | 21% |
| 2.0% | 5% | 0.0% | 20% | 400% | 20% | 2.0 | 0.930 | 35% |
| 2.0% | 5% | 0.0% | 20% | 600% | 20% | 2.0 | 0.910 | 38% |
| 2.0% | 5% | 0.0% | 20% | 400% | 40% | 2.0 | 0.750 | 63% |
| 2.0% | 5% | 0.0% | 20% | 600% | 40% | 2.0 | 1.080 | 19% |
| 2.0% | 15% | 0.0% | 20% | 400% | 20% | 2.0 | 0.740 | 64% |
| 2.0% | 15% | 0.0% | 20% | 600% | 20% | 2.0 | 1.100 | 18% |
| 2.0% | 15% | 0.0% | 20% | 400% | 40% | 2.0 | 0.700 | 71% |
| 2.0% | 15% | 0.0% | 20% | 600% | 40% | 2.0 | 0.680 | 74% |
| 8.0% | 20% | 5.0% | 40% | 400% | 20% | 1.0 | 0.620 | 84% |
| 8.0% | 20% | 5.0% | 40% | 800% | 20% | 1.0 | 0.640 | 81% |
| 8.0% | 20% | 5.0% | 40% | 400% | 40% | 1.0 | 1.060 | 21% |
| 8.0% | 20% | 5.0% | 40% | 800% | 40% | 1.0 | 0.990 | 28% |

r_f = riskless rate of interest; μ = diffusion statistical drift; σ = diffusion volatility; q = dividend yield;
 λ = arrival rate of jumps (per year); $\gamma = -\sigma^2/2$ = jump amplitude; δ = volatility of jump amplitude;
T = simulation horizon (years)
Experiment parameters: time step number = 90; simulation number = 50,000;
resampling pool number = 10; resampling simulation number = 100,000; down-sampling ratio = 10

The p -values in Table 55 are evenly distributed and some of them are close to the critical region. As a result, repeated trials are warranted to check the sensitivity of these results. Each of the twenty test cases is now repeated thirty times leading to a grand total of six hundred trials. The p -values are computed for each trial and the aggregate distribution of p -values is reported in Figure 65. The distribution of p -values looks evenly distributed. Of interest, only sixteen trials in the 600 trials have p -values below the 5% level of significance which represents less than 2.7% of all tests. Consequently, the null hypothesis cannot be rejected at the 5% level of significance.




Figure 65: Distribution of p -values for 600 Kolmogorov-Smirnov tests (Jump Diffusion processes)

The combination of simulation, non-parametric Esscher transformation, and bootstrap represents a “serial” experimentation. It is expected that the quality of results degrades as the initial sample from the simulation is being “worked on” and that the final sample quality is impacted by the quality of the transformation applied at each intermediate step. To assess the criticality of each step on the final outcome, the coefficient of correlation between the final p -values and the intermediate p -values is

computed for the 600 trials and reported in Table 56. It appears that the bootstrap procedure is most critical.

Table 56: Correlation between final results and intermediate results

| | Correlation between final p -values and p -values after initial Esscher transform | Correlation between final p -values and p -values after bootstrap resampling |
|---------------------------|---|--|
| Geometric Brownian Motion | -0.040 | 0.620 |
| Merton Jump Diffusion | -0.047 | 0.493 |

Statistical tests – Testing the mean with z-tests and t-tests

Switching to the first moment of these distributions, the expected values of the empirical distributions obtained via simulation, non-parametric Esscher transformation, and bootstrap are compared to the expected values of corresponding theoretical distributions. Statistical testing is performed by repeating each test case thirty times yielding an experimental average mean and an experimental standard error which are used to carry out two-tailed z -tests and two-tailed t -tests. The null hypothesis for the z -tests and t -tests is that the experimental average mean equals the expected value of the known theoretical distribution.

The first batch of experiments deals with geometric Brownian motions using different combinations of parameters (risk free rate, drift, volatility, dividend yield, and maturities) representative of what would be used for the pricing of real options. The risk-free rate of return is varied between 2% and 8%, the drift rate is varied between 5% and 20%, the dividend yield is varied between 0% and 15%, the volatility is varied between 20% and 40%, and finally, the time step is varied between four and eight days (90 time steps with a maturity varied between one and two years). Table 57 highlights the results of twenty experiments: approximate p -values are computed and all of them are above 5%.

Consequently, the null hypothesis cannot be rejected at the 5% level of significance. As expected, the t -test and the z -test yield similar approximations of the p -values since both tests are applicable “in the limit” (due to the large sample size).

Table 57: z -tests and t -tests for the mean returns of twenty cases of geometric Brownian motions

| r_f | μ | q | σ | T | Experimental Sample Mean Return | Experimental Sample Standard Error | Theoretical Mean Return | z -test and t -test statistic | z -test p -value | t -test p -value |
|-------|-------|-----|----------|---|---------------------------------|------------------------------------|-------------------------|-----------------------------------|----------------------|----------------------|
| 2% | 5% | 0% | 20% | 1 | -4.34E-05 | 2.75E-05 | -3.85E-20 | 1.579 | 11% | 13% |
| 2% | 20% | 0% | 20% | 1 | 2.24E-05 | 3.33E-05 | -3.85E-20 | 0.674 | 50% | 51% |
| 2% | 5% | 0% | 40% | 1 | -6.60E-04 | 4.98E-05 | -6.67E-04 | 0.135 | 89% | 89% |
| 2% | 20% | 0% | 40% | 1 | -7.08E-04 | 5.58E-05 | -6.67E-04 | 0.750 | 45% | 46% |
| 8% | 5% | 0% | 20% | 1 | 6.72E-04 | 2.97E-05 | 6.67E-04 | 0.173 | 86% | 86% |
| 8% | 20% | 0% | 20% | 1 | 6.89E-04 | 2.29E-05 | 6.67E-04 | 0.974 | 33% | 34% |
| 8% | 5% | 0% | 40% | 1 | -6.59E-05 | 5.02E-05 | -1.54E-19 | 1.314 | 19% | 20% |
| 8% | 20% | 0% | 40% | 1 | -9.56E-05 | 7.28E-05 | -1.54E-19 | 1.315 | 19% | 20% |
| 2% | 5% | 0% | 20% | 2 | 3.59E-05 | 4.77E-05 | -7.71E-20 | 0.752 | 45% | 46% |
| 2% | 20% | 0% | 20% | 2 | 2.94E-05 | 3.51E-05 | -7.71E-20 | 0.837 | 40% | 41% |
| 2% | 5% | 0% | 40% | 2 | -1.36E-03 | 8.63E-05 | -1.33E-03 | 0.271 | 79% | 79% |
| 2% | 20% | 0% | 40% | 2 | -1.30E-03 | 7.95E-05 | -1.33E-03 | 0.442 | 66% | 66% |
| 8% | 5% | 0% | 20% | 2 | 1.27E-03 | 3.87E-05 | 1.33E-03 | 1.760 | 8% | 9% |
| 8% | 20% | 0% | 20% | 2 | 1.37E-03 | 3.34E-05 | 1.33E-03 | 0.985 | 32% | 33% |
| 8% | 5% | 0% | 40% | 2 | -4.90E-05 | 7.41E-05 | -3.08E-19 | 0.661 | 51% | 51% |
| 8% | 20% | 0% | 40% | 2 | 4.49E-05 | 7.28E-05 | -3.08E-19 | 0.617 | 54% | 54% |
| 2% | 5% | 5% | 20% | 1 | -5.66E-04 | 3.05E-05 | -5.56E-04 | 0.339 | 73% | 74% |
| 2% | 20% | 5% | 20% | 1 | -5.58E-04 | 3.02E-05 | -5.56E-04 | 0.067 | 95% | 95% |
| 2% | 5% | 15% | 20% | 1 | -1.71E-03 | 3.38E-05 | -1.67E-03 | 1.385 | 17% | 18% |
| 2% | 20% | 15% | 20% | 1 | -1.67E-03 | 3.45E-05 | -1.67E-03 | 0.095 | 92% | 93% |

r_f = riskless rate of interest; μ = diffusion statistical drift; σ = diffusion volatility;
 q = dividend yield; T = simulation horizon (years)
Experiment parameters: time step number = 90; simulation number = 80,000;
resampling pool number = 4; resampling draws = 20,000; down-sampling ratio = 1

The second batch of experiments deals with the return distributions induced by Merton jump diffusion processes. Table 58 highlights the results of twenty experiments with emphasis put on varying parameters governing jumps: jump arrival rate is varied between 200% and 800% while jump size volatility is varied between 20% and 40%. The

risk-free rate of return is still varied between 2% and 8%, the drift rate is still varied between 5% and 20%, the dividend yield is varied between 0% and 5%, and the time step is varied between four and eight days. Each experiment consists of thirty repeated trials which enable the computation of the sample mean, the standard error, and in turn the z -statistic and t -statistic. The null hypothesis for the z -test and t -test is that the experimental sample mean is equal to the theoretical expected value.

Table 58: z-tests and t-tests for the mean returns of twenty cases of Merton jump diffusion processes

| r_f | μ | q | σ | λ | δ | T | Exp. Sample Mean Return | Exp. Sample Standard Error | Theo. Mean Return | t -test and z -test statistics | z -test p -value | t -test p -value |
|-------|-------|-----|----------|-----------|----------|---|-------------------------|----------------------------|-------------------|------------------------------------|----------------------|----------------------|
| 2% | 5% | 0% | 20% | 400% | 20% | 1 | -9.03E-04 | 2.31E-05 | -8.93E-04 | 0.434 | 66% | 67% |
| 2% | 5% | 0% | 20% | 600% | 20% | 1 | -1.38E-03 | 3.46E-05 | -1.34E-03 | 1.217 | 22% | 23% |
| 2% | 5% | 0% | 20% | 400% | 40% | 1 | -3.57E-03 | 5.32E-05 | -3.57E-03 | 0.066 | 95% | 95% |
| 2% | 5% | 0% | 20% | 600% | 40% | 1 | -5.38E-03 | 5.97E-05 | -5.35E-03 | 0.509 | 61% | 61% |
| 2% | 15% | 0% | 20% | 400% | 20% | 1 | -9.26E-04 | 2.41E-05 | -9.23E-04 | 0.107 | 92% | 92% |
| 2% | 15% | 0% | 20% | 600% | 20% | 1 | -1.37E-03 | 3.12E-05 | -1.36E-03 | 0.218 | 83% | 83% |
| 2% | 15% | 0% | 20% | 400% | 40% | 1 | -3.63E-03 | 5.71E-05 | -3.64E-03 | 0.197 | 84% | 84% |
| 2% | 15% | 0% | 20% | 600% | 40% | 1 | -5.47E-03 | 3.77E-05 | -5.41E-03 | 1.657 | 10% | 11% |
| 2% | 5% | 0% | 20% | 400% | 20% | 2 | -1.78E-03 | 3.35E-05 | -1.79E-03 | 0.242 | 81% | 81% |
| 2% | 5% | 0% | 20% | 600% | 20% | 2 | -2.75E-03 | 4.66E-05 | -2.67E-03 | 1.641 | 10% | 11% |
| 2% | 5% | 0% | 20% | 400% | 40% | 2 | -7.21E-03 | 7.41E-05 | -7.14E-03 | 0.939 | 35% | 36% |
| 2% | 5% | 0% | 20% | 600% | 40% | 2 | -1.06E-02 | 8.96E-05 | -1.07E-02 | 1.281 | 20% | 21% |
| 2% | 15% | 0% | 20% | 400% | 20% | 2 | -1.88E-03 | 3.97E-05 | -1.85E-03 | 0.852 | 39% | 40% |
| 2% | 15% | 0% | 20% | 600% | 20% | 2 | -2.78E-03 | 3.69E-05 | -2.73E-03 | 1.551 | 12% | 13% |
| 2% | 15% | 0% | 20% | 400% | 40% | 2 | -7.21E-03 | 8.28E-05 | -7.28E-03 | 0.854 | 39% | 40% |
| 2% | 15% | 0% | 20% | 600% | 40% | 2 | -1.07E-02 | 7.84E-05 | -1.08E-02 | 1.347 | 18% | 19% |
| 8% | 20% | 5% | 40% | 400% | 20% | 1 | -1.41E-03 | 3.56E-05 | -1.46E-03 | 1.275 | 20% | 21% |
| 8% | 20% | 5% | 40% | 800% | 20% | 1 | -2.34E-03 | 4.02E-05 | -2.35E-03 | 0.201 | 84% | 84% |
| 8% | 20% | 5% | 40% | 400% | 40% | 1 | -4.08E-03 | 5.44E-05 | -4.17E-03 | 1.751 | 8% | 9% |
| 8% | 20% | 5% | 40% | 800% | 40% | 1 | -7.64E-03 | 5.18E-05 | -7.73E-03 | 1.605 | 11% | 12% |

r_f = riskless rate of interest; μ = diffusion statistical drift; σ = diffusion volatility; q = dividend yield;
 λ = arrival rate of jumps (per year); $\gamma = -\sigma^2/2$ = jump amplitude; δ = volatility of jump amplitude;
T = simulation horizon (years)
Experiment parameters: time step number = 90; simulation number = 50,000; resampling pool number = 10;
resampling simulation number = 100,000; down-sampling ratio = 10

Approximate p -values are computed and all of them are above 5%. Consequently, the null hypothesis cannot be rejected at the 5% level of significance. As expected, the t -tests and the z -tests yield similar approximations of the p -values.

7.3.4 Least-squares Monte Carlo for trigger boundary generation

The difficulty with analyzing path-dependent options is that: it requires the knowledge of the early-exercise boundary position; this, in turn, needs backward-type of analysis to estimate continuation values, but Monte Carlo simulation is a forward-type of analysis generating possible evolutions of the underlying from a starting point. The least-squares Monte Carlo method of Longstaff-Schwartz was previously identified as a promising method that resolves the aforementioned problem: the method first performs a regression of the continuation value using cross-sectional information from the simulation; this regressed continuation value is then compared with the immediate exercise value to derive the early-investment policy. In a real options framework, the continuation value of the business prospect is the value attained if the decision to invest is delayed by at least one time step. This value is compared to the value of the business prospect if the investment were to be made immediately. By performing a comparison for each trajectory and at each time step in the simulation, it is possible to construct the trigger policy and identify regions where early-exercise is optimal.

Research Question 1.2.1 – Early-investment boundaries to detect trigger events

How can early-investment boundaries be defined for real options featuring early-exercise possibilities and how can they be used for the identification of precursors of successful development programs?

Hypothesis 1.1.3 – Least-squares Monte Carlo approach for real options

Real options with early-exercise properties may be analyzed using a least-squares Monte Carlo approach to both estimate their value and derive the early-investment policy and the optimal-investment boundary. The optimal-investment boundary may be used by decision-makers to substantiate the need to invest now, delay, or abandon an investment.

Verification process and criteria for success

The verification procedure for the generation of the early-exercise boundary is different from what was done in the preceding sections. Indeed, the verification no longer tests the equality of distributions or the equality of properties of distributions but rather the similarity between curves. To test the similarity between the early-exercise boundary generated by the proposed method and a reference boundary obtained via a finite-difference scheme, four metrics are used for the different test cases. Two metrics are used to compare the global shape of the two curves; one metric tracks the maximum error between the curves; and one metric assesses the initial error which is the present time error. The first metric is the Hausdorff distance which measures how similar two curves are. This metric yields an absolute distance (not a relative one) which may be difficult to gauge. However, all test cases involve options with strike prices of one and thus trigger boundaries are never far from one. This means that there are no scaling issues and that normalization is not required. A test is considered successful if the distance between the two curves does not exceed 0.05. The second metric is the more usual root mean square error (RMSE) which is again not normalized due to the absence of scaling issues. A test is considered successful if the RMSE between the two curves is less than 0.05. The third metric is the maximum relative error which tracks the largest error between the two

curves. A test is considered successful if the absolute value of the maximum relative error is less than 5%. Finally, the last metric is the initial relative error between the two curves. Tracking this error is important since this error is contemporary to the immediate investment decision-making process. A test is considered successful if the absolute value of the initial relative error is less than 5%. All tests are performed for the geometric Brownian motion.

Similarity tests

The experimental early-exercise boundaries obtained via combined simulation, non-parametric Esscher transformation, bootstrapping, and least-squares Monte Carlo are compared to the corresponding reference early-exercise boundaries obtained via the solving of partial differential equations with a finite-difference solver. The geometric Brownian motion is simulated to generate 30,000 trajectories. These trajectories induce samples of 30,000 returns at each time step. Four of these samples are pooled together to yield a larger sample of 120,000 returns and a change of probability measure is performed using the non-parametric Esscher transform. The sample of 120,000 weighted returns is bootstrapped using the sampling wheel algorithm to yield 30,000 new trajectories. These trajectories are used in the improved least-squares Monte Carlo algorithm to generate the trigger boundary. Twenty different test cases of geometric Brownian motions are studied using typical parameters of real options analyses: the maturity is set at one year, the risk-free rate is set at 2%, the dividend yield is set at 5%, the drift rate is varied between 5% and 20%, the volatility is varied between 20% and 40%, and finally the spot to strike ratio is varied between 0.8 and 1.2. Each of these test cases is repeated thirty times to assess the variability of results, resulting in six hundred

trials. For each trial, the Hausdorff distance, the root mean square error, the maximum relative error magnitude, and the initial relative error magnitude are reported respectively in Figure 66, Figure 67, Figure 68, and Figure 69. Looking at these results indicates that for each similarity metric retained, all outcomes from the trials are below the verification threshold. With no trial failing any of the tests, the verification is considered successful.




Figure 66: Hausdorff distance between experimental and reference trigger boundaries for 600 cases of geometric Brownian motions




Figure 67: RMSE of the experimental trigger boundary for 600 cases of geometric Brownian motions




Figure 68: Maximum relative error of the experimental trigger boundary for 600 cases of geometric Brownian motions




Figure 69: Initial relative error of the experimental trigger boundary for 600 cases of geometric Brownian motions

7.3.5 Least-squares Monte Carlo for option valuation

Besides using the updated least-squares Monte Carlo method to construct the early-exercise boundary, the method can be used to carry out the valuation of real options featuring early-exercise possibilities, the valuation of real options without early-exercise possibilities, and the valuation of real options on underlying assets following exotic stochastic processes. Indeed, Monte Carlo simulations can generate many kind of trajectories including some representing the realization of complex stochastic processes such as those featuring jumps. At the same time, simpler Monte Carlo based methods can value vanilla European options, while the Monte Carlo based method of Longstaff-Schwartz resolves the problem of the continuation value estimation for path-dependent American options.

Research Question 1.1.1 — Enlarging the domain of applicability of real options

How can the domain of application of current state-of-the-art real options methodologies be extended to include corporate investments with value processes that do not follow classic geometric random walks?

Hypothesis 1.1.1 – Monte Carlo methods for real options analyses

Monte Carlo methods and lattice-based methods present the most promising approaches to solve for the arbitrage-free value of corporate investments featuring flexibility. Within the context of the aerospace industry, Monte Carlo methods offer the ability to integrate well with other probabilistic methods.

Verification process and criteria for success

The hypothesis states that real options can be evaluated using Monte-Carlo methods. Furthermore, the hypothesis states that using Monte Carlo simulation enables

the use of more sophisticated models that may be better suited to represent a complex reality. This is purely a technical claim that can be verified using canonical examples by checking the prices of options obtained with Monte Carlo based methods to the prices of options obtained with more established methods. These more established methods include the Black Scholes closed-form solution for European options on assets following a geometric Brownian motion, its modified version for European options on assets following a jump-diffusion process, as well as the solution obtained with a finite-difference scheme for the valuation of American options on assets following a geometric Brownian motion. Thus, this is a quantitative verification and a successful verification is achieved if the pricing of European and path-dependent American options is accurate and exhibits low variability during repeated experiments. The accuracy test is performed by straightforward comparisons. For the test to be successful, option prices have to be within 5% of the reference option price obtained with established methods. The repeatability is checked by repeating each test several times which by virtue of Monte Carlo simulations will lead to slightly different option prices. Obviously, less variability is better and for the test to be successful, the standard error (i.e. the standard deviation of the sample of option prices) must yield a confidence interval no larger than 10% of the option price. This will ensure that the real option price can be approximated consistently with the proposed approach. In order to prove that Monte Carlo based methods are more versatile and can handle a wider variety of stochastic processes than traditional methods, the tests are performed for two classes of stochastic processes: the traditional geometric Brownian motion as well as a more sophisticated jump-diffusion process. Indeed, typical real options methods are unable to handle processes with jumps or require extensive and

complex modifications while Monte Carlo based algorithm can indifferently handle cases with and without jumps.

Results

The experimental real option prices obtained via combined simulation, non-parametric Esscher transformation, bootstrapping, and least-squares Monte Carlo are compared to reference prices. A geometric Brownian motion is first simulated to generate 30,000 trajectories. These trajectories induce samples of 30,000 returns at each time step. Four of these samples are pooled together to yield a larger sample of 120,000 returns and a change of probability measure is performed using the non-parametric Esscher transform. The sample of 120,000 weighted returns is bootstrapped using the sampling wheel algorithm to yield 30,000 new trajectories. These trajectories are used in the improved least-squares Monte Carlo algorithm to obtain the value of the option. Twenty different test cases of geometric Brownian motions are studied using typical parameters of real options analyses: the maturity is set at one year, the risk-free rate is set at 2%, the dividend yield is set at 5%, the drift rate is varied between 5% and 20%, the volatility is varied between 20% and 40%, and finally the spot to strike ratio is varied between 0.8 and 1.2. Each of these test cases is repeated thirty times to assess the variability of results, resulting in six hundred trials. For each trial, the relative error of the European option price, the corresponding relative width of the 95% confidence interval of the European option price, the relative error of the American option price, and finally the relative width of the 95% confidence interval of the American option price are reported respectively in Figure 70, Figure 71, Figure 72, and Figure 73.




Figure 70: Relative error of European call option prices with underlying following a geometric Brownian motion (reference Black Scholes formula)




Figure 71: Relative width of 95% confidence interval for European call option prices on underlying following a geometric Brownian motion




Figure 72: Relative error of American call option prices with underlying following a geometric Brownian motions (reference finite-difference solver)

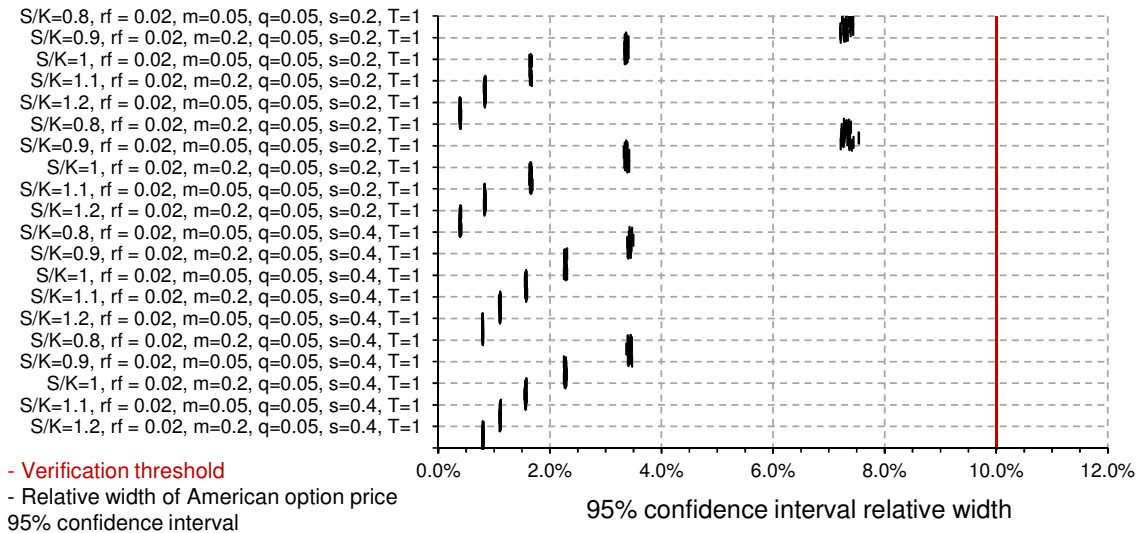


Figure 73: Relative width of 95% confidence interval for American call option prices on underlying following a geometric Brownian motion

First, a cursory check at the figures indicates striking similarities for the relative error of European and American options. This is to be expected as the same Monte Carlo simulation and therefore the same set of trajectories are used for the pricing of both types of options: because the prices of American and European options are naturally correlated

and because the same trajectories are used for the pricing exercise, results are not independent and some correlation in both the option prices and their errors is expected.

In addition, it appears that for the European and American options, there are two trials for which the option price relative error is beyond the 5% accuracy threshold. This is not much given the other 598 successful trials.

As for the analysis of variability, all trials exhibit 95% confidence intervals that have a relative width of less than 10% of the option price and therefore meet the verification criteria. Interestingly, the combination of low volatility and small spot to strike ratio seems to exhibit the widest relative confidence interval. The reason for this is that these test cases result in few trajectories being in the money and the Monte Carlo simulation based pricing scheme becomes less efficient (the more trajectories end up in the money, the better as more data points can be used to compute the expectation of the option payoffs). In any case, owing to the number of trials that pass the accuracy and the variability tests, the verification is considered successful.

Another batch of test is provided in Table 59 for European types of option. This time, the real option value is reported for twenty experiments with stochastic processes following geometric Brownian motions. For these experiments, the risk free rate of return is varied between 2% and 8%, the drift rate is varied between 5% and 20%, the dividend yield is varied between 0% and 5%, the volatility is varied between 20% and 40%, and the time step is varied between four and eight days. In each experiment, Monte Carlo simulations are used to generate 80,000 trajectories. A sample of 320,000 weighted returns is constructed from the pooling of 80,000 simulated returns from four time steps and subsequent weighting via the non-parametric Esscher transform. This weighted

distribution is bootstrapped to generate 20,000 trajectories. European option payoffs are estimated at maturity for each of these trajectories and discounted back to the present time. Most of the experimental results are very close to the expected theoretical results and the relative error never exceeds 2%.

Table 59: European option prices for twenty cases of geometric Brownian motions

| $\frac{S}{K}$ | r_f | μ | q | σ | T | Experimental European Option Price | Theoretical European Option Price | Relative Difference |
|---------------|-------|-------|-----|----------|-----|------------------------------------|-----------------------------------|---------------------|
| 1 | 2.0% | 5% | 0% | 20% | 1.0 | 0.0705 | 0.0694 | 1.71% |
| 1 | 2.0% | 20% | 0% | 20% | 1.0 | 0.0691 | 0.0694 | -0.33% |
| 1 | 2.0% | 5% | 0% | 40% | 1.0 | 0.1471 | 0.1472 | -0.09% |
| 1 | 2.0% | 20% | 0% | 40% | 1.0 | 0.1466 | 0.1472 | -0.46% |
| 1 | 8.0% | 5% | 0% | 20% | 1.0 | 0.0435 | 0.0442 | -1.50% |
| 1 | 8.0% | 20% | 0% | 20% | 1.0 | 0.0445 | 0.0442 | 0.66% |
| 1 | 8.0% | 5% | 0% | 40% | 1.0 | 0.1155 | 0.1170 | -1.27% |
| 1 | 8.0% | 20% | 0% | 40% | 1.0 | 0.1183 | 0.1170 | 1.11% |
| 1 | 2.0% | 5% | 0% | 20% | 2.0 | 0.0927 | 0.0917 | 1.05% |
| 1 | 2.0% | 20% | 0% | 20% | 2.0 | 0.0912 | 0.0917 | -0.64% |
| 1 | 2.0% | 5% | 0% | 40% | 2.0 | 0.1991 | 0.1993 | -0.08% |
| 1 | 2.0% | 20% | 0% | 40% | 2.0 | 0.2005 | 0.1993 | 0.60% |
| 1 | 8.0% | 5% | 0% | 20% | 2.0 | 0.0466 | 0.0463 | 0.59% |
| 1 | 8.0% | 20% | 0% | 20% | 2.0 | 0.0466 | 0.0463 | 0.67% |
| 1 | 8.0% | 5% | 0% | 40% | 2.0 | 0.1396 | 0.1403 | -0.49% |
| 1 | 8.0% | 20% | 0% | 40% | 2.0 | 0.1417 | 0.1403 | 1.00% |
| 1 | 2.0% | 5% | 5% | 20% | 1.0 | 0.0924 | 0.0923 | 0.14% |
| 1 | 2.0% | 20% | 5% | 20% | 1.0 | 0.0924 | 0.0923 | 0.09% |
| 1 | 2.0% | 5% | 15% | 20% | 1.0 | 0.1460 | 0.1480 | -1.30% |
| 1 | 2.0% | 20% | 15% | 20% | 1.0 | 0.1498 | 0.1480 | 1.23% |

S = underlying asset price; K = strike price; r_f = riskless rate of interest; μ = diffusion statistical drift; σ = diffusion volatility; q = dividend yield; T = option maturity (years)
 Experiment parameters: time step number = 90; simulation number = 80,000; resampling pooling = 4; resampling draws = 20,000; down-sampling ratio = 1

Monte Carlo simulations introduce some variability in the computation of option prices as new pseudo random number sequences are used each time an experiment is carried out. To check the robustness of these results, each experiment is now made of

thirty repeated trials which enable the computation of the sample mean, standard error, z -statistic and t -statistic. The null hypothesis for the z -test and t -test is that the experimental average option price is equal to the theoretical option price. Approximate p -values are computed and all but one of them are above the 5% significance level.

Table 60: z -test and t -test for European option – Repeated cases of geometric Brownian motions

| $\frac{S}{K}$ | r_f | μ | q | σ | T | Exp. Sample Mean European Option Price | Exp. Sample Standard Error | Theo. European Option Price | z -test and t -test statistic | z -test p -value | t -test p -value |
|---------------|-------|-------|-----|----------|---|--|----------------------------|-----------------------------|-----------------------------------|----------------------|----------------------|
| 1 | 2% | 5% | 0% | 20% | 1 | 6.95E-02 | 1.38E-04 | 6.94E-02 | 1.347 | 18% | 19% |
| 1 | 2% | 20% | 0% | 20% | 1 | 6.92E-02 | 1.01E-04 | 6.94E-02 | 1.508 | 13% | 14% |
| 1 | 2% | 5% | 0% | 40% | 1 | 1.47E-01 | 1.73E-04 | 1.47E-01 | 1.249 | 21% | 22% |
| 1 | 2% | 20% | 0% | 40% | 1 | 1.47E-01 | 2.13E-04 | 1.47E-01 | 0.173 | 86% | 86% |
| 1 | 8% | 5% | 0% | 20% | 1 | 4.41E-02 | 8.51E-05 | 4.42E-02 | 1.208 | 23% | 24% |
| 1 | 8% | 20% | 0% | 20% | 1 | 4.43E-02 | 8.98E-05 | 4.42E-02 | 1.461 | 14% | 15% |
| 1 | 8% | 5% | 0% | 40% | 1 | 1.17E-01 | 1.71E-04 | 1.17E-01 | 0.139 | 89% | 89% |
| 1 | 8% | 20% | 0% | 40% | 1 | 1.17E-01 | 2.04E-04 | 1.17E-01 | 1.819 | 7% | 8% |
| 1 | 2% | 5% | 0% | 20% | 2 | 9.15E-02 | 1.64E-04 | 9.17E-02 | 1.415 | 16% | 17% |
| 1 | 2% | 20% | 0% | 20% | 2 | 9.16E-02 | 1.64E-04 | 9.17E-02 | 0.700 | 48% | 49% |
| 1 | 2% | 5% | 0% | 40% | 2 | 1.99E-01 | 2.44E-04 | 1.99E-01 | 0.720 | 47% | 48% |
| 1 | 2% | 20% | 0% | 40% | 2 | 1.99E-01 | 2.85E-04 | 1.99E-01 | 0.608 | 54% | 55% |
| 1 | 8% | 5% | 0% | 20% | 2 | 4.64E-02 | 1.25E-04 | 4.63E-02 | 1.079 | 28% | 29% |
| 1 | 8% | 20% | 0% | 20% | 2 | 4.64E-02 | 1.13E-04 | 4.63E-02 | 0.652 | 51% | 52% |
| 1 | 8% | 5% | 0% | 40% | 2 | 1.40E-01 | 2.27E-04 | 1.40E-01 | 0.017 | 99% | 99% |
| 1 | 8% | 20% | 0% | 40% | 2 | 1.40E-01 | 2.37E-04 | 1.40E-01 | 0.068 | 95% | 95% |
| 1 | 2% | 5% | 5% | 20% | 1 | 9.21E-02 | 1.30E-04 | 9.23E-02 | 1.185 | 24% | 25% |
| 1 | 2% | 20% | 5% | 20% | 1 | 9.26E-02 | 1.39E-04 | 9.23E-02 | 2.086 | 4% | 5% |
| 1 | 2% | 5% | 15% | 20% | 1 | 1.48E-01 | 1.44E-04 | 1.48E-01 | 0.237 | 81% | 81% |
| 1 | 2% | 20% | 15% | 20% | 1 | 1.48E-01 | 1.89E-04 | 1.48E-01 | 0.175 | 86% | 86% |

r_f = riskless rate of interest; μ = diffusion statistical drift; σ = diffusion volatility;
 q = dividend yield; T = simulation horizon (years)
Experiment parameters: time step number = 90; simulation number = 80,000;
resampling pooling = 4; resampling draws = 20,000; down-sampling ratio = 1

How to deal with this failed experiment? Is this a fluke or a symptom of something more profound? In order to answer this question, the history surrounding this experiment is revisited. It indicates that the p -values for the t -test and z -test comparing

the sample mean expected value to the theoretical expected value were not abnormally low (95% for both tests). The average option price is also not really off with an error of just over 0.3%. However, the standard error is rather low (standard deviation of 8.03E-04 leading to a standard error of 1.47E-04 for an option price of 9.26E-02). Consequently, this failed experiment is the combination of a high biased option price and low variability between the thirty results. It is believed that this is a one-case event due to a single outlying option price within the thirty trials probably resulting from a bad seed in one Monte Carlo simulation. To verify this assertion, the experiment consisting of thirty trials is run again and the new p -values of 53% is indeed much larger than the 5% level of significance as highlighted in Table 61. In addition, it is not statistically unlikely that one experiment fails out of a total of twenty experiments since that represents exactly 5% of all experiments.

Consequently, the null hypothesis is not rejected at the 5% level of significance. As expected, the t -test and the z -test yield similar approximations of the p -values since both tests are applicable “in the limit”.

Table 61: New z -test and t -test for the European option price of the repeated experiment

| $\frac{S}{K}$ | r_f | μ | q | σ | T | Exp. Sample Mean European Option Price | Exp. Sample Standard Error | Theo. European Option Price | z -test and t -test statistic | z -test p -value | t -test p -value |
|---|-------|-------|-----|----------|---|--|----------------------------|-----------------------------|-----------------------------------|----------------------|----------------------|
| 1 | 2% | 20% | 5% | 20% | 1 | 9.24E-02 | 1.33E-04 | 9.23E-02 | 0.633 | 53% | 53% |
| r_f = riskless rate of interest; μ = diffusion statistical drift; σ = diffusion volatility; q = dividend yield; T = simulation horizon (years) Experiment parameters: time step number = 90; simulation number = 80,000; resampling pooling = 4; resampling draws = 20,000; down-sampling ratio = 1 | | | | | | | | | | | |

Next, a jump-diffusion process is first simulated to generate 50,000 trajectories. These trajectories induce samples of 50,000 returns at each time step. Ten of these samples are pooled together to yield a larger sample of 500,000 returns and a change of probability measure is performed using the non-parametric Esscher transform. The sample of 500,000 weighted returns is bootstrapped using the sampling wheel algorithm to yield 50,000 new trajectories. These trajectories are used in the improved least-squares Monte Carlo algorithm to obtain the value of the option. Twenty different test cases of jump-diffusion are studied with emphasis on parameters governing jumps: the maturity is set at one year, the risk-free rate is set at 2%, the dividend yield is set at 5%, the drift rate is varied between 5% and 20%, the volatility is varied between 20% and 40%, and finally the spot to strike ratio is varied between 0.8 and 1.2. Each of these test cases is repeated thirty times to assess the variability of results, resulting in six hundred trials. For each trial, the relative error of the European option price, the corresponding relative width of the 95% confidence interval of the European option price, and finally the relative width of the 95% confidence interval of the American option price are reported respectively in Figure 74, Figure 75 and Figure 76.

Despite the traditional difficulty in simulating rare events with Monte Carlo simulations, the results are very good. A vast majority of tests exhibits relative errors below the threshold retained and no test exceeds the threshold retained for the relative width of the 95% confidence interval. One of the reasons for this successful verification is the use of Monte Carlo simulations with a high number of replications to ensure that jumps are reasonably well represented in the sample that is used for the bootstrap resampling procedure.




Figure 74: Relative error of European real call option prices with underlying following a Merton jump diffusion process (reference modified Merton-Black Scholes formula)




Figure 75: Relative width of 95% confidence interval for European real call option prices on underlying following a Merton jump-diffusion process




Figure 76: Relative width of 95% confidence interval for American call option prices on underlying following a Merton jump diffusion process

While this set of experiments aimed primarily at quantifying errors, another set of experiments is run with a large number of simulations. Table 62 corresponds to a first batch of twenty experiments with stochastic processes following a Merton jump diffusion process. Emphasis is put on varying parameters governing jumps with a jump arrival rate varied between 200% and 800% and a jump size volatility varied between 20% and 40%.

The risk free rate of return is still varied between 2% and 8%, the drift rate is still varied

between 5% and 20%, the dividend yield is varied between 0% and 5%, and the time step is varied between four and eight days. In each experiment, Monte Carlo simulations are used to generate 50,000 trajectories. A sample of 500,000 weighted returns is constructed from the pooling of 50,000 simulated returns from ten time steps and subsequent weighting via the non-parametric Esscher transform. This weighted distribution is bootstrapped to generate 100,000 trajectories. European option payoffs are estimated at maturity for each of these trajectories and discounted back to the present time.

Table 62: European option prices for twenty cases of Merton jump diffusion processes

| $\frac{S}{K}$ | r_f | μ | q | σ | λ | δ | T | Experimental European Option Price | Theoretical European Option Price | Relative Difference |
|---------------|-------|-------|------|----------|-----------|----------|-----|------------------------------------|-----------------------------------|---------------------|
| 1 | 2.0% | 5% | 0.0% | 20% | 400% | 20% | 1.0 | 0.1607 | 0.1618 | -0.67% |
| 1 | 2.0% | 5% | 0.0% | 20% | 600% | 20% | 1.0 | 0.1929 | 0.1935 | -0.29% |
| 1 | 2.0% | 5% | 0.0% | 20% | 400% | 40% | 1.0 | 0.2977 | 0.2962 | 0.51% |
| 1 | 2.0% | 5% | 0.0% | 20% | 600% | 40% | 1.0 | 0.3571 | 0.3604 | -0.92% |
| 1 | 2.0% | 15% | 0.0% | 20% | 400% | 20% | 1.0 | 0.1634 | 0.1637 | -0.21% |
| 1 | 2.0% | 15% | 0.0% | 20% | 600% | 20% | 1.0 | 0.1927 | 0.1947 | -1.05% |
| 1 | 2.0% | 15% | 0.0% | 20% | 400% | 40% | 1.0 | 0.2985 | 0.2978 | 0.25% |
| 1 | 2.0% | 15% | 0.0% | 20% | 600% | 40% | 1.0 | 0.3588 | 0.3612 | -0.66% |
| 1 | 2.0% | 5% | 0.0% | 20% | 400% | 20% | 2.0 | 0.2240 | 0.2217 | 1.03% |
| 1 | 2.0% | 5% | 0.0% | 20% | 600% | 20% | 2.0 | 0.2646 | 0.2645 | 0.03% |
| 1 | 2.0% | 5% | 0.0% | 20% | 400% | 40% | 2.0 | 0.4065 | 0.4048 | 0.41% |
| 1 | 2.0% | 5% | 0.0% | 20% | 600% | 40% | 2.0 | 0.4853 | 0.4845 | 0.15% |
| 1 | 2.0% | 15% | 0.0% | 20% | 400% | 20% | 2.0 | 0.2239 | 0.2242 | -0.13% |
| 1 | 2.0% | 15% | 0.0% | 20% | 600% | 20% | 2.0 | 0.2661 | 0.2661 | 0.00% |
| 1 | 2.0% | 15% | 0.0% | 20% | 400% | 40% | 2.0 | 0.4066 | 0.4066 | 0.00% |
| 1 | 2.0% | 15% | 0.0% | 20% | 600% | 40% | 2.0 | 0.4857 | 0.4855 | 0.05% |
| 1 | 8.0% | 20% | 5.0% | 40% | 400% | 20% | 1.0 | 0.1954 | 0.1937 | 0.87% |
| 1 | 8.0% | 20% | 5.0% | 40% | 800% | 20% | 1.0 | 0.2392 | 0.2386 | 0.25% |
| 1 | 8.0% | 20% | 5.0% | 40% | 400% | 40% | 1.0 | 0.3029 | 0.3029 | 0.00% |
| 1 | 8.0% | 20% | 5.0% | 40% | 800% | 40% | 1.0 | 0.4029 | 0.4030 | -0.03% |

r_f = riskless rate of interest; μ = diffusion statistical drift; σ = diffusion volatility; q = dividend yield;
 λ = arrival rate of jumps (per year); $\gamma = -\sigma^2/2$ = jump amplitude; δ = volatility of jump amplitude;
T = simulation horizon (years)
Experiment parameters: time step number = 90; simulation number = 50,000; resampling pooling = 10;
resampling draws = 100,000; down-sampling ratio = 10

Most of the experimental results are very close to the expected theoretical results and the relative error never exceeds 1.1%. Monte Carlo simulations introduce some variability in the computation of option prices and therefore the robustness of these results is checked. Each experiment is now made of thirty repeated trials which enable the computation of the sample mean, standard error, z -statistic and t -statistic. The null hypothesis for the z -test and t -test is that the experimental average option price is equal to the theoretical option price. Approximate p -values are computed and all of them are above the 5% significance level. Consequently, the null hypothesis cannot be rejected at the 5% level of significance. As expected, the t -test and the z -test yield similar approximations of the p -values since both tests are applicable “in the limit”.

From these tables, the results are better and there are several reasons for this observation: first, the number of replications has been significantly increased, going from 30,000 replications to 100,000 replications in order to increase the likelihood of capturing jumps in the simulation, second, the arrival rate of jumps has been increased to go from between one and two in the first experiment to between four and eight. In the end, if jumps are to be modeled it is probably because these jumps are quite frequent and it is questionable whether having a jump-diffusion process with a single jump per year is worth investigating. In any cases, increasing the number of simulation trajectories has resolved the accuracy and variability issues first encountered. As a result, the verification of the hypothesis is successful.

Table 63: z-tests and t-tests for European options – Repeated cases of Merton jump diffusion processes

| $\frac{S}{K}$ | r_f | μ | q | σ | λ | δ | T | Exp. Sample Mean European Option Price | Exp. Sample Standard Error | Theo. European Option Price | z-test and t-test statistic | z-test p-value | t-test p-value |
|---------------|-------|-------|------|----------|-----------|----------|-----|--|----------------------------|-----------------------------|-----------------------------|----------------|----------------|
| 1 | 2.0% | 5% | 0.0% | 20% | 400% | 20% | 1.0 | 1.62E-01 | 1.67E-04 | 1.62E-01 | 0.651 | 52% | 52% |
| 1 | 2.0% | 5% | 0.0% | 20% | 600% | 20% | 1.0 | 1.94E-01 | 1.90E-04 | 1.93E-01 | 0.795 | 43% | 43% |
| 1 | 2.0% | 5% | 0.0% | 20% | 400% | 40% | 1.0 | 2.96E-01 | 3.10E-04 | 2.96E-01 | 0.249 | 80% | 80% |
| 1 | 2.0% | 5% | 0.0% | 20% | 600% | 40% | 1.0 | 3.60E-01 | 3.24E-04 | 3.60E-01 | 0.102 | 92% | 92% |
| 1 | 2.0% | 15% | 0.0% | 20% | 400% | 20% | 1.0 | 1.63E-01 | 1.44E-04 | 1.64E-01 | 1.532 | 13% | 14% |
| 1 | 2.0% | 15% | 0.0% | 20% | 600% | 20% | 1.0 | 1.94E-01 | 1.93E-04 | 1.95E-01 | 1.257 | 21% | 22% |
| 1 | 2.0% | 15% | 0.0% | 20% | 400% | 40% | 1.0 | 2.97E-01 | 2.30E-04 | 2.98E-01 | 1.806 | 7% | 8% |
| 1 | 2.0% | 15% | 0.0% | 20% | 600% | 40% | 1.0 | 3.61E-01 | 3.22E-04 | 3.61E-01 | 0.474 | 64% | 64% |
| 1 | 2.0% | 5% | 0.0% | 20% | 400% | 20% | 2.0 | 2.22E-01 | 1.72E-04 | 2.22E-01 | 0.274 | 78% | 79% |
| 1 | 2.0% | 5% | 0.0% | 20% | 600% | 20% | 2.0 | 2.65E-01 | 1.64E-04 | 2.64E-01 | 0.735 | 46% | 47% |
| 1 | 2.0% | 5% | 0.0% | 20% | 400% | 40% | 2.0 | 4.05E-01 | 2.73E-04 | 4.05E-01 | 0.512 | 61% | 61% |
| 1 | 2.0% | 5% | 0.0% | 20% | 600% | 40% | 2.0 | 4.84E-01 | 2.55E-04 | 4.85E-01 | 1.562 | 12% | 13% |
| 1 | 2.0% | 15% | 0.0% | 20% | 400% | 20% | 2.0 | 2.24E-01 | 2.18E-04 | 2.24E-01 | 0.149 | 88% | 88% |
| 1 | 2.0% | 15% | 0.0% | 20% | 600% | 20% | 2.0 | 2.66E-01 | 1.66E-04 | 2.66E-01 | 0.470 | 64% | 64% |
| 1 | 2.0% | 15% | 0.0% | 20% | 400% | 40% | 2.0 | 4.06E-01 | 3.02E-04 | 4.07E-01 | 0.343 | 73% | 73% |
| 1 | 2.0% | 15% | 0.0% | 20% | 600% | 40% | 2.0 | 4.85E-01 | 2.79E-04 | 4.85E-01 | 0.289 | 77% | 77% |
| 1 | 8.0% | 20% | 5.0% | 40% | 400% | 20% | 1.0 | 1.94E-01 | 1.54E-04 | 1.94E-01 | 0.275 | 78% | 79% |
| 1 | 8.0% | 20% | 5.0% | 40% | 800% | 20% | 1.0 | 2.39E-01 | 1.58E-04 | 2.39E-01 | 1.632 | 10% | 11% |
| 1 | 8.0% | 20% | 5.0% | 40% | 400% | 40% | 1.0 | 3.03E-01 | 3.34E-04 | 3.03E-01 | 0.440 | 66% | 66% |
| 1 | 8.0% | 20% | 5.0% | 40% | 800% | 40% | 1.0 | 4.03E-01 | 2.89E-04 | 4.03E-01 | 0.208 | 84% | 84% |

r_f = riskless rate of interest; μ = diffusion statistical drift; σ = diffusion volatility; q = dividend yield;
 λ = arrival rate of jumps (per year); $\gamma = -\sigma^2/2$ = jump amplitude; δ = volatility of jump amplitude;
T = simulation horizon (years)
Experiment parameters: time step number = 90; simulation number = 50,000; resampling pooling = 10;
resampling draws = 100,000; down-sampling ratio = 10

7.3.6 Real options for staggered development valuation

Besides using the updated least-squares Monte Carlo method to analyze path-dependent options, the method can also be modified and used to carry out the valuation of compound or nested real options. A compound option is an option on another option: this means that the exercise of one option gives the possibility to exercise another option

later. Compound options are suitable for the analysis of staggered development programs with decision tollgates where the decision to fund one phase of the program opens the possibility of funding the following phase of the program. Compound options can be of the European or American type. For real options applications in the aerospace industry, compound options of the European-European and American-European types are probably more appropriate as the initial funding decision may be rushed but subsequent phases (detailed design, certification and testing) cannot really be started before the previous phase is completed. This leads to the following method hypothesis. However, only a single aspect of this method hypothesis is verified in this section, namely that real options can be used to analyze staggered investments with several decision tollgates. This partial verification provides the ground for a proper validation in Chapter 9.

Hypothesis 1 — Real options for valuation with flexibility and uncertainty

Within the context of aerospace research and development programs, real options methods enable the development of value-based design frameworks accounting for the staggered nature of investments and the value created by managerial flexibility in uncertain environments.

Verification process and criteria for success

One technical aspect of this hypothesis is that real options can be used to evaluate staggered investments. This technical claim can be verified using canonical examples by checking the prices of options obtained with the proposed method against the prices of options obtained with a more established method. In particular, Geske [203] provides a

formula to value European-European compound options on assets following a geometric Brownian motion and this formula is used for verification purposes.

Unfortunately, American-European options seem to have been left out of the literature and no benchmark can be found to check the results of the proposed model. Similarly, no formula is available for the pricing of compound options on underlying assets following a Merton jump diffusion process. Therefore, a heuristic verification is used for these cases. The absence of analytical formula for these cases is probably due to the complexity of compound options. Indeed, unlike simpler options where the underlying can be modeled directly, the underlying used for the first option is the nested option process which is a complex stochastic process. In particular, this process is not stationary since its volatility increases as the option maturity gets closer. This means that the proposed methodology needs to be updated to account for the non-stationarity of the underlying process¹. In addition, since the underlying for the first option is the nested option process, its value (i.e. the value of the nested option) must be computed at each time step for every trajectory in the Monte Carlo simulation. This is equivalent to performing nested Monte Carlo simulations which are notoriously computationally demanding.

To circumvent this problem, the continuity of option prices with respect to the underlying is invoked to reduce the number of computations performed. At each time cross-section in the Monte Carlo simulation, the minimum and maximum values of the nested option underlying asset are identified. These values of the underlying asset yield a

¹ Sample of returns under the equivalent martingale measure can no longer be pooled across different time steps as the properties of the samples are different. The non-parametric Esscher transform and the bootstrap resampling procedure must be applied individually at each and every time step.

range which is then evenly discretized to get evenly-spaced points. At each of these points, the value of the nested option is computed. Once the value of the nested option is known for each of the evenly-distributed points, the option price for all other points can be interpolated using a simple linear regression. The procedure is repeated next for each and every time step in the simulation.

However, for European-European options on assets following a geometric Brownian motion, the availability of a formula enables a quantitative verification and a successful verification is achieved if the pricing is accurate and exhibits low variability during repeated experiments. The accuracy test is performed by straightforward comparisons. For the test to be successful, option prices have to be within 5% of the reference option price obtained with the Geske analytical formula. The repeatability is checked by repeating each test several times which, by virtue of Monte Carlo simulations, leads to slightly different option prices. Obviously, less variability is better and for the test to be successful, the standard error must yield a confidence interval no larger than 10% of the option price. This will ensure that the real option price can be approximated consistently with the proposed approach.

Results

The experimental real option prices obtained via combined simulation, non-parametric Esscher transformation, bootstrapping, and least-squares Monte Carlo are compared to reference prices. A geometric Brownian motion is first simulated to generate 100,000 trajectories. These trajectories induce samples of 100,000 returns at each time step. A change of probability measure is performed using the non-parametric Esscher transform. The sample of 100,000 weighted returns is bootstrapped using the sampling

wheel algorithm to yield 50,000 new trajectories. These trajectories are used in the improved least-squares Monte Carlo algorithm to obtain the value of the option. Twenty five different test cases of geometric Brownian motions are studied using typical parameters of real options analyses: the maturities are set to one year and two years, the risk-free rate is set at 2%, the dividend yield is set at 5%, the drift rate is varied between 5% and 20%, the volatility is varied between 20% and 40%, the spot is varied between 1.7 and 2.3, and the strikes are set between 0.8 and 1.2 for the two options. Each of these test cases is repeated thirty times to assess the variability of results, resulting in seven hundred and fifty trials. For each trial, the relative error of the European option price and the corresponding relative width of the 95% confidence interval of the European option price are reported respectively in Figure 77 and Figure 78.

Investigations of the results indicate that no trial yields a relative error larger than 5%. As for the analysis of variability, all trials exhibit 95% confidence intervals that have a relative width of less than 10% of the option price and therefore meet the verification criteria. Interestingly, small ratios of the spot price to the sum of the two strike prices seem to exhibit the widest relative confidence interval. The reason for this is that these test cases have few trajectories being in-the-money and the Monte Carlo simulation-based pricing scheme becomes less efficient. In any case, owing to the number of trials that pass the accuracy and the variability tests, the verification is considered successful.

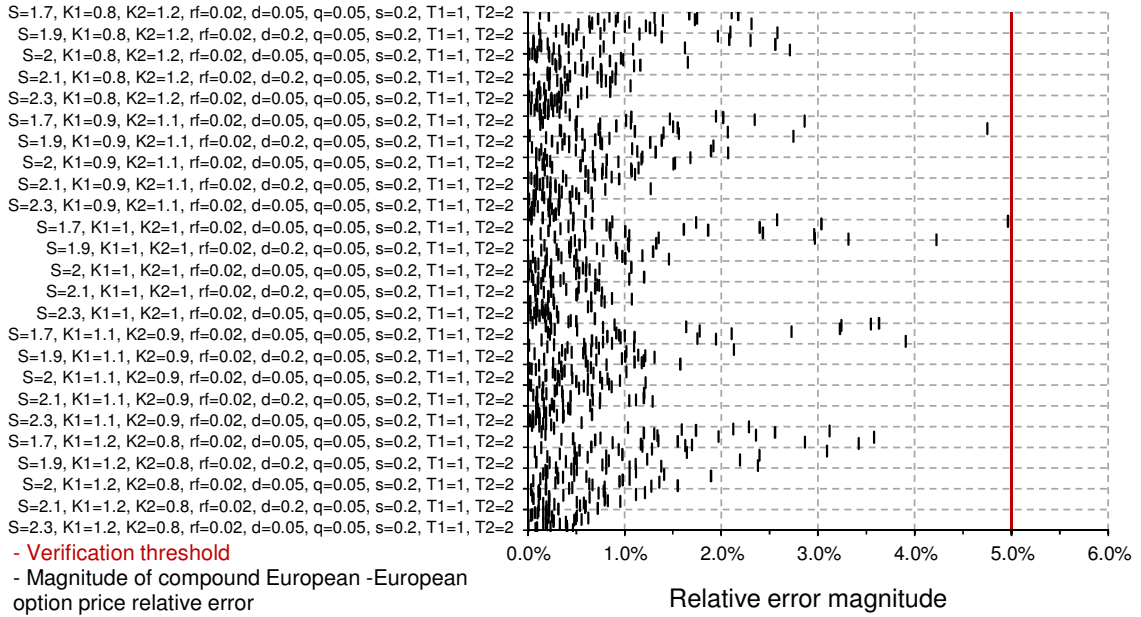


Figure 77: Relative error of compound European-European call option prices with underlying following a geometric Brownian motion (reference Geske formula)

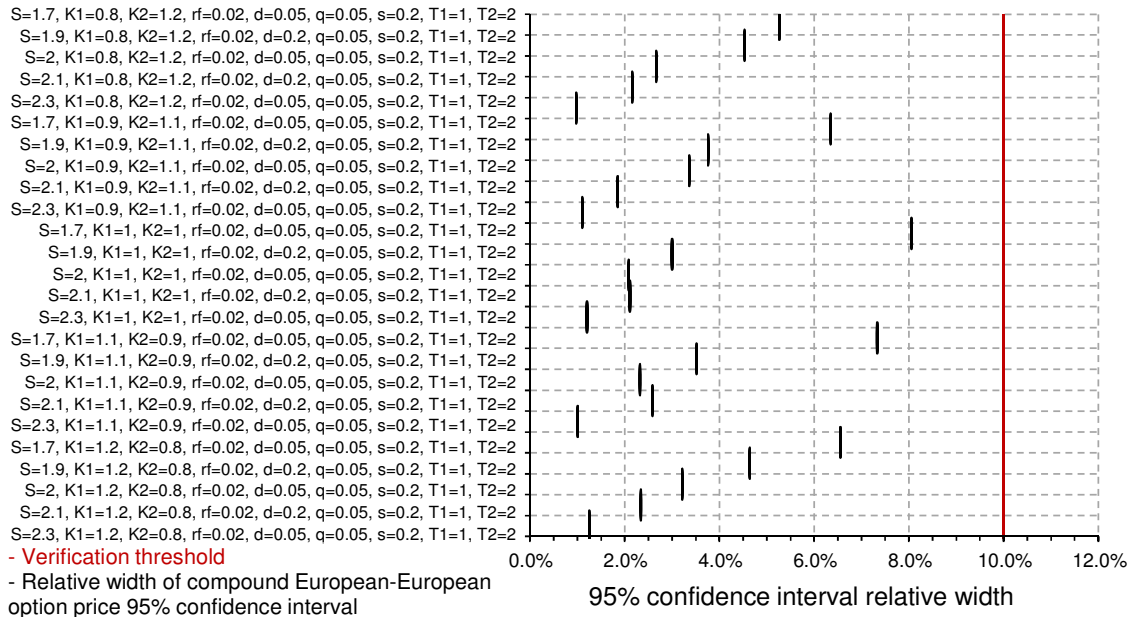


Figure 78: Relative width of 95% confidence interval for compound European-European call option prices on underlying following a geometric Brownian motion

CHAPTER 8: PROOF OF CONCEPT DESCRIPTION AND IMPLEMENTATION

Following the development of a novel real options methodology and the verification of the technical hypotheses, this section introduces a proof-of-concept study that is used to validate the method. The aim of this study is to demonstrate the applicability of the proposed methodology to real-world problems to solve relevant issues faced by decision-makers. Doing so, the proof-of-concept study is used to validate some of the method and modeling hypotheses formulated previously. The study highlights some typical issues faced by decision-makers when developing technology retrofits such as Performance Improvement Packages (PIP). Performance improvement packages are proposed by manufacturers as a means to improve the operating economics of currently out of production aircraft. This chapter starts with a description of the technologies featured in the package and describes briefly the development timeline and the competitive setting. Next, to analyze the economic viability of the technology retrofit, the operating economics of an aircraft and engine combination are studied using the point of view of aircraft operators. Owing to the lack of publicly available tool to perform this analysis, a new aircraft and engine evaluation model dubbed *i-CARE* for Integrated Cost And Revenue Estimation is developed. Because *i-CARE* is at the core of the proof-of-concept application, significant efforts are made to ensure that the model is properly calibrated and verified. The research proceeds next with the development of a simple market model to estimate the adoption of the performance improvement package by operators worldwide. Finally, the section concludes with the identification and calibration

of market uncertainties that impact most the viability of the performance improvement package development.

8.1 Presentation of an industry problem to be investigated

In this pilot study, a Performance Improvement Package (PIP) is being proposed as a means to improve the operating economics of a currently out of production aircraft engine. The engine manufacturer has identified a gap in its development stream which makes it possible to develop, certify, and produce the package. Decision-makers have to identify whether the conditions are currently optimal for the commercial launch of this product and whether it makes sense to commit resources to this development now. If not, there is a wide window to actually launch the development. The manufacturer can then delay the launch of the package development in order to wait for trigger events that will ensure that the development program has a high likelihood of commercial success.

8.1.1 Performance Improvement Package (PIP)

Performance Improvement Packages are nothing new in the aircraft and engine manufacturing industry and have been often proposed to operators as a stop-gap measure to improve the economics of aircraft currently on the market. For instance, *McDonnell Douglas* introduced a series of PIP [161] in the 1990's to improve the aerodynamics, reduce the drag, and improve the fuel-burn of its flagship MD-11 aircraft as the aircraft and its engines were not meeting promised specifications at entry into service. These packages can also be used to rejuvenate a design that is slowly aging by infusing some refinements to keep it competitive. Following experimentations with the *Tech56* technology demonstrator, *CFM International* announced in 2007 the first delivery of a

Tech Insertion package, and in 2011, announced the availability of a new performance improvement package for the CFM56-5B3 [162]. Both of these aimed at reducing NO_x emissions, improving fuel-burn of the engine, and extending its time on-wing (flight time without any engine shop visit). This was achieved by replacing some modules in the engine with newer ones: these changes included a new combustor, new 3D-shaped turbine blades, and tighter tolerances for the fan skew angles. Similarly in 2009, *Boeing* introduced refinements to its 777 flagship aircraft that airlines could buy as a package to increase the aircraft range and payload [163]. This was achieved by reshaping vortex generators on the upper surface of the wing, optimizing the ram air intake system to reduce drag, as well as drooping ailerons by two degrees while in flight. More recently, *Airbus* launched in 2013 the *Sharklet* retrofit [164] for already in-service aircraft of the A320 family. This retrofit consists in new advanced wingtip devices to reduce fuel-burn by up to four percent, reduce carbon emissions, and increase the operating life of the aircraft.

In the pilot study under investigation, an engine manufacturer is investigating the potential development of a performance improvement retrofit package for one of its engine. The package includes several technologies to reduce maintenance costs and decrease fuel-burn and carbon emissions. There are three reasons motivating this proposed development: demand by airlines for more efficient aircraft and engines to reduce their exposure to fluctuating energy prices, desire by manufacturers to increase the operating life of their engines by making them more competitive with other offerings from the competition, and identification of a gap in the development stream that needs to be filled.

8.1.2 Performance Improvement Package development timeline

Like many developments in the aerospace industry, the program is articulated around several phases, each separated by a milestone during which a technical and market review is performed. The purpose of these milestones is to monitor the development program and decide whether funding the subsequent phase of development is economically viable. In this research, the development is articulated around four distinct phases: a marketing and conceptual study phase, a preliminary design phase, a detailed design phase, and finally the production phase. Because of the numerous uncertainties surrounding the development, the manufacturer hesitates between two developments: a light performance improvement package (dubbed PIP-Light) featuring mostly mature off-the-shelves technologies, and a more involved performance improvement package (dubbed PIP-Involved) featuring new state-of-the-art technologies. The two packages require different development efforts and therefore have different development timelines as highlighted in Figure 79.




Figure 79: Development timeline for PIP-Light and PIP-Involved

8.1.3 Technology description

Two performance improvement packages are under analysis. The first retrofit package, denoted PIP-Light, is a lower scope improvement. The PIP-Light retrofit benefits from a short time to market and a cheaper acquisition cost by infusing into the engine new technologies that have matured recently. It is rushed to the market in order to counter the perceived threat of other engine manufacturers willing to enter the lucrative engine maintenance and part replacement business. Indeed, engine manufacturers make significant amount of money by servicing and providing replacement parts for the engines they manufacture. However, other manufacturers and maintenance providers are technically able to design replacement parts and sell them at a cheaper price in order to undercut the original equipment manufacturer. Indeed, under the Parts Manufacturer Approval (PMA), any manufacturer can get design and production approvals for modification and replacement parts to be sold and installed on type-certificated aircraft and engines. In fact, as long as PMA manufacturers do not infringe on patents, these replacement parts can be installed in lieu of the original, provided proof is given that the part is as good as or better than the original one. Furthermore, *Aircraft Commerce* [208] quotes Rob Baumann of HEICO Parts Group who indicates that PMA parts “*start at 50-60% of the OEMs [price] but sometimes are as low as 25%, meaning a discount of 75%*”.

As a result, the PIP-Light is introduced as a quick upgrade of the turbofan engine to pre-empt potential competition by offering replacement parts that improve the operating economics of the turbofan engine. The limited-scope retrofit brings noticeable, yet limited, improvements by infusing off-the-shelf technologies that can be brought to

the market within two years depending on market conditions. In some sense, the retrofit is similar to the “*Time On Wing*” upgrade package [209] offered by *CFMI* to CFM56 operators and to the “*Select One*” upgrade package [210] offered by *IAE* to V2500 operators during the first decade of the XXIst century.

The second retrofit package, denoted PIP-Involved, is a larger scope improvement. It brings significant improvements to the turbofan engine by infusing state-of-the-art technologies. However, technologies for this retrofit package are not yet mature and it is estimated that the package will be available in five to six years. This retrofit package requires more development time, is more expensive, but also provides significant maintenance cost reductions, fuel savings, and reduces carbon and nitrous oxides emissions. In some sense, this retrofit is similar to the “*Tech Insertion*” package [209] offered by *CFMI* to CFM56 operators or to the “*Phoenix Standard*” package [210] offered by *IAE* to V2500 operators.

Performance Improvement Package “Light” description

The PIP-Light retrofit is lower scope with a short time to market and a cheaper acquisition cost. The retrofit is made of several technologies that can be installed during a major maintenance event, such as a shop visit when the engine is removed from the wing. Its purpose is to improve the economics of the turbofan engine at a minimum capital expenditure for the airline. Improving the economics consists in reducing the direct operating cost by reducing fuel-burn, emissions, and maintenance costs. The focus of the modification is on the high pressure compressor and high pressure turbine. The high pressure compressor benefits from improved blade leading edges, improved surface coatings, and 3-D airfoil designs. This enables an increase in the exhaust gas temperature

margin. The high pressure turbine benefits from a new material, and from redistributed internal cooling which enables another modest increase of the exhaust gas temperature margin. In addition, an improved tip shelf and a stronger material to produce blades resistant to airfoil untwist and airfoil distortion enable a further reduction of the deterioration rate of the exhaust gas temperature margin over time. All in all, the performance improvements offered by the light package are summarized in Table 64.

Table 64: PIP-light key metrics with respect to current baseline turbofan engine

| | |
|-----------------------------|-------|
| Initial/Restored EGT margin | +5°C |
| EGT margin degradation rate | -5% |
| Specific Fuel Consumption | -0.8% |
| PIP-Light price (2014-US\$) | 0.3M |

Performance Improvement Package “Involved” description

The PIP-Involved retrofit is a significant improvement that increases substantially the value of the turbofan engine by infusing several cutting edge technologies to reduce fuel-burn, cut carbon emissions, and stretch the on-wing time so as to reduce maintenance expenditures. The retrofit can be installed during a major maintenance event, such as a shop visit during which the engine is removed from the wing. The focus of the modification is on several aspects: updates to the high pressure compressor, updates to the high pressure turbine, and updates to the low pressure turbine. The retrofit builds upon the technologies developed for the PIP-Light but goes further. It includes improved materials for the turbine in order to increase the exhaust gas temperature margin, an improved cooling scheme in order to further increase the exhaust gas temperature margin, but also improved blade leading edges, new 3-D airfoil designs, as well as a new multilayer erosion-resistant thermal barrier coating with improved rub-in capabilities.

In addition to these improvements, the PIP-Involved tries to address the loss of efficiency over time due to increasing blade tip clearances. Indeed, increasing clearance between the tip of blades and the shroud leads to a reduced stage efficiency and therefore to a the reduction in the exhaust gas temperature margin over time [211]. In order to mitigate the increase in blade tip clearance, an improved clearance control system is installed. A clearance control system consists of a valve adjusted in real-time by the engine Full Authority Digital Engine Control (FADEC) to control the mix of hot air from the compressor and cold air from the bypass duct. This airflow is subsequently circulated in tubes surrounding the casing of each turbine stage in order to control the casing temperature. By doing so, the air expands or contracts the turbine casing which helps control the clearances between the casing and the tip of blades. This improved model-based control software uses the engine operating information to estimate clearances and to improve the gas path sealing [212]. All of this leads to better overall performance and better performance retention over time. The performance improvements offered by the PIP-involved package are summarized in Table 65.

Table 65: PIP-Involved key metrics with respect to current baseline turbofan engine

| | |
|--------------------------------|-------|
| Initial/Restored EGT margin | +15°C |
| EGT margin degradation rate | -8% |
| Specific Fuel Consumption | -1.8% |
| PIP-Involved price (2014-US\$) | 1.6M |

Competing Performance Improvement Package using PMA

The original equipment manufacturer anticipates that a competing maintenance, repair, and overhaul provider will start offering a competing package available for retrofit starting in 2019. The competing package falls somewhere in between the PIP-Light and the PIP-Involved packages, although at a significant discount.

8.1.4 Competitive setting and technology development timeline

The previous paragraph described the two performance improvement packages that may be developed by the engine manufacturer. Due to limited resources, these two development streams are exclusive of one another and decision-makers need to assess the merit of each alternative in order to select one development strategy. Three potential moves have been identified by the engine manufacturer: do nothing and run the risk of losing market share on profitable servicing of the fleet of turbofan engines, develop the low-cost PIP-Light in order to try to preempt competition but still face the risk of having a competitor with a better product in the future, or finally, develop the more expensive PIP-Involved but face a longer development phase during which a significant opportunity cost will be experienced owing to reduced sales during development and reduced market size once the retrofit package reaches the market. A timeline representing the different scenarios and the different moves that the engine manufacturer and its competitor may make is proposed in Figure 80.

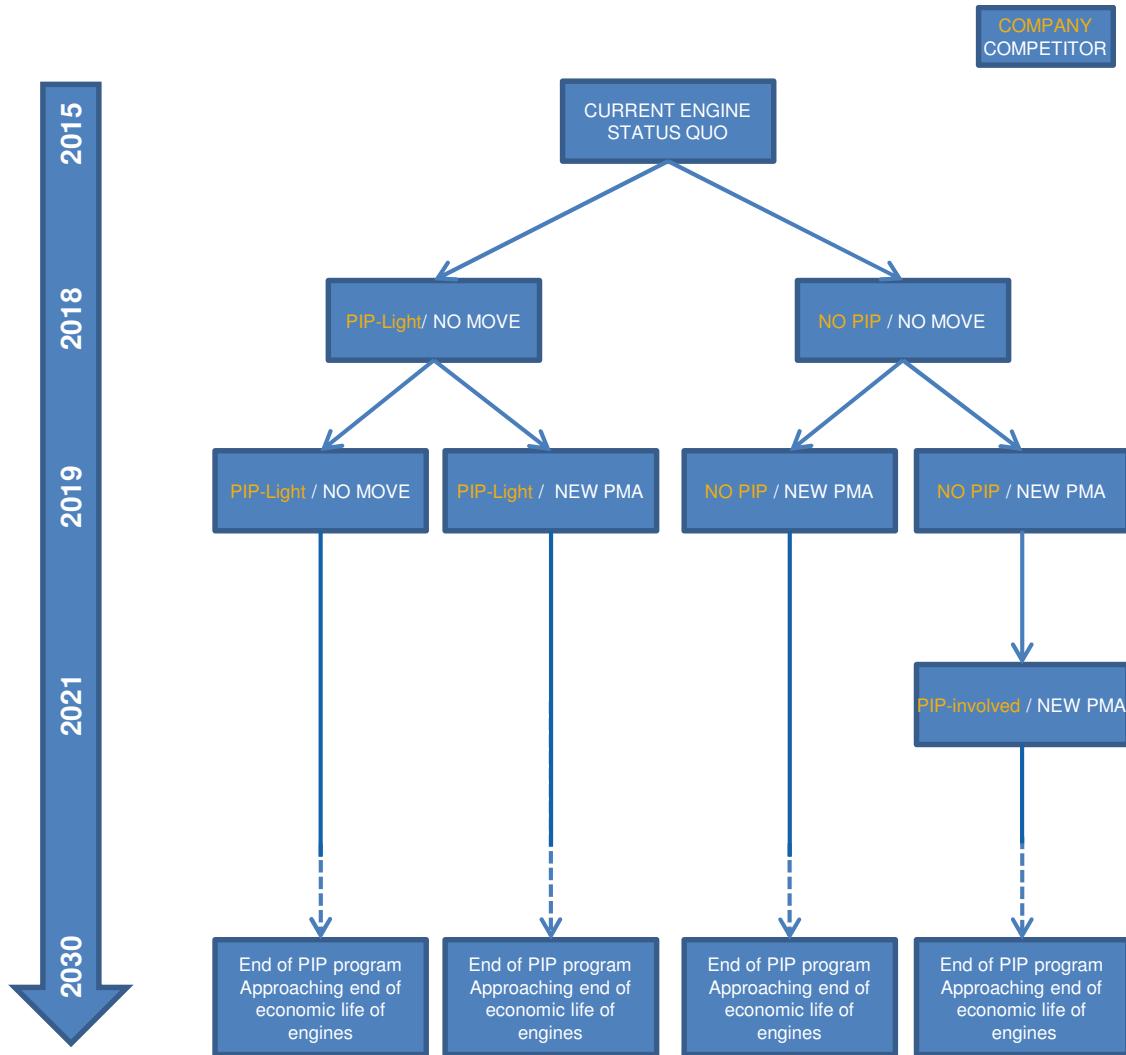


Figure 80: PIP development strategies in a competitive environment

8.2 Aircraft and engine operating economics

The performance improvement package is a set of technologies that aims at improving the operating economics of the aircraft. Ahead of the design, certification, and production, the manufacturer builds a business case by investigating the market reaction and by quantifying the benefits of retrofitting existing aircraft with the new technologies. Before delving further into the economic analysis of the package, a review of the operating economics of aircraft and turbofan engines is performed. It is a first step based

on the paper by Justin and Mavris [213] that helps understand what the technology package physically does to the engine, and how the physical changes map to economic improvements and increased airline profitability. Indeed, aircraft and engine economic evaluations are at the core of the proposed proof-of-concept study. Their goal is to help understand the behavior of airlines when faced with choices concerning the fleet renewal and fleet upgrading processes.

To do so, this research temporarily takes the point of view of airlines and investigates how airlines assess the economic performance of aircraft and engines, and how they eventually choose between the offerings of various manufacturers. To differentiate between the different aircraft available, two overall evaluation criteria are chosen to summarize the economic performance of the aircraft. These are the total airplane-related operating cost (TAROC) and the total airplane-related operating revenues (TAROR). The airplane-related operating costs and revenues are metrics that focus only on those costs and revenues that are incurred because of the operation of the aircraft. Consequently, they are a direct translation in monetary units of the operating performance of the plane and its suitability for the airlines' network.

The TAROR metric reflects the ability of the aircraft to generate revenues and consists of the revenues generated by the cabin, the cargo-holds, as well as the ancillary fees such as baggage fees. The TAROC metric reflects the costs incurred by airlines when operating the aircraft and includes both the indirect and direct operating costs. These are made of the acquisition costs, the financing costs, the insurance costs, the spare acquisition costs, the maintenance costs, the labor costs, the fuel costs, and various fees, taxes and ground handling charges. The choice of these two metrics is motivated by their

industry-wide acceptance [214]. Combined together and discounted properly over the entire operating life of the asset, they yield an estimate of the value of the aircraft. Indeed, in a similar application, Thokala et al. [215] indicate that the entire asset life-cycle needs to be accounted for when trade-off studies are performed for design selection. In this study, an airframe and engine valuation methodology is developed and it allows analysts to quickly evaluate the intrinsic value of an aircraft – or fleet of aircraft – by estimating both the TAROR and TAROC experienced by airlines.

The proposed aircraft and engine evaluation methodology called *i-CARE* for Integrated Cost And Revenue Estimation is articulated around four steps as highlighted in Figure 81. The first step is a network analysis that uses the schedule of flights of an airline to perform flight performance estimations. The second step is a TAROR evaluation based on the payload computations assessed during the first step. The third step is a TAROC evaluation that also uses the outputs from the network analysis. In the final step, results from the previous analyses are gathered to estimate present and expected future cash flows that are both aircraft and airline-specific. For the purpose of this thesis, more emphasis is put on the estimation of the TAROC since the performance improvement package will mostly impact costs. On the other hand, no significant changes are expected to the TAROR metric as the performance improvement package is assumed to have a negligible effect on the payload-carrying capability of the aircraft.

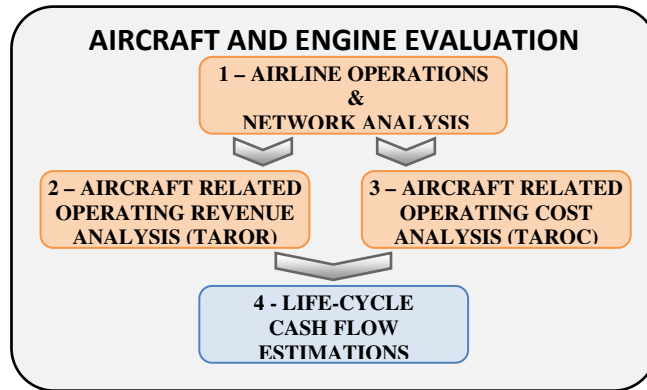


Figure 81: Aircraft evaluation methodology

8.2.1 Enabling the aircraft and engine economic evaluation

Several software are available in the industry which enable economic analyses of aircraft and engine in addition to the myriad of in-house calculators developed by original equipment manufacturers, airlines, research agencies, and other third parties. Popular names includes the *Jeppesen* Airline Optimization Suite, the *Pace* Mission software, the NASA FLOPS and ALCCA software, as well as the PIANO-X software. These analysis tools can be categorized into two categories: flight operations analysis tools which aim at optimizing flight operations for airlines and aircraft operators, and mission analysis tools geared towards aircraft designers and research entities performing conceptual and preliminary analyses. If it is easier to have access to the latter, they unfortunately do not go into the level of details required to account for the aging process of the engine and the resulting impact on the aircraft operating economics. As a result, a new aircraft and engine evaluation tool is developed with significant emphasis put on engine degradation processes, on engine maintenance analysis, and on engine performance degradation over time.

8.2.2 Airline operations and network analysis

The analysis is performed using a mission analysis software and the inputs for this analysis are shown in Figure 82. These inputs consist of an airline schedule, which is a list describing the departure and destination for each flight in the network, and of an aircraft performance file. Depending on the type of analysis performed, the airline's schedule might be a single generic route repeated several times a day, a typical schedule with a realistic mix of generic short and long segments, or a true airline schedule. In the latter case, the airline schedule can be retrieved directly from the airline website or from the OAG¹ database and is subsequently processed to extract a database of flights operated by the fleet of aircraft under review.

Two types of investigations can be performed using real airline schedules: fleet replacement simulations and fleet upgrade simulations. The first type of investigation simulates the analysis performed by airlines studying a possible replacement of their fleet with newer and more efficient aircraft models. The second type of investigation simulates upgrades to aircraft in the fleet currently operating the network. Under both assumptions, the network is operated by a fleet of new or upgraded aircraft replacing on a one-to-one basis the existing fleet such that no net growth in the network occurs.

¹ OAG: Official Airline Guide




Figure 82: Network Analysis

The output of the analysis is a table showing the leg-distance, block time, block fuel, and payload for each and every flight. Optional outputs include the optimal flight level, average expected wind, and derate factor for each flight in the network. These outputs are used to complete the remaining steps of the analysis and to estimate standard airline statistics such as yearly aircraft utilization and flight hour to flight cycle (FH:FC) ratios.

8.2.3 Revenue analysis

The revenue analysis consists in computing the operating revenues for each flight and then summing these revenues for each week of the period studied. The revenues come primarily from two sources: the passengers in the cabin and the cargo in the holds. Following recent trends in the airline industry, ancillary fees can also be accounted for [216] [217]. To estimate the number of passengers and the amount of cargo carried on each flight, an algorithm is used to fill up the plane with payload. For each flight in the network, the algorithm checks whether there is sufficient payload-carrying capability available to fill the cabin and the cargo holds given assumed or historical passenger and cargo load factors. If there is enough room to accommodate both passenger and cargo loads, the aircraft will be filled up. If not, cargo is removed first and if this is still not

sufficient, passengers are also removed. This ensures that the whole carrying capability of an aircraft design is used when performing comparative studies.

Once the passenger and cargo loads are set, revenue estimations are made using historical or forecasted yields for the airline under review. These estimations may be modulated on a monthly basis by a seasonal factor to account for cyclic variations in passenger and cargo demand. The revenues computed for each flight in the network are next aggregated to provide an estimate of the revenues generated by the operation of the aircraft.

8.2.4 Cost analysis

The cost analysis is carried out next by breaking down the costs into different areas, each representing one source of cost. These different areas are not totally independent and some relationships are defined in between them. For instance, turbofan engines deteriorate over time and fuel-burn is impacted by the aging status of the engine. These different costs are described in the following paragraphs with a greater emphasis on those dealing specifically with the aircraft and the engine where product differentiation may occur.

Acquisition and Financing Cost Analysis

There are many ways airlines can get access to aircraft. Some airlines may benefit from state-sponsored export credits while some others benefit from government credit guarantees. However, most of the financial schemes revolve around a handful of methods spanning from outright purchases of the asset to simpler operating leases. The most common acquisition methods are acquisition using cash, acquisition using debt, and

financial lease for which the asset is leased during a fraction of its operating life and then purchased by the operator. Another popular way of getting access to aircraft is through operating lease, although this is not an acquisition of the asset per se [218] [219].

Although the acquisition method does not really help with product differentiation, its peculiarities can affect the evaluation of other expenditures. For instance, operating lease contracts generally stipulate that the lessee must perform maintenance checks for both the aircraft and its engines before returning them to the lessor. Therefore, these additional checks must be accounted for during the evaluation of the maintenance costs.

Inputs for this analysis range from the lease contract structure or the loan structure to the expected residual value of the aircraft at the end of the study to account for positive cash flows when operators resell or scrap the aircraft for parts [220].

Aircraft Maintenance Cost Analysis

This analysis deals with the computation of the maintenance costs for both the airframe and some heavy components such as the auxiliary power unit (APU), the tires, the wheels, the brakes, the landing gear, and the thrust reversers. Both the airframe and the heavy components have their own specific maintenance programs and these usually consist of a list of parts for which inspection is due and a list of items for which replacement is due at specific calendar times, flight hours, or flight cycle intervals. These programs usually leave enough room for airlines to slightly adapt and optimize the schedule of maintenance for their own specific operations.

The airframe maintenance programs prescribed by either OEMs or independent maintenance, repair, and overhaul facilities (MROs) and reported in *Aircraft Commerce* [221] are used alongside the aircraft utilization statistics to derive the most likely

calendar of line maintenance, base checks, and heavy structural checks. Using the work scope associated with each of these events, the airframe maintenance costs can be estimated. The maintenance burden for heavy components is estimated in a similar way by first deriving a calendar of likely maintenance events and then estimating the costs of each of these maintenance events.

To simulate what is happening during the day-to-day operations of an airline, an airframe maintenance optimizer is implemented as depicted in Figure 83. Its purpose is to design a maintenance schedule which optimizes aircraft availability [222] by grouping maintenance events that are expected to be performed within a short time-period.

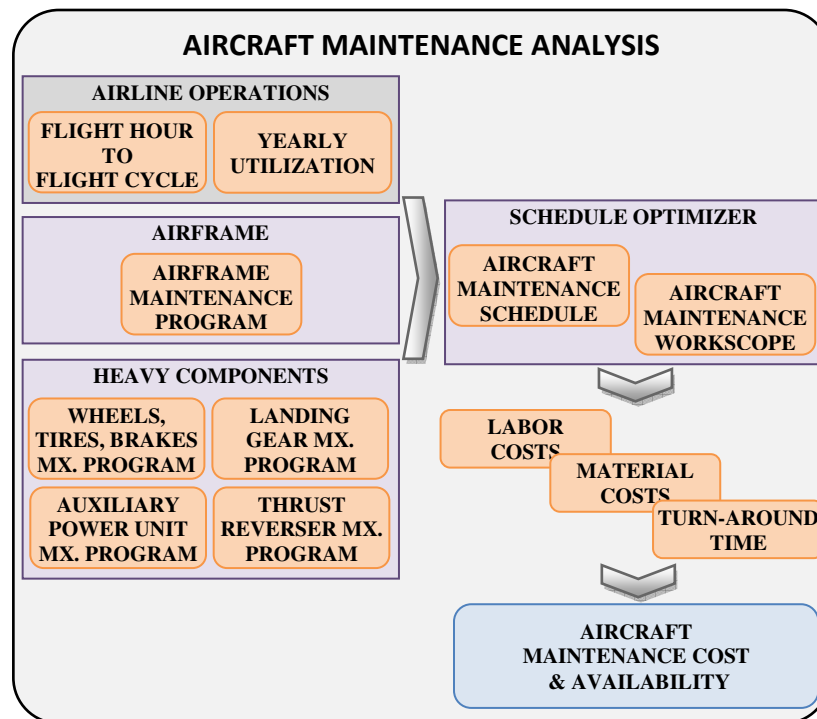


Figure 83: Elements of airframe and heavy component maintenance

The detailed algorithm used to forecast airframe and heavy components maintenance is given in Figure 84.

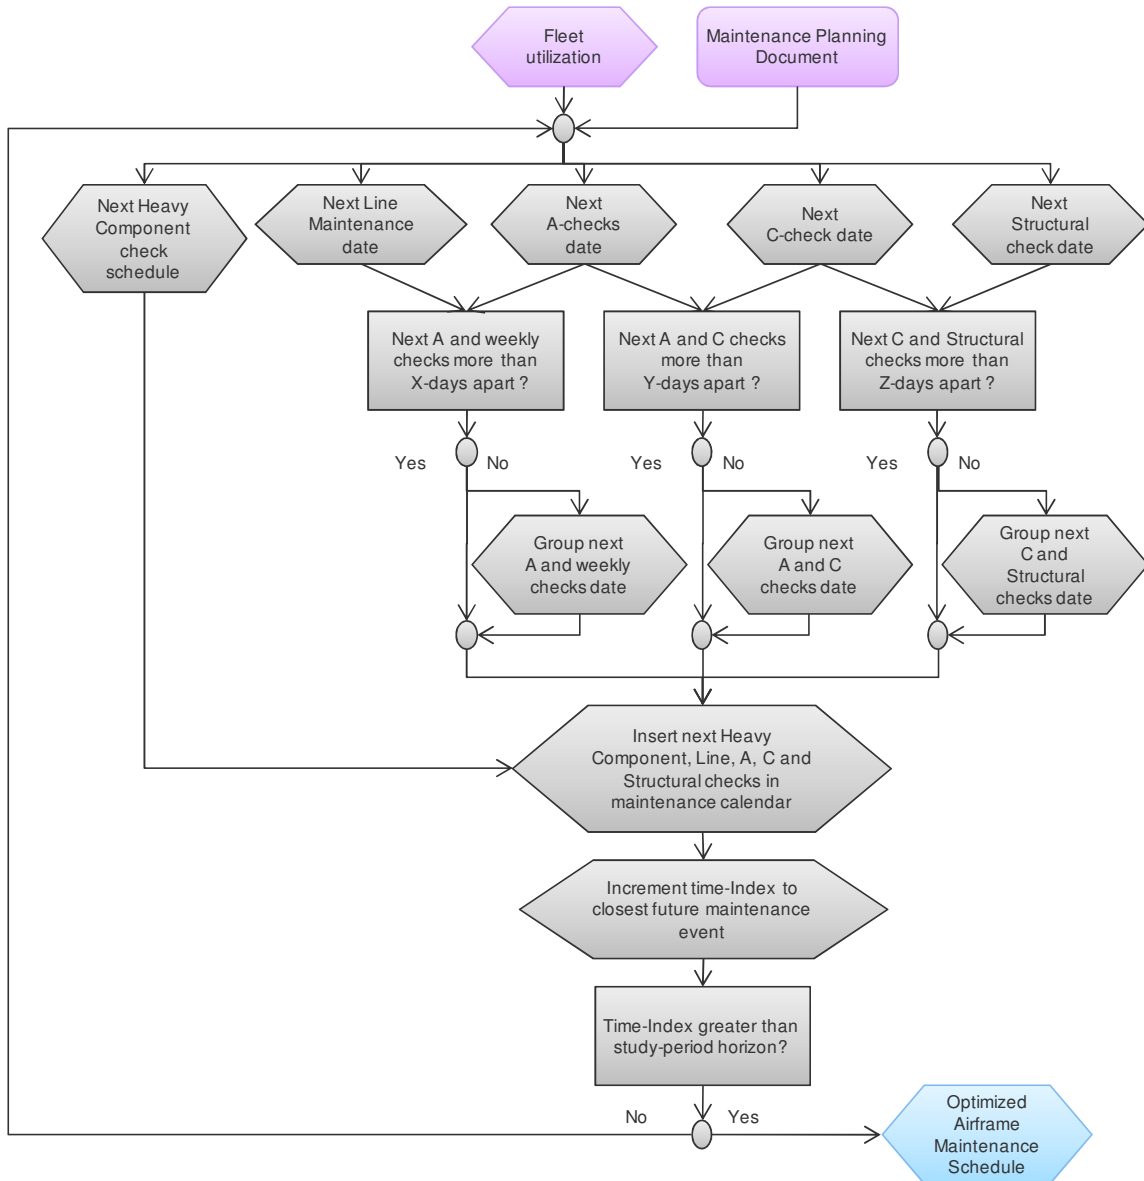


Figure 84: Forecasting airframe and heavy components maintenance events

Engine Maintenance Cost Analysis


The purpose of the engine maintenance analysis is twofold: first to estimate the expenses to keep the engine airworthy, and then to estimate the deterioration status of the engine (wear stemming from the operations and the surrounding environment) and the resulting effect on fuel-burn. Engine maintenance is usually divided into three main tasks: the first concerns the replacement of life-limited parts, the second concerns the

monitoring and restoration of the exhaust gas temperature margin, and the third concerns unscheduled engine removals due to aging, faulty engine hardware, excessive oil consumption, excessive vibrations, or external factors such as foreign object induced damage.

Life-limited parts (LLP) are parts of the engine for which failure cannot be contained. To ensure these parts do not fail during operations, they are completely removed from the engine and replaced with new ones at specific time intervals expressed in flight cycles. There are different sets of life-limited parts in an engine depending on the level of stress under which they operate: LLP for the fan, LLP for the low pressure compressor and turbine, and LLP for the high pressure compressor and turbine. Information regarding the cost and the certified lives of life-limited parts of typical narrowbody aircraft are published in *Aircraft Commerce* [222] [223] [224] [207] [208].

The exhaust gas temperature (EGT) is a temperature measured in the engine exhaust which indicates how efficient the engine is at producing its design thrust [225]. As the engine ages, it becomes less efficient, must burn more fuel, and thus runs hotter to provide the same amount of thrust as a new engine. There is however an upper limit – a *physical limit* – as to how hot an engine can operate since high temperatures will adversely affect the engine integrity. A new or recently overhauled engine can produce its rated thrust at an EGT well below a design reference, called the EGT red line, thus providing a large EGT margin. As the turbofan ages, the EGT margin decreases until the engine cannot produce its rated thrust without exceeding this reference red line. At this point, it becomes necessary to remove the engine and send it for a performance restoration overhaul which will restore some of the original EGT margin. The concepts of

EGT redline, initial EGT margin, and reduced EGT margin due to aging are described in Figure 85. The EGT margin erosion over time is a complex process and some research is carried out in the literature to model it. Using EGT margin erosion data-points for typical narrowbody aircraft reported by *Aircraft Commerce* [208] [227], Justin et al. [44] derive the EGT margin erosion power-law model displayed in Eq. 55: it represents the EGT margin lost as a function of the number of flight cycles. The expense incurred during these performance restoration shop visits includes the cost to replace and repair parts as well as the labor expense related to the removal, opening, and inspection of the engine.



$$EGT_{LOST} = a \cdot \left(\frac{CYCLE}{NBR} \right)^b \quad \text{Eq. 55}$$

$$a = 0.104 ; b = 0.659$$

Figure 85: Effect of engine wear on the EGT margin

Besides these expensive shop visits, airlines perform cheaper engine-wash to remove some of the contamination that is deposited inside the engine over the course of normal operations. Indeed, modern turbofan engines are very efficient pieces of machinery, but this comes at the cost of increased sensitivity to disturbances. Over time, the accumulation of deposits (of anti-icing and de-icing fluids to name a few [228]) inside the engine leads to a performance deterioration which can be mitigated by regular engine wash. According to Ackert [226], engine washing is an on-wing and ground-based

process that pumps water and cleansing additives into the engine intake while the engine is operating. The process fully penetrates the compressor and turbine to clean the airfoil surfaces. By removing deposits, regular engine wash increases the compressor efficiency, which restores some of the EGT margin and results in longer on-wing times and reduced fuel-burn. Figure 86 highlights the slower erosion and slower fuel-burn degradation of a typical narrowbody turbofan having regular engine wash and periodic overhauls.

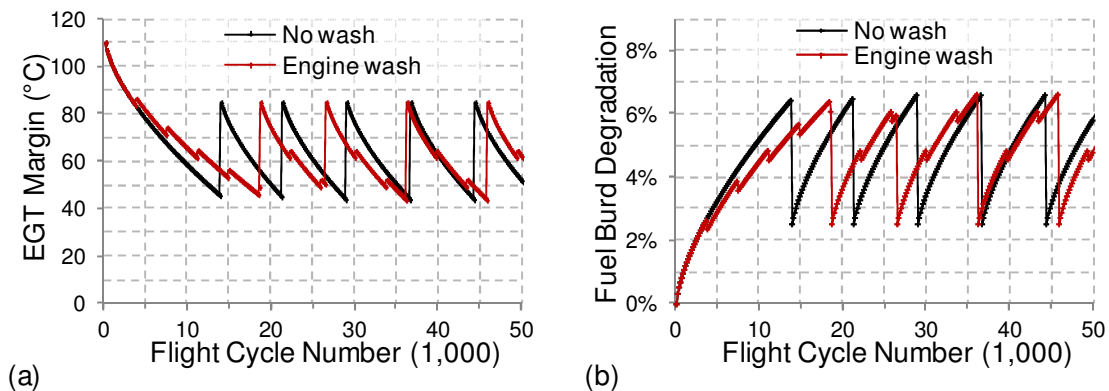


Figure 86: Effect of regular engine wash (a) on the EGT margin erosion and (b) on the fuel-burn degradation as a function of flight cycles

In addition to the shop visits, another source of maintenance expenses is related to unscheduled engine removals. According to Kleinhanas as reported in *Aircraft Commerce* [229], unscheduled removals are split between engine and non-engine related events and encompass a wide variety of maintenance events stretching from aging issues, faulty engine hardware leading to high oil consumption and excess vibrations, to extensive repairs following the ingestion of foreign objects.

Traditionally, this analysis is done using standard estimates for both the EGT margin erosion and the unscheduled maintenance models. This means that the effects of idiosyncratic operations are not accounted for in the reliability models. According to Henning quoted in *The Engine Yearbook* [230], the effects of flight operations need to be

accounted for using an operating severity factor which assesses the harshness of the operations and the resulting level of wear and tear on the engine above or below a standard. This engine-specific standard is defined as a triplet (flight hour to flight cycle ratio, engine derate, and outside air temperature) [231] which helps establish baselines for both the EGT margin erosion and the unscheduled maintenance models. Rupp [232] asserts that the severity modulates the maintenance cost by accounting for the engine derate used at take-off, climb, and cruise as well as for the flight length measured with the FH:FC ratio. Ackert [226] mentions that turbofans operated with larger derate factors use less thrust and therefore run cooler, which translates into lower deterioration rates and longer on-wing lives. Similarly, turbofans used for longer flights are spending proportionally less time at the take-off and climb power settings that are most demanding to the engine. This leads again to lower deterioration rates and longer on-wing lives as indicated in Figure 87.




Figure 87: EGT margin erosion for engines operating with low severity (red) and high severity (black) factors

A composite severity factor is estimated by multiplying an operational severity factor and an environmental factor. Both are estimated using standard tables and severity surfaces released either in the public domain [226] such as the one depicted in Figure 88, or in documentations provided by engine manufacturers. A methodology to compute the

severity factor is also proposed by Hanumanthan et al. [231]. The operational severity accounts for the harshness of the operations and uses the derate factor as well as the average length of flight (FH:FC) as inputs.




Figure 88: Operational severity factor

The environmental factor accounts for the atmospheric conditions and captures the additional wear and tear due to higher than standard outside air temperatures and more pronounced corrosive-erosive levels resulting from salty air and dust, sand, and gravel particles. A map representing the impact of the environment on the degradation of the engine is shown in Figure 89. The resulting composite severity is used next as an inverse multiplier for the characteristic life of Weibull distributions used to model the occurrence of unscheduled maintenance events [233] [234] and as a multiplier for the slope of the EGT erosion models.

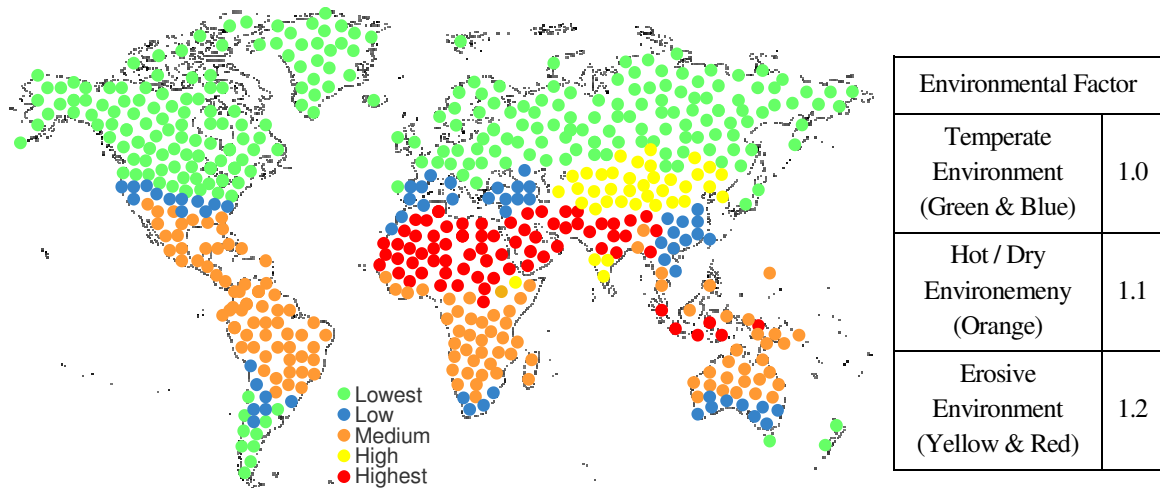


Figure 89: Environment harshness and impact on severity (adapted from [235])

The EGT margin erosion and the unscheduled maintenance models are however not valid over the entire engine lifetime. Instead, engine manufacturers usually distinguish between the first life, defined as the time up to the first shop visit, and the mature life, defined as the time after this first shop visit. Therefore, different models are retained for these two periods. For instance, the initial EGT margin for a brand new engine is higher than the EGT margin for an overhauled engine fresh out of maintenance because shop visits cannot restore an engine to like-new conditions [236].

Having defined the various causes of turbofan engine maintenance, the aim of engine operators is to schedule the maintenance events to maximize on-wing time, minimize maintenance costs, and minimize fuel-burn. A schedule optimizer is therefore used to maximize engine availability by grouping maintenance events while constraining the removal of life-limited parts that still have some operating potential. For this, the concept of life-limited part maximum *stub-life* is defined as the maximum number of cycles remaining for a life-limited part to be considered for replacement. Thus, at each maintenance event, the remaining potential of life-limited parts is compared to the

maximum stub-life to determine whether replacement is necessary. The stub-life concept is inimically linked to the concept of *build standard* which defines the expected on-wing time of the engine before its next shop visit and therefore the required EGT margin [235]. This helps define the workscope associated with each maintenance event so as to achieve a fine balance between less frequent maintenance events and more expensive maintenance.

As a result, the schedule optimization first assesses the workscope associated with each maintenance event and then estimates the associated cost. If the shop visit results from the need to replace the fan LLP, the low pressure LLP, or from other unscheduled maintenance reasons, a performance restoration may be performed concurrently depending on the remaining EGT margin. However, if the core LLP must be replaced, then the engine needs to be taken apart and a performance restoration is always performed. By tracking the EGT margin of each engine over time, the optimizer selects whether a comprehensive but expensive overhaul is required or whether a cheaper limited-scope performance restoration is sufficient. This leads to the maintenance model depicted in Figure 90 which properly accounts for both the operating and environmental aspects.

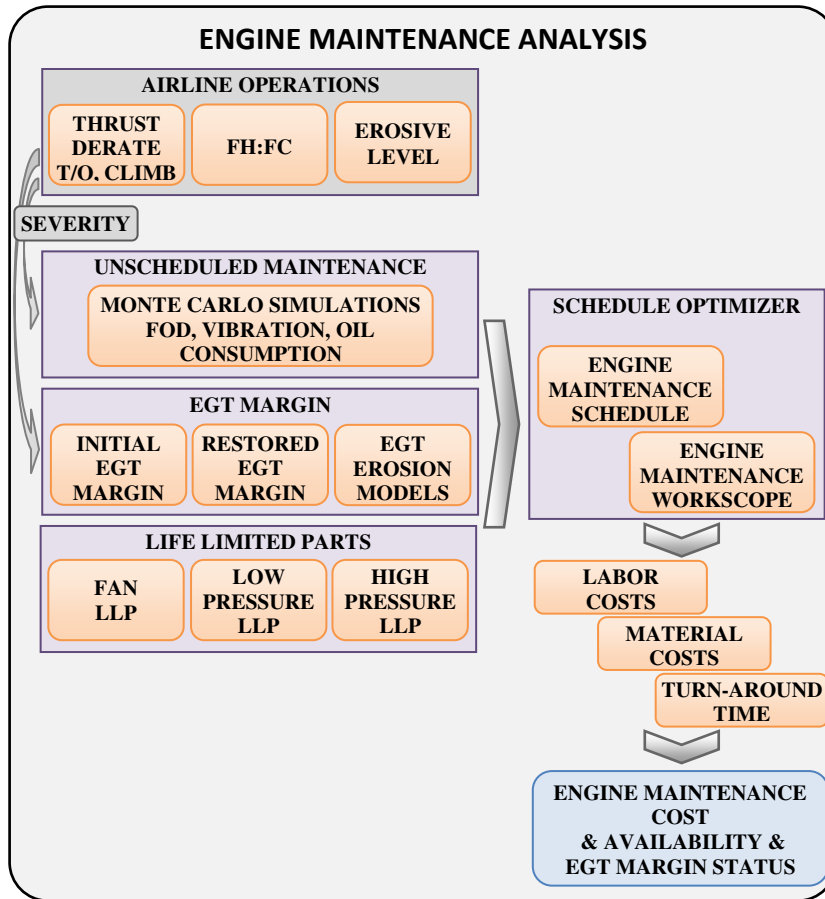


Figure 90: Engine maintenance analysis

The detailed algorithm used to estimate the engine maintenance schedule and engine maintenance expenditures is described in Figure 91.

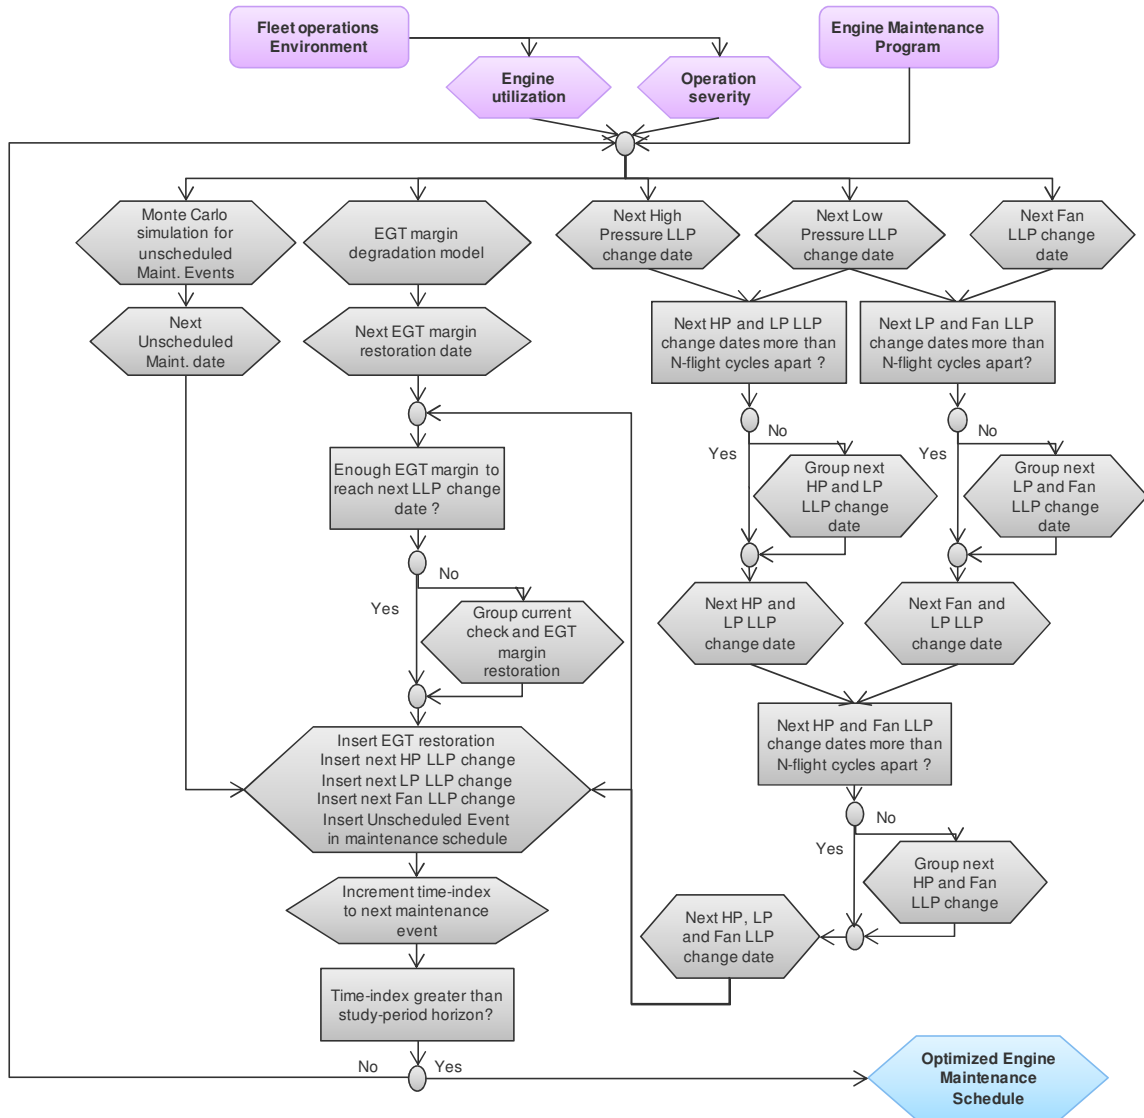


Figure 91: Forecasting engine maintenance events

Spare Cost Analysis

The acquisition of spares is also part of the overall life-cycle cost of an aircraft. The spares are subdivided into three categories: spare parts for the airframe, spare parts for the engines, and spare engines for the fleet of aircraft. In line with standard practices of the industry reported by IATA [237], it is assumed that the spare parts for the airframe and engine represent ten percent of the overall value of the aircraft to achieve smooth operations. The spare part and spare engine model is illustrated in Figure 92.




Figure 92: Spare part analysis

The number of spare engines is more complex to assess. First, a requirement regarding the engine dispatch availability is set. The dispatch availability is the probability that a spare engine is available whenever an engine experiences a failure and needs to be removed from the fleet. Next, the mature Shop Visit Rate (SVR: number of engine-caused shop visits per one thousand engine flight-hours) and the Turn-Around Time (TAT: time during which the engine is unavailable due to maintenance) are used to estimate the number of spare engines required. To count the number of engine removals (and therefore the need for spare engines) during a given time period, a Poisson counting process with intensity set to the mature SVR is used. This number of spare engines is estimated by iteratively computing the probability that one engine, then two, then three, and so forth go to maintenance simultaneously. This probability is compared next to the target dispatch availability to assess the number of spare engines required.

Mathematically, when the mature SVR and the TAT are expressed with respect to a common factor (one thousand flight-hours), this is equivalent to solving for the lowest number N_SPARE such that the inequality in Eq. 56 holds. Both the mature SVR and the TAT are outputs from the engine maintenance module with the mature SVR computed using statistics of simulated maintenance events after the first shop visit and the TAT estimated from the workscope associated with these events.

$$\sum_{n=0}^{N_{SPARE}} \frac{e^{-SVR \cdot TAT} \cdot (SVR \cdot TAT)^n}{n!} \geq \frac{\text{Target Dispatch Availability}}{\text{Dispatch Availability}} \quad \text{Eq. 56}$$

The detailed algorithm used to perform the computation of the number of spare engines required to achieve a target dispatch availability is given in Figure 93.

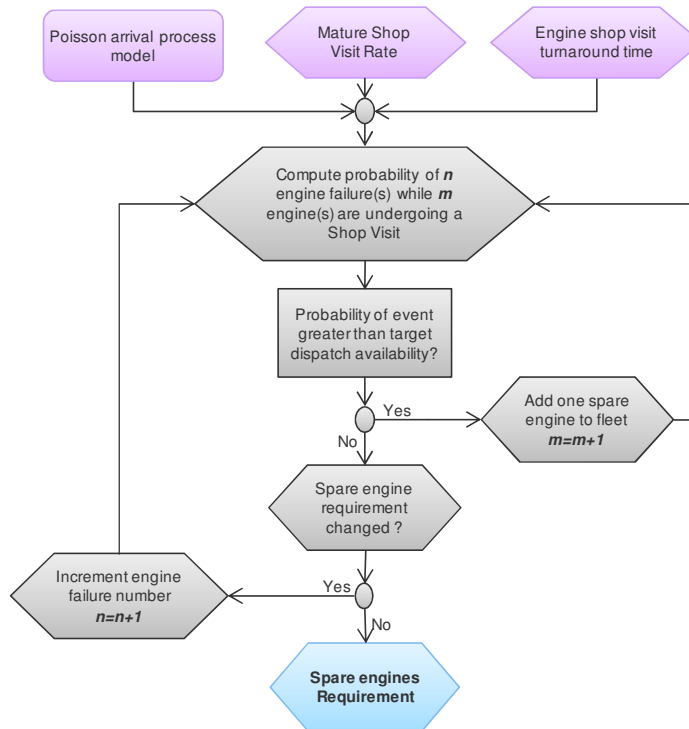


Figure 93: Spare engine requirement computation

Fuel Cost Analysis

The fuel expenses incurred by the aircraft operations consist of the cost related to the purchase of fuel for the powerplants and auxiliary power unit. This estimation is done in three steps as summarized in Figure 94.

The first step is the computation of the block fuel required to perform each and every flight in the network. This is done using a mission analysis software and this provides a first estimate of the amount of fuel needed. However, this does not account for the engine degradation status and the resulting impact on fuel-burn.

In the second step, the engine status is taken into account to refine the initial fuel-burn estimate. The engine degradation is tracked using the EGT margin erosion modeled in the engine maintenance analysis. Wiseman et al. [238] indicate that for large civil turbofans, a 1°C increase in EGT corresponds to as much as 0.1% increase in specific fuel consumption (SFC) in cruise. This estimation is corroborated by Yilmaz [239] for CFM56-7B turbofan engines mounted on typical narrowbody aircraft who estimate the mapping to be 0.81% increase in fuel flow for each 1°C increase in EGT at max continuous thrust and 0.87% increase in fuel flow for each 1°C increase in EGT at take-off thrust. Using this relationship, an estimate of the extra fuel consumed by an aging turbofan engine can be computed. To extend this computation to all engines within the fleet, the EGT margin erosion status of each engine is tracked and the corresponding excess fuel-burn is assessed. The excess fuel-burn of each engine is averaged to yield the instantaneous fleet-wide excess fuel-burn.

In the third step, the amount of jet-fuel required for the airline operations is converted into an expense by using a fuel price model. Because the price of jet-fuel is

volatile and hardly predictable over long periods of time, analysts can investigate different fuel price scenarios or use different stochastic models to represent the evolution of the spot price of jet-fuel over time.

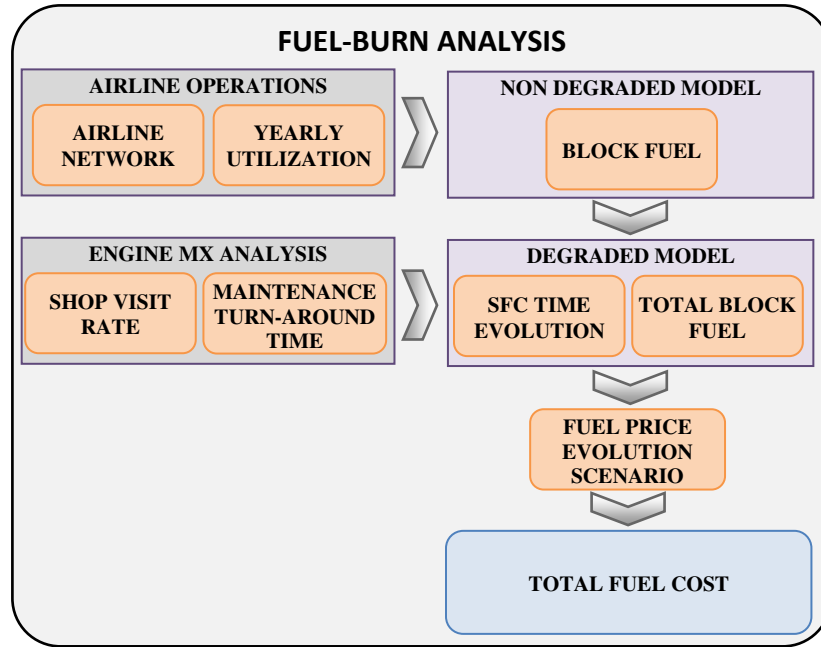


Figure 94: Fuel-burn analysis

Emission Cost Analysis

Aircraft release pollutants in the atmosphere while flying and these may be subject to taxation. Indeed, even though aviation only accounts for about 2.5% of global carbon emissions, the carbon footprint of aviation has increased by 98% between 1990 and 2006 as reported by the International Emissions Trading Association (IETA) [240]. IETA further indicates that if the carbon emissions growth forecast of 667% from 2006 to 2050 materializes, then the aviation sector may become a significant source of greenhouse gas emissions unless mitigation policies are undertaken. Since it is unlikely that the airline industry is able to improve fuel efficiency at the same rate as air traffic expands, the European Commission decided to include aviation into the emission trading

scheme on the basis that it is the most cost-efficient and environmentally-effective option for controlling aviation industry emissions [241].

Tradable emission permits are one way such taxation may be implemented as explained by Pearce and Pearce [242]. In its June-July 2010 issue, *Aircraft Commerce* [243] reports that airlines affected by the European Union's Emission Trading Scheme (ETS) must start collecting fuel consumption data to prepare for the trading of emission allowances. Although the taxation of international flight emissions has recently been postponed, emission expenses are still relevant for flights within the European Economic Area, and may be useful for scenario investigations used in long-term fleet planning exercises. The emission analysis is designed to accommodate carbon dioxide and nitrous oxides taxation schemes. The fuel consumption estimated during the fuel-burn analysis is used to estimate the amounts of carbon dioxide and nitrous oxides released in the atmosphere. In turn, these emissions are used to estimate the number of allowances and taxes that airlines have to pay. For carbon allowances, airlines have to buy about 15% of their emissions through auctions of carbon emission allowances [244]. For other emissions, emission quantities as well as tax rates [245] are used. These allowances and tax rates need to be adjusted for flights in different geographical areas, where allowances and taxation may be different or even non-existent. The emission analysis model is illustrated in Figure 95.

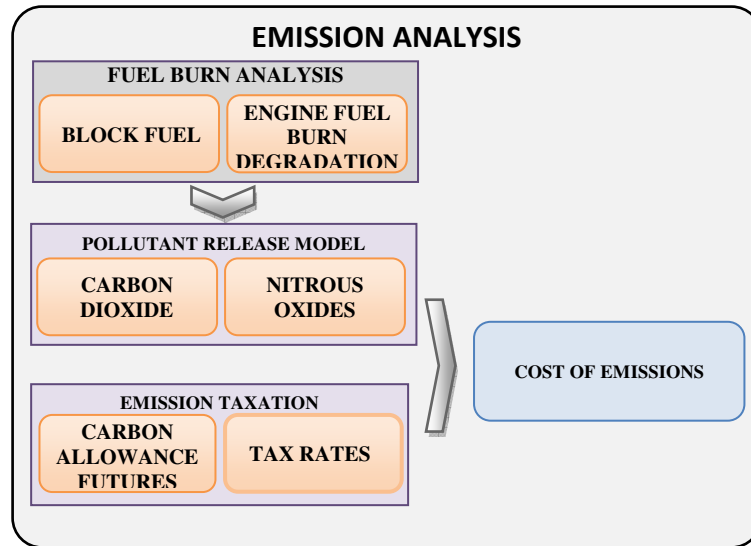


Figure 95: Emission taxation analysis

Labor Cost Analysis

The labor costs entering into the total airplane-related direct operating cost computation consist of the flight crew and the flight attendant costs. In order to perform this computation, the block time for each flight in the network as well as the capacity of the aircraft operating each flight are used. These two metrics are sufficient to estimate the required number of flight crews and relief crews as well as the number of flight attendants as per the Federal Air Regulations [246] [247]. With these two estimates, the overall numbers of man-hours for both categories are estimated using flight crew and cabin crew hourly rates escalated by the allowances, pension contributions, transportation, accommodation, and training as reported by *Aircraft Commerce* [248]. Extra staffing to provide a better cabin experience can also be taken into account by increasing the number of flight attendants on duty above the minimum number required.

Tax, Fee, and Charge Analysis

The tax, fee, and charge analysis is quite complex because of the diversity of tax schemes and fees used worldwide. The charges captured in this analysis include the navigation fees, the landing fees, and the handling charges such as gate fees, parking fees, lighting fees as well as the noise tax for environmentally-sensitive airports. For each flight in the network of the airline under review, the various charges collected by airports and air traffic control service providers are computed using the aircraft physical characteristics such as maximum take-off weight, approach and take-off noise levels, the network schedule (parking time, turn-around-time) as well as a customized database of airport and enroute charges built on published information (EUROCONTROL [249], IATA airport database [250]).

Insurance Cost Analysis

The insurance analysis deals with the computation of the expenses incurred to insure the aircraft hull. As a result, the inputs for this analysis are the fleet composition at any given point in time, as well as the market values for each aircraft in the fleet. If not publicly available, the market-driven aircraft values can be computed using the same residual value regressions employed during the estimation of the ownership costs (and resale price) using data from Kelly [251]. Next, typical hull insurance premiums for the airline industry [252] are used to assess the monthly insurance premiums.

8.2.5 Calibration of the engine maintenance modules in i-CARE

Calibration of life-limited part replacement

Replacement of life-limited parts is one of the main causes of maintenance for turbofan engines. The literature is reviewed to determine the different sets of life-limited parts as well as their certified lives. Typical short- to medium-haul narrowbody aircraft have three sets of life-limited parts [208]: one set for the fan and booster module, one set for the high pressure compressor and turbine, and one set for the low pressure turbine. The corresponding lives expressed in flight cycles are given in Table 66.

Table 66: Life-limited part lives and costs for typical narrowbody aircraft (Aircraft Commerce [208])

| | Lives (Flight Cycles) | Replacement Costs (2014-US\$) |
|--|--------------------------|----------------------------------|
| <i>FAN LLP:</i> Fan disk Booster spool Fan Shaft | 30,000 | 495,000 |
| <i>LPC & LPT LLP:</i> Stage 1 disk Stage 2 disk Stage 3 disk Stage 4 disk Shaft Conical support | 25,000 | 610,000 |
| <i>Core LLP:</i> Forward shaft Stage 1-2 spool Stage 3 disk Stage 4-9 spool Compressor CDP seal Front shaft Rear air seal High pressure turbine disk Rear shaft | 20,000 | 1,050,000 |

Calibration of EGT margin erosion process

One task for the construction of the engine maintenance cost model is to estimate the exhaust gas temperature margin erosion over time to determine the point in time where no EGT margin is left and the engine must be removed for an overhaul. Once

again, the specialized literature is used in order to retrieve EGT margin erosion data to fit a model. Unfortunately, rich and consistent sets of data are not plentiful in the public domain. Two sets of data have nevertheless been reported in *Aircraft Commerce*: Singer [253] reports EGT degradation rates for several narrowbody aircraft turbofans with lower thrust (23,000 lbs) and mean flight times of 1.2 hours, while Jesus [208] reports EGT degradation rates for mean flight times of 1.8 hours for lower thrust engines (23,000 lbs), and Karhumaki [208] reports EGT degradation for higher thrust engines (32,000 lbs). Two power law regressions are performed and yield the graphs of Figure 96 where the left exhibit is for shorter flights with lower thrust engines, while the right exhibit is for longer flights with higher thrust engines.




Figure 96: EGT margin erosion is dependent on the thrust rating and severity of operations

Besides the EGT margin erosion curves, it is also of interest to estimate the probabilistic distribution which yields the time distribution between performance restoration overhauls. Several probabilistic models could be used but Weibull distributions are popular to model the aging process of pieces of machinery, and have been used to model the aging process of turbofans. According to Nowlan and Heap [254], “Weibull distributions are candidates for representing items that have moderately high probability of failures at low ages and demonstrate monotonically increasing failure

probabilities thereafter.” Following Hanumanthan [231], a Weibull distribution is calibrated to match the EGT margin erosion curve. First, the expected time to erode the initial EGT margin, denoted $EGT_MARGIN_{t_0}$, is estimated using the erosion regression. This expected time, named mean time between removal and denoted MTBR, is expressed in flight cycles to be consistent with the rest of the engine maintenance program which uses mostly flight cycle units. When the outside air temperature is close to or hotter than the engine corner point, some operators want to have some extra EGT margin padding to ensure that the full thrust can be used without exceeding the red-line. Therefore, the entire EGT margin is seldom used and a minimum EGT margin, denoted EGT_MARGIN_{MIN} , is retained during operations. This means that only $EGT_MARGIN_{t_0} - EGT_MARGIN_{MIN}$ will be consumed during operations before a performance restoration overhaul is performed. Finally, if the engine is in its mature life, which means that an overhaul has already been performed, only a fraction of the original EGT margin has been restored. This fraction is denoted ξ and the restored margin is defined by $\xi \cdot EGT_MARGIN_{t_0}$. Using the power law previously regressed, this leads to the definition of the MTBF expressed in Eq. 57.

$$\xi \cdot EGT_MARGIN_{t_0} - EGT_MARGIN_{MIN} = a * MTBR^b \quad \text{Eq. 57}$$

With the estimate of the MTBR, the shape parameter denoted β as well as the scale parameter denoted η of the Weibull distribution are estimated next. Following Hanumanthan [231], Nowlan [255], Yu et al. [256], and Pascovici et al. [257], the shape parameter of the Weibull distribution is set to 5. The scale parameter is then used to match the probability distribution expected value, which is the expected failure time or mean time between failures, according to the formula in Eq. 58.

$$\eta = \frac{MTBR}{\Gamma\left(1 + \frac{1}{\beta}\right)}$$

Eq. 58

This leads to the parameterization and set of graphs displayed in Figure 97. These represent the probabilistic failure times (not a failure per se but rather the number of cycles accumulated when no EGT margin remains) for a lower thrust engine and a higher thrust engine for both the initial engine life and the mature engine life. As could be expected, the failure times happen sooner for higher thrust engines because they have less EGT margin to start with. They also occur much sooner during mature lives because the previous shop visit restored only about 75% of the EGT margin of a mint engine.

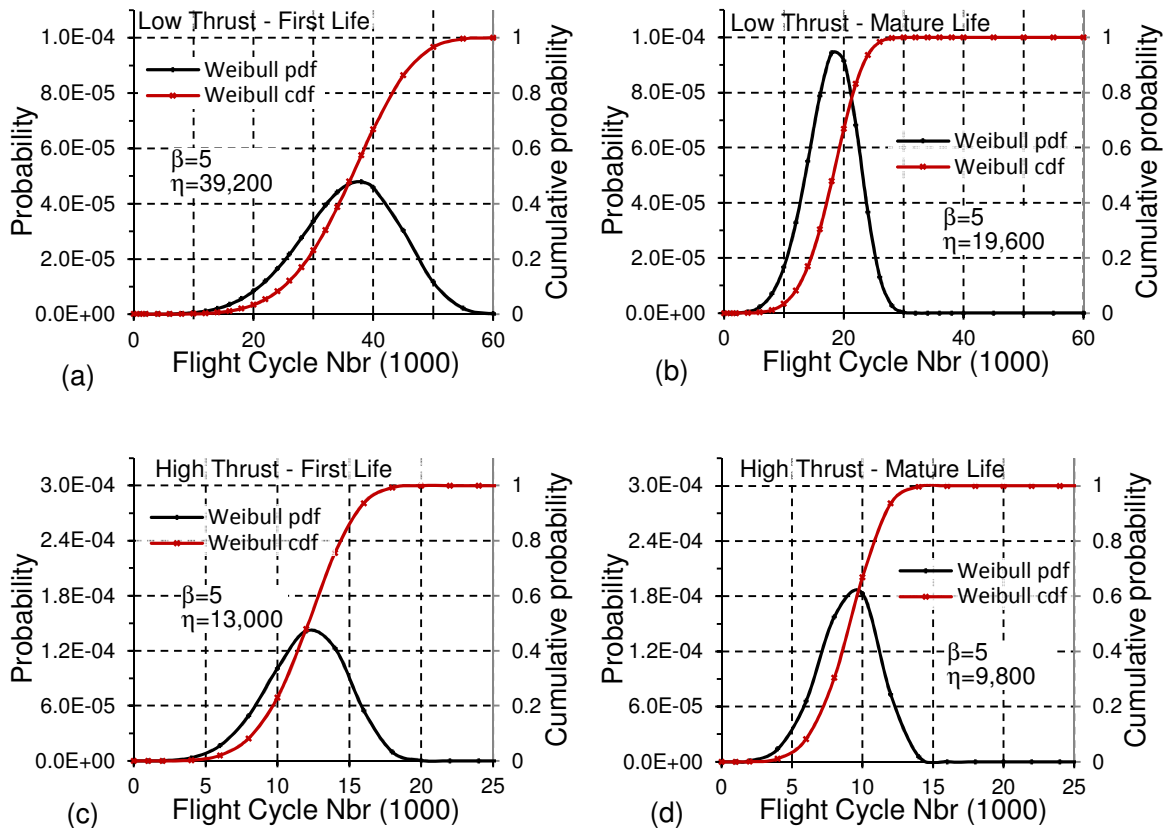


Figure 97: EGT margin erosion using Weibull probability distributions for the first and mature lives of low thrust engines (exhibits (a) and (b)) and high thrust engines (exhibits (c) and (d))

Calibration of aging process

Besides the erosion of the exhaust gas temperature margin, wear and tear on the engine affect other components which may result in different failure modes. Modeling the failures of each and every component in the engine, from regulators to pumps, is beyond the scope of this research. A new generic failure mode, denoted aging, is introduced. Some of these failures may be engine-related such as oil leaks, bearing failures and faulty hardware, while some others are non-engine related such as foreign object damage and bird-strikes. In *Aircraft Commerce* [229], Kleinhans indicates that “*non-engine related events [...] occur at a rate of about 0.005 per 1,000 engine flight hour while engine-related events occur at a rate of 0.028 per 1,000 engine flight hour*”. In another article, *Aircraft Commerce* [209] reports that “*all unscheduled removals occur at an average of once every 30,000 engine flight hour*” which is equivalent to a rate of 0.033 per 1,000 engine flight hours. Following Hanumanthan [231], “*the ageing curve as per the Maintenance Repair and Overhaul practice is a Weibull curve with a slope of 1.5, spreading to the entire life in service*”. This is further corroborated by Pascovici et al. [257] who indicate that the shape parameter of Weibull distributions representing the aging process of a turbofan engine is 1.5. With Eq. 58, the scale parameter is determined using the shape parameter by matching the expected time between failures to the estimates reported in the literature. This leads to the parameterization and the graph displayed in Figure 98.




Figure 98: Aging process requiring unscheduled engine removals using Weibull probability distribution

Calibration of engine removal costs

Lastly, the engine removal costs are estimated. Removal costs vary widely depending on the workscope to be performed and whether the engine needs to be opened (i.e. if major mating engine flanges must be separated). The turbofan engine Workscope Planning Guide usually suggests the amount of work to be performed [226] and three levels of workscope are suggested for engine shop visits: minimum level, performance level, and full overhaul level. Accordingly, engine removal costs are given in Table 67.

Table 67: Engine removal workscope and costs

| | Shop Visit | | | |
|---|---|--|--|---|
| | Minimum Level | Performance Level | Full Overhaul Level | Unscheduled Removal |
| Description | Limited time since last overhaul. Mostly external inspection and minor repairs. | Airfoils, vanes, seals, and shrouds are inspected. EGT margin needs partial restoration to meet min-build. | Degraded hardware condition forces full disassembly of engine. EGT margin is low and restoration required to meet min-build. | Non EGT margin related engine removals. |
| Criteria in Engine Maintenance Analysis | Fan LLP LPC LLP LPT LLP | EGT Weibull Sampling | EGT Weibull Sampling | Aging Weibull Sampling |
| Threshold | LLP lives up $EGT_MARGIN \geq 0.5 EGT_MARGIN_{t_0}$ | $EGT_MARGIN \leq 0.5 EGT_MARGIN_{t_0}$ | $EGT_MARGIN \leq 0.5 EGT_MARGIN_{t_0}$ | Whenever event occurs |
| Cost | US\$ 300,000 | US\$ 1,150,000 | US\$ 1,400,000 | US\$ 1,600,000 |

8.2.6 Verification and validation of aircraft and engine analysis method

Implementation verification

Having reviewed the various implementation choices and the different algorithms used, the aircraft and engine evaluation methodology is now verified. Contrary to other parts of this research, the methodology verification is not in the Experimental Plan chapter of this dissertation because it is not part of the proposed methodology to analyze program developments facing market and competitive uncertainties, and therefore it does not contribute to the verification or validation of hypotheses set forth in this research. Nonetheless, validation of the proposed methodology requires a proof-of-concept application to a relevant industry problem, which in this case relies on a proper economic assessment of aircraft and engine combinations. Therefore, a full-fledged verification of the aircraft and engine evaluation methodology is necessary.

A test case is used for verification purposes. It is designed to verify the correctness of the approach undertaken and to check for abnormal patterns in the results by comparing those with existing published analyses. The case is set-up with a single narrowbody medium-range aircraft operating on a network with an average flight length of 1.5 hour. This setup is very similar to a set of analyses published in 2006 by *Aircraft Commerce* [221] [258] [259] [260] [261]. This published study used a short to medium range narrowbody A319 aircraft manufactured by *Airbus* and the study was carried out using specialized software as well as field data collected from airlines and maintenance facilities worldwide. Table 68 summarizes the input parameters for these two analyses. Both analyses are very similar which enables a direct comparison of their outcomes.

Table 68: Inputs for the verification test case

| OPERATING METRIC INPUTS | | |
|--|--------------------------|--------------------|
| 125 Passengers – Medium Range – Narrowbody Aircraft | | |
| | <i>Aircraft Commerce</i> | <i>i-CARE</i> |
| Aircraft | | |
| Aircraft Type | A319 | 125-pax Narrowbody |
| Engine Type | CFM56-5B6 | 23K Thrust |
| Acquisition Method | Operating Lease | Operating Lease |
| Network Operations | | |
| Business Environment | Europe | Europe |
| Airline Type | Domestic Carrier | Domestic Carrier |
| Network Type | Generic Route | 35 routes |
| Aircraft Utilization | | |
| Yearly Block Time (BH) | | 3,186 |
| Yearly Flight Hours (FH) | 2,800 | 2,765 |
| Yearly Flight Cycles (FC) | 1,830 | 1,826 |
| FH:FC | 1.51 | 1.51 |
| Average Mission Length ESAD (nm) | 627 | 604 |
| Average Take-Off Derate (%) | | 14 |
| Environment | | Hot and Dry |
| Severity Factor | | 1.18 |
| Analysis Input Parameters | | |
| Aircraft Lease Rent (% Value) | 6% per year | 6% per year |
| Maintenance Program | MPD28 | Similar to MPD28 |
| Target Engine Dispatch Availability (%) | | 97.5 |
| Shop Visit at End of Lease | | Yes |
| Typical Life-Limited Parts Stub-Life (%) | | 8 |
| Maintenance Minimum Build (FC) | | 5,000 |
| Fuel Price (2014US\$/Gal) | 2.00 | 2.00 |
| Insurance Premiums (% Value) | | 1.5 |
| Study Length (Years) | 25 | 25 |

Results from this verification test case are summarized in Table 69. It appears that the results from the proposed methodology closely match both the direct and indirect operating costs reported in the specialized literature. The main sources of difference are

related to the fuel costs, the airframe maintenance costs, and the taxes, and may be explained as follows:

- One reason to explain the 5.6% difference in the projected fuel costs is the fact that the proposed analysis takes into account the fuel-burn degradation over time while it is not clear whether the *Aircraft Commerce* analysis models this. If the analysis is run without the fuel degradation model, the fuel cost per flight cycle decreases to US\$ 2,363 which is much closer to the *Aircraft Commerce* estimate (-0.8% difference).
- The remaining difference can be attributed to the choice of alternate airports for diversion purposes. The choice of an alternate airport and the additional distance to fly there may affect the block fuel estimate as it takes fuel in order to carry fuel to fly to the alternate airport. While the distance to the alternate airport is not stipulated in the *Aircraft Commerce* analysis, it is an average of 419 nm in the proposed analysis.
- For the taxes, fees, and user charges difference, one explanation may be the sensitivity of those to the type of operations investigated. The *Aircraft Commerce* number is an average estimate for European operations whereas the proposed analysis uses real charges for a small “equivalent” network in Southern Europe where airport taxes are lower than the European average.

Finally, the two analyses differ by 0.8% for the overall TAROC figure. When the fuel-burn degradation over time is not accounted for, the two analyses are only 0.2% apart.

Table 69: Comparison of results for verification test case

| AIRLINE COST PER FLIGHT CYCLE – 2014 US\$ | | | |
|--|--------------------------|--------------------------|-------------------------|
| 125 Passengers – Medium Range – Narrowbody Aircraft | | | |
| | <i>Aircraft Commerce</i> | <i>i-CARE</i> | % Difference |
| Fuel Costs [258] (Fuel Costs No Deterioration [258]) | 2,334 | 2,465 (2,363) | 5.6% (1.2%) |
| Flight and Cabin Crew Costs [262] | 1,753 | 1,740 | -0.8% |
| Airframe Maintenance Costs [221] | 1,114 | 1,103 | -1.0% |
| Engine Maintenance Costs [221] | 584 | 580 | -0.8% |
| Total Maintenance Costs | 1,699 | 1,683 | -0.9% |
| Spare Costs | | 96 | |
| Aircraft Leasing Costs [259] | 1,614 | 1,566 | -3.0% |
| Total Operating Costs (No deterioration) | 7,400 | 7,550 (7,456) | 2.0% (0.8%) |
| Emission Costs [263] | 47 | 45 | -4.1% |
| Taxes, Fees, and User Charges [248] | 1,246 | 1,171 | -6.0% |
| Insurance Costs [262] | 141 | 143 | 1.4% |
| TAROC (TAROC No deterioration) | 8,835 | 8,909 (8,815) | 0.8% (-0.2%) |

Since most of the inputs used to construct the aircraft and engine evaluation methodology are sourced from data published by *Aircraft Commerce* regarding a specific mission of one and a half hour, the comparison in Table 69 provides a means to check the consistency between the outputs of the analyses from *Aircraft Commerce* and the outputs from the *i-CARE* implementation. In other words, what is being checked in this section is the proper functioning of the methodology implementation and this constitutes the verification step. Given the excellent correlation between the outputs of these two analyses, the verification process turns out to be successful.

Implementation validation: fuel-burn and maintenance costs

In the previous section, the verification of the *i-CARE* implementation is performed by checking the outputs of the analysis on a specific mission using inputs

corresponding to this mission. Yet, many different types of missions may be investigated over the course of this dissertation in order to capture the whole spectrum of aircraft operators and owners. As a consequence, a wide variety of missions are analyzed next and compared with corresponding results published in the literature. The main difference with the preceding exercise is that the published data serving as reference in these new analyses was not used during the implementation of the aircraft and engine evaluation tool. In some sense, the entire domain of application of the proposed methodology is now checked with off-design analyses.

Since the proof-of-concept application requires mainly fuel-burn estimates and maintenance cost estimates, special emphasis is put on the validation of the maintenance cost and the fuel-burn estimations. Several missions ranging from one hour to slightly over two and a half hour are analyzed using *i-CARE* and compared with new analyses published by *Aircraft Commerce* [223] [208] [236]. The results are displayed in Figure 99 where the maintenance costs are expressed per flight cycle (in black) and per flight hour (in blue) for the different missions.




Figure 99: Comparison of published (+) and computed (◊) maintenance costs per flight cycle and per flight hour for different missions

Results from the published literature are depicted using (+) signs while outputs from the *i-CARE* calculator are depicted with rhombus signs. The graph shows excellent correlations between these results and the validation is successful.

For the fuel-burn estimates, the different missions are defined with the Equivalent Still Air Distance (ESAD) metric since wind assumptions may otherwise introduce noise in the data. Validation is performed for missions ranging from an ESAD of 373 nm to an ESAD of 1266 nm as shown in Figure 100. Results are provided for the model with fuel-burn degradation depicted using blue rhombus signs as well as without the fuel-burn degradation depicted using red rhombus signs. Again, computed results match very well the published results by *Aircraft Commerce* [258] [264] [265] and the validation is successful.




Figure 100: Comparison between flight fuel-burn published (+) and computed (◇) for different missions

Implementation validation: airline subfleet network analysis

The aircraft and engine evaluation methodology is tested next on a bigger scale using the narrowbody sub-fleet of an airline operating in Southern Europe. The fleet consists of five narrowbody aircraft delivered over a three-year period that are operated

for twenty six years from a base in Malta. The aircraft are assumed to be financed using debt. The airframe maintenance program used is derived from the *Lufthansa Technik's* progressive maintenance program for *Airbus* narrowbody aircraft as described in *Aircraft Commerce* [221].

Table 70 describes the assumptions related to the airline and highlights some relevant airline statistics. For this analysis, the environment is assumed to be erosive due to the proximity of the Mediterranean Sea and the likelihood of sand and dust in the air. The mean derate for the different flights in the network is set to 10%. The average outside air temperature at the Malta airport is assumed to be 24°C or 75°F. According to Hanumanthan [231] and Ackert [226], these network statistics yield an operational severity of 0.9 and an average composite severity of 1.12 when compared to the reference short-haul mission of unit severity (reference mission defined as 1.4 hour flight with 10% derate, and 18°C outside air temperature).

Table 70: Network description and operating statistics

| OPERATING METRIC INPUTS | | | |
|--|---------|---------------------------|-------|
| 140 Passengers – Medium Range – Narrowbody Aircraft | | | |
| <i>Network Description</i> | | <i>Airline Statistics</i> | |
| Yearly Aircraft Block Hours | 3,790 | Fleet Size | 5 |
| Yearly Aircraft Flight Hours | 3,379 | Average Load Factor | 76.1% |
| Yearly Aircraft Flight Cycles | 1,764 | Average Yield (\$/pax/nm) | 0.129 |
| Flight Hours to Flight Cycles | 1.92 | Cargo Load Factor | 45% |
| Average Flight Length (nm) | 780 | Average Derate | 10% |
| Operating Environment | Erosive | Composite Severity | 1.33 |

The chart in Figure 101 describes the repartition of costs over 25 years of operations. As may be expected, fuel expenditures and labor expenditures constitute the

main sources of costs with asset depreciation and airframe and engine maintenance following next. The computation is sensitive to the price of jet-fuel and its evolution over time. For this analysis, the price of jet-fuel is set to US\$1.50 per gallon (the 2014 US\$ equivalent of the price of jet-fuel when the first aircraft was delivered to the airline in 2004). This price is then escalated by 3% every year. This repartition of expenditures is in-line with typical results from the industry.





Figure 101: Fleet-wide operating cost breakdown **Figure 102: Life-cycle cash inflows and outflows**

Next is Figure 102 which shows the monthly discounted cash flows over time as dark red bars and the cumulative cash flows as the light-colored surface. Peaks of cash outflows are initially related to the deposits made during aircraft deliveries, while later on, they are tied to maintenance events: these events are either related to the engines (full overhaul with low pressure and core LLP change around year 2015 and fan LLP change around year 2021) and to the airframe (C-4 check starting around year 2010, C-8 check starting around year 2016). The periodicity of cash flows exhibited in Figure 102 is caused by the seasonal nature of yields and load factors.

The chart in Figure 103 exhibits the evolution of the fleet-average EGT margin and the resulting fleet-average fuel-burn degradation over time. Indeed, the specific fuel consumption increases as the engine ages until a maintenance overhaul restores some of its original performance. This is primarily due to increasing blade tip clearances, airfoil erosion, and contamination [212]. The subtle jaggies along the fuel consumption degradation curve are due to the fuel-burn improvements following regular engine-wash. These jaggies are however smoothed-out a bit due to the fleet-wide averaging effect.




Figure 103: Fleet-wide EGT margin degradation (red, decreasing) and SFC degradation (blue, increasing)

Another figure of interest to assess the plausibility of the engine evaluation analysis is the shop visit rate (SVR). The shop visit rate estimates the frequency of engine removals by computing the number of shop visits per one thousand engine flight hours. Because the engine maintenance expenses vary significantly over the life of the engine, the shop visit rate is usually computed for the first life and for the mature lives. First, an engine life is the time interval between two engine shop visits. The first engine life is the time interval up to the very first engine removal and shop visit, and the mature lives are life thereafter. The chart in Figure 104 displays the evolution of the running shop visit rate (cumulative metric since entry into service, in red) as well as the mature shop visit

rate (straight red line). The first life engine shop visit rate is 0.018 per 1,000 engine flight hours or EFH (value of the running SVR right before the first jump), while the mature life engine shop visit rate is 0.052 per EFH. Both compare well with numbers published by Ackert [226] and the International Aviation Service Group [266] for similar engines.




Figure 104: Engine Shop Visit Rates and engine dispatch availability with two spare engines

In addition to the shop visit rates, Figure 104 provides information regarding the engine dispatch availability. The engine dispatch availability is related to the shop visit rate since it measures the probability that one spare engine is available whenever an engine is removed and in maintenance. Airlines set up a target engine dispatch availability as a corporate policy and determine the amount of spare engines needed to achieve this goal. In this analysis, a target of 97.5% which is typical for the industry is used and this results in the need for two spare engines. This is slightly over the recommended amount of spare engines set forth by engine manufacturers (about 8% to 9% as reported by *Flightglobal* [267]) and other industry stakeholders (about 10 to 15% according to *Willis Lease Engine Corporation* [268]), but the airline fleet under analysis

is small (five aircraft, ten mounted engines, and two spare engines) and the number of spare engines must be a whole number!

8.3 Developing a market behavior model for the PIP adoption

This section deals with the development of a market model to estimate the adoption of the performance improvement package by engine operators worldwide. The purpose of the market model is several-fold: the first objective is to divide the market into homogeneous market segments while trying to capture the wide variety of operators worldwide; the second objective is to quantify the operating benefits of retrofitting an engine with performance improvement packages taking the point of view of the engine operator; the final objective is to estimate how likely operators are to purchase the retrofit for subsequent installation on their fleet of turbofan engines.

8.3.1 Market segmentation

The first step is the market segmentation whose purpose is to create different customer profiles representing various types of airlines, each with its own set of requirements, preferences, and operations. Thus, a market segment is a homogeneous entity which represents one specific type of customer. This market segmentation paves the way for the estimation of the economic utility of the technology retrofit for each customer profile. As highlighted in Figure 105, there is nonetheless a wide variety of customer profiles since the airline industry encompasses operators as diverse as ultra-low-cost airlines, low cost airlines, domestic airlines, legacy airlines, and premium international airlines. For the purpose of the proof-of-concept application, it is assumed that this diversity of customers can be captured using two metrics: the average flight

length and the ease of access to capital. Indeed, amongst the myriad of parameters that can be used to describe airline operations, the average length of flight and the aircraft utilization are the two metrics that impact most the total aircraft-related operating costs, which is how the effect of technology infusions will be measured. Consequently, several market segments are constructed using representative airline networks with given flight-hour to flight-cycle ratios and corresponding yearly aircraft utilizations.




Figure 105: Similar aircraft, widely different flight operations

Besides, each of these market segments is further split into two sub-segments: one sub-segment represents airlines in good financial standing with ease of access to capital, while the other sub-segment represents airlines in financial hardship with limited access to capital. The degree of access to capital does not impact directly the aircraft and engine economics. Instead, it impacts the economic analysis of investments and capital expenditure decision such as the purchase of performance improvement packages. All in all, a representative airline network and an airline discount rate are used to define a particular market segment, the size of which is commensurate with the size of the narrowbody fleet of the airlines represented.

Using the Bureau of Transportation Statistics (BTS) data [269] publicly available, the distribution of flight length operated by short- to medium-haul narrowbody aircraft is generated and used to split the whole market into eleven different market segments. The yearly utilization corresponding to each of these segments is estimated by regressing data published by *Boeing* [270]. These are depicted in Figure 106.




Figure 106: Typical narrowbody utilization (a) and yearly utilization (source *Boeing* [270]) (b)

For the segmentation based on the ease of access to capital, a Weighted Average Cost of Capital (WACC) corresponding to established and financially sound airlines is used in conjunction with another weighted average cost of capital corresponding to financially strained airlines. The values are provided in Table 71 following a quick survey of the financial reports of a dozen airlines.

Table 71: Selection of the WACC for market segmentation


| Easy access to capital | | Limited access to capital | |
|------------------------|-------------|---------------------------|--------------|
| Air France | 7.1% | Avianca | 10.5% |
| Delta Airlines | 8.9% | Alitalia | 9.2% |
| Hawaiian Airlines | 6.5% | China Eastern | 10.8% |
| Jetblue | 6.8% | Iberia | 12.2% |
| Lufthansa | 6.2% | Jazeera Airlines | 12.5% |
| Ryanair | 6.1% | LATAM | 9.2% |
| Southwest | 8.7% | Thai Airways | 9.1% |
| Average | 7.2% | Average | 10.5% |

8.3.2 Simplified market behavior model

The second step is the modeling of the market behavior which aims at trying to forecast the market preference or the probability of purchase by each operator within the eleven different market segments. This is a very complex task that is well beyond the scope of the dissertation since publicly available data to calibrate these models is not available. Some methods to perform this prediction include utility-based methods such as the choice models proposed by Luce [271], Lesourne [272], and McFadden [273] [274], as well as the popular conjoint analysis [275].

Owing to the inability to get relevant data to calibrate these models, a simplified approach is undertaken which relies on the estimation of the economic surplus generated by the infusion of the technology retrofit. Indeed, the net present value of the gains obtained from operating the retrofitted engines (using the discount factor representing the type of market segment under review) is compared to the initial capital expenditure to yield the dollar surplus of the investment for operators. A logit model is used next to assess the customer purchasing intentions of the retrofit. Two alternatives are offered to the customer: either purchase the retrofit if it makes sense economically or keep operating the engine as it is. In the case of purchase or no purchase decision by the operator, the attribute retained is the difference in net present value between the two alternatives normalized by the purchasing price of the retrofit. The logit model is only “heuristically calibrated” for lack of other means. On the one hand, if the net present value of the retrofit is negative for the operator, there is little chance that the retrofit will be sold (unless the operator views the extra efficiency brought by the retrofit as a hedge against future uncertain energy prices). As a result, the β value must be large enough to

have small purchasing probability in these cases. On the other hand, if the net present value of the retrofit is positive, a significant number of operators will purchase the retrofit. For instance, John Leahy of Airbus mentioned that the “Airbus A320 New Engine Option will sell at a premium of US\$7-8 million , or one half the net present value of the 15% fuel savings the aircraft would deliver over today’s generation of Airbus A320 and Boeing 737s” [276]. This is particularly relevant for the current proof-of-concept study as the application is very similar. Given the commercial success of the A320 New Engine Option (most of the new orders for Airbus narrowbody aircraft have been for this new engine option, estimated to be 90% over the period 2010-2012) and a price tag of about US\$55 million for a new A320 aircraft, this indicates that a difference in net present value of US\$8 million normalized by the price of a new A320 aircraft brings a 90% choice probability. This translates into a β value of 15 as shown in Figure 107 and Eq. 59.



$$P(Alt_1) = \frac{1}{1 + \exp(-\beta(z_1 - z_2))} \quad \text{Eq. 59}$$

Figure 107: Probability of choice between two alternatives with attributes z_1 and z_2

8.3.3 Decreasing market size and value leakages

The third step is the estimation of the market size. Since the performance improvement package retrofit is applied to a turbofan engine currently powering a fleet of narrowbody aircraft, the historical deliveries of narrowbody aircraft over time may be

used as a first estimate of the entire market. It is also assumed that about half of this market is powered by turbofan engines that can be retrofitted with the performance improvement package.

The performance improvement package is offered as a retrofit of technologies to be infused into turbofan engines that are currently in operations. It is not offered for newly built engines that are more efficient because they already benefit from a different set of non-retrofitable technologies. In addition, the economic life of aircraft and engines is not infinite: as time goes by, aircraft reach the end of the structural life for which they have been certified and new competitive offerings become available to the market. For these reasons, airlines usually retire or sell their aircraft after operating them between twenty and thirty years. Retirement usually occurs right before heavy maintenance when significant investment is due to keep operating the aircraft. This often makes continued operations either not economically sustainable or not competitive. Owing to the finite life of aircraft and engines, the target market size for the performance improvement package is expected to decrease over time.

In order to quantify this *value leakage*, the retrofit manufacturer may use the survival curves of aircraft. Survival curves indicate how many aircraft of a given age are still in operation somewhere in the world. To derive the survival curve for the *Airbus* A320 family, the entire production list of A318, A319, A320 and A321 aircraft is first collected. The publicly available database published by *Airlinerlist*¹ contains detailed information about each and every tail number and indicates whether the aircraft is in operation, stored, or withdrawn from use. Using this information, the fraction of aircraft

¹ Available at www.planelist.net and retrieved in August 2015

still in operations in a given age cross-section is estimated. Repeating this exercise for every possible age cross-section yields the survival curve. The survival curve for the Airbus A320 family of aircraft is given in Figure 108. Since the survival curve is quite noisy, it is smoothed using a custom implementation of a non-parametric local regression algorithm (locally weighted scatterplot smoothing or LOWESS [277]).

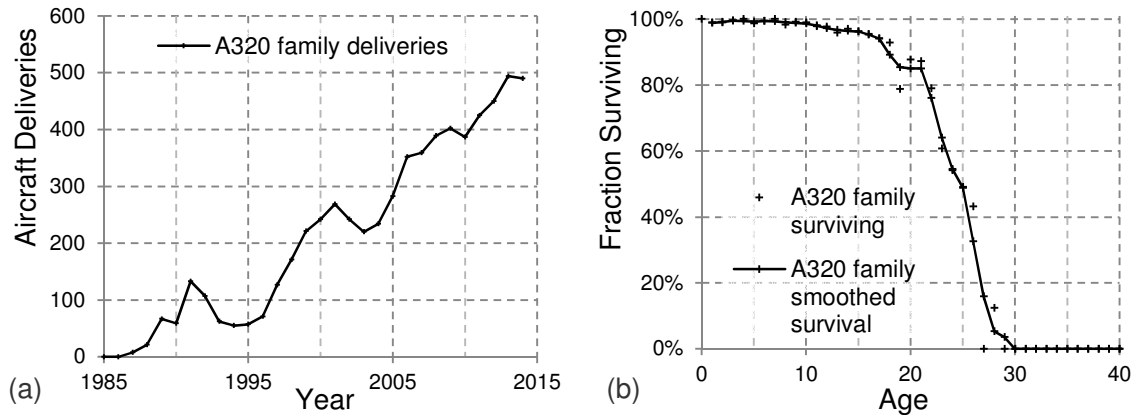


Figure 108: Yearly deliveries in (a) and survival curve in (b) for the Airbus A320 family of aircraft

The A320 family of aircraft being rather recent and to make sure that “boundary effects” due to the newness of the jets do not truncate the survival curves, another set of deliveries and current operating status is analyzed for the *Boeing 737 Classic* family of aircraft (737-200, 737-300, 737-400, 737-500) and for the *McDonnell Douglas MD80* family of aircraft (MD80, MD81, MD82, MD83, MD87, MD88). The end-result is showcased in Figure 109 and exhibits somewhat similar patterns.

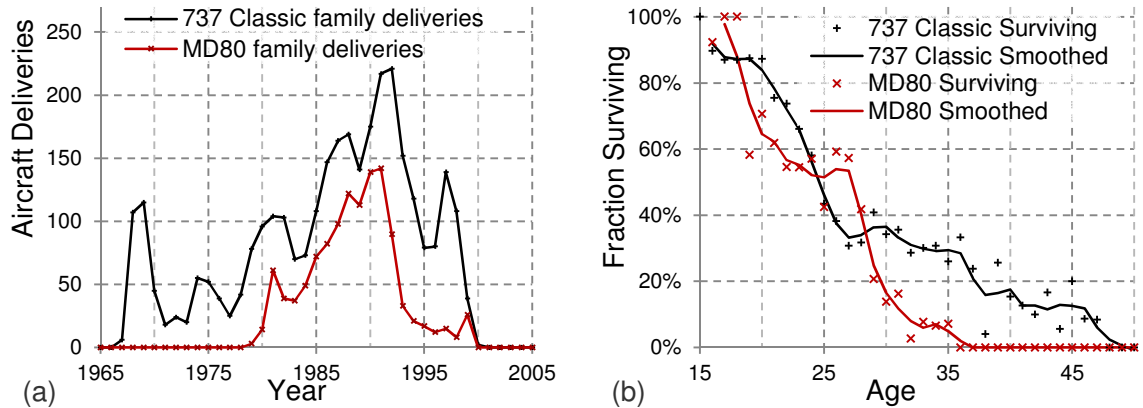


Figure 109: Yearly deliveries in (a) and survival curves in (b) for the 737 Classic family and MD80 family

Using the survival curves for the three fleets of narrowbody aircraft, information can be obtained regarding the market for the performance improvement package retrofit. Indeed, it seems that the target market is decreasing at a rate of 8% per year after the first fifteen years of operations as highlighted in Figure 110. Indeed, during these initial fifteen years, the aircraft are new and efficient, and operators do not retire any. After these initial fifteen years, increased maintenance costs as well as obsolescence lead to a rapid retirement of aircraft. How does this retirement behavior affect the target market for the performance improvement package? Let's first investigate when operators would purchase the retrofit. The performance improvement package is installed into engines during engine maintenance shop visits during which the engine is completely removed from the aircraft. For typical modern turbofan engines powering short- to medium-haul narrowbody aircraft, the first engine shop visit usually occurs after 13,000 to 15,000 flight cycles for the highest thrust ratings (powering the larger A321 and 737-900) and after 18,000 to 20,000 flight cycles for the lowest thrust ratings (powering the A319 and 737-700). This corresponds to about nine to eleven years of on-wing operations for the higher thrust engines and to thirteen to fifteen years for the lower thrust engines.

Consequently, few if any retrofits will be installed during the first twelve years of operations (on average) and the first retrofit installations will roughly coincide with the age cross-section that starts to experience retirement at an average rate of 8% per year.




Figure 110: Calibration of the target market size shrink over time

8.4 Market uncertainties driving PIP demand

This section deals with the selection of the market uncertainties that affect most the development program under review. As could be expected, there is a myriad of market uncertainties that can potentially affect the development of the performance improvement package. These range from uncertainties related to demand (demand for air transportation), uncertainties related to energy prices (market price of jet-fuel), uncertainties related to maintenance labor (hourly rate of maintenance workers), uncertainties related to material costs (market price of titanium and nickel), to uncertainties related to new environmental taxation (cap and trade schemes for carbon emissions). Unfortunately, it is not possible to include all of these uncertainties as the dimensionality of the problem becomes too large and results become less tractable anyways.

8.4.1 Down-selecting uncertainties

The price of jet-fuel has traditionally been one significant source of risk for airlines for two reasons: first, airlines are particularly exposed to the price of jet-fuel as twenty to forty percent of the total aircraft-related operating costs are related to fuel expenditures; second, the high volatility of jet-fuel price compounds this exposure and has historically proven to be a problem for airlines. As a result, airlines have shown a lot of creativity in order to limit their exposure to the fluctuations of jet-fuel price over time. For instance, some airlines have undersigned derivative contracts to mitigate the impact of unknown future fuel price by locking it ahead of time in exchange of the payment of a premium, while a major American airline purchased an oil-refinery to “hedge” some of its fuel expenditures [278]. Still, the most efficient hedge against the volatility of jet-fuel price is probably the use of more efficient aircraft and engines, which is why airlines are constantly pushing aircraft and engine manufacturers to design more efficient aircraft and engines [279]. As a result, the price of jet-fuel is a major driver impacting many aerospace developments and is retained for the analysis.

Airlines in Europe are facing a new form of taxation based on their carbon emissions. Given the novelty of this new form of taxation and the uncertainty regarding the price of future carbon emissions allowances as well as the uncertainty regarding the amount of carbon emissions that will be subject to the trading scheme, this presents an interesting area to investigate. In particular, including the carbon emission price uncertainty into the analysis may help understand whether the trading scheme as currently implemented in Europe is a strong enough incentive for airlines to purchase new technology packages improving the fuel-burn of their existing assets.

8.4.2 Calibration of jet-fuel price model

The jet-fuel price analysis is performed using data from the United States Energy Information Administration representing the historical time series of U.S. Gulf Coast kerosene-type jet-fuel spot price [34]. The time series is plotted in Figure 111 and looks similar to many financial time series with high volatility and no obvious autocorrelation structure.




Figure 111: Closing price of jet fuel (left graph); Continuously compounded daily jet-fuel price returns (right graph)

Inspection of the continuous returns of the price time series indicates a bell-shaped distribution of the return centered on zero with some clustering of high volatility as shown in Figure 112. Despite this heteroscedasticity, a stochastic model similar to a random walk, the Geometric Brownian Motion (GBM), is hypothesized.




Figure 112: Distribution of daily jet-fuel price returns

Several statistical tests are run to check whether the geometric Brownian motion hypothesis can be rejected at the usual 5% level of significance. The first set of tests involves the Variance Ratio test as described by Campbell et al. [280]. These tests check

the correlation structure of increments which should be uncorrelated under the geometric Brownian motion assumption. The null hypothesis for these tests is that increments are uncorrelated. The autocorrelation structure is studied at lags 2, 4, 8, and 16 days. As shown in Table 72, the p -values are all above the 5% level of significance and therefore all tests fail to reject the geometric Brownian motion hypothesis.

Table 72: Variance ratio tests for jet-fuel price time series

| | Lag 2 | Lag 4 | Lag 8 | Lag 16 |
|--|---------------|---------------|--------------|---------------|
| Daily return average (%) | -0.005 | -0.005 | -0.005 | -0.005 |
| Daily return variance (%) | 0.00077 | 0.00077 | 0.00077 | 0.00077 |
| Lagged variance (%) | 0.00076 | 0.00072 | 0.00064 | 0.00067 |
| Variance Ratio (2) Standardized Test Statistic | -0.162 | -1.095 | -1.765 | -0.928 |
| Variance Ratio (2) Standardized Test Critical Value (5%) | 1.960 | 1.960 | 1.960 | 1.960 |
| p-value | 87.2 % | 27.3 % | 7.8 % | 35.3 % |

The second test is the Cowles-Jones Ratio test described again in Campbell et al. [279]. It checks the dependency of the increments which should be independently and identically distributed under the geometric Brownian motion assumption. The null hypothesis for this test is that increments are independently and identically distributed. As shown in Table 73, this test also fails to reject the geometric Brownian motion hypothesis at the 5% level of significance. This means that the apparent heteroscedasticity previously observed is not significant.

Table 73: Cowles-Jones (CJ) ratio test for jet-fuel price


| | |
|-----------------------------|---------------|
| CJ Value | 1.00 |
| CJ Test Value | 0.06 |
| Critical Value (5%) | 1.960 |
| p-value | 95.1 % |

Based on these results, a geometric Brownian motion is used to model the stochastic process driving the price of jet-fuel. The stochastic differential equation is given in Eq. 60 with the Wiener process (W_t), the spot price (S_t), the yearly drift (μ) that measures the trend, and the yearly volatility (σ) that measures the variability of the process over time.

$$dS_t = \mu \cdot S_t \cdot dt + \sigma \cdot S_t \cdot dW_t \quad \text{Eq. 60}$$

$$S_0 = 2.75 \text{ US\$}; \mu = .005\%; \sigma = 43.9\%$$

The geometric Brownian motion model is fine to represent the past evolution of the spot price of jet-fuel. In particular, it yields good information regarding the volatility of the jet-fuel price process. However, it does not give any information regarding the price of jet-fuel in the future and there is no reason to hypothesize that the future will repeat the past or follow the same trends as those recently experienced. In order to look into the future so as to estimate the drift of the process, derivative contracts called futures are used. Futures are contract that allow the purchase ahead of time of a specific amount of a commodity to be delivered at a specific date for a pre-specified price. In an efficient market, the market price of futures on jet-fuel reflects all the knowledge of the market and therefore the expectation of the market regarding its evolution in the future. Consequently, the initial value and the drift of the geometric Brownian motion process are calibrated using market information concerning futures available from the *Chicago Mercantile Exchange* [281] as shown in Figure 113.



$$dS_t = \mu \cdot S_t \cdot dt + \sigma \cdot S_t \cdot dW_t$$

$$S_0 = 1.38 \text{ US\$}; \mu = 7.5\%; \sigma = 43.9\%$$

Figure 113: Jet-fuel futures quotes and stochastic process retained

8.4.3 Calibration of carbon emission allowances price model

The Emissions Trading Scheme requires airlines to buy permits for about fifteen percent of the airlines' carbon dioxide emissions. These permits are in limited quantity and may be purchased on the carbon market in the form of European Union Allowances (EUA). For instance, Szabo reports that one large European airline started using the *BlueNext* exchange platform in 2012 to buy allowances on the spot market [282]. The carbon emission analysis is therefore performed using the time series *BNS EUA 08-12* available on the exchange website for data from February 2008 to June 2012 [283]. This time series is plotted in Figure 114 and, like in the previous example, it exhibits high volatility with no obvious autocorrelation structure, but with a downward trend.




Figure 114: Closing price of EUA (left graph); Continuously compounded daily returns (right graph)

Inspection of the continuous returns distribution displayed in Figure 115 indicates a bell-shaped distribution of the returns centered on zero with some clustering of high volatility.




Figure 115: Distribution of daily EUA price returns

Based on these observations, a geometric Brownian motion is again hypothesized. The same statistical tests described in the previous section for jet-fuel prices are run to check whether the GBM assumption for emission costs can be rejected at the 5% level of significance. The first set of tests is the Variance Ratio test which is run for lags 2, 4, 8, and 16 days. Each time, the geometric Brownian hypothesis cannot be rejected at the 5% level of significance as shown in Table 74.

Table 74: Variance ratio test for CO₂ emissions

| | Lag 2 | Lag 4 | Lag 8 | Lag 16 |
|--|---------------|---------------|---------------|---------------|
| Daily return average (%) | -0.114 | -0.114 | -0.114 | -0.114 |
| Daily return variance (%) | 0.00075 | 0.00075 | 0.00075 | 0.00075 |
| Lagged variance (%) | 0.00078 | 0.00075 | 0.00073 | 0.00076 |
| Variance Ratio (2) Standardized Test Statistic | 1.211 | 0.101 | -0.320 | 0.055 |
| Variance Ratio (2) Standardized Test Critical Value (5%) | 1.960 | 1.960 | 1.960 | 1.960 |
| p-value (under H₀) | 22.6 % | 91.9 % | 74.9 % | 95.6 % |

Next is the Cowles-Jones Ratio test which also fails to reject the geometric Brownian motion hypothesis at the 5% level of significance, as shown in Table 75.

Table 75: Cowles-Jones (CJ) ratio test for CO₂ emission allowances


| | |
|---|--------------|
| CJ Value | 1.02 |
| CJ Test Statistic Value | 0.28 |
| Critical Value at 5% | 1.960 |
| <i>p</i> -value (Hypothesis: IID Increments) | 78.3% |

Based on these results, a Geometric Brownian Motion is used to model the stochastic process driving the price of carbon allowances and its parameters are provided in Eq. 61:

$$dS_t = \mu \cdot S_t \cdot dt + \sigma \cdot S_t \cdot dW_t \quad \text{Eq. 61}$$

$$S_t = 8.25 \text{ EUR} ; \mu = -28.8\% ; \sigma = 43.5\%$$

As before, to estimate the initial price and the drift of the geometric Brownian process, futures contract retrieved from the *European Energy Exchange* [284] with delivery date during the period of interest are used and the results are reported in Figure 116.



$$dS_t = \mu \cdot S_t \cdot dt + \sigma \cdot S_t \cdot dW_t$$

$$S_t = 9.62 \text{ US\$} ; \mu = 1.53\% ; \sigma = 43.5\%$$

Figure 116: E.U. Allowance futures quotes and stochastic process retained for subsequent analyses

8.4.4 Treatment of correlations

The two stochastic processes retained for the modeling of jet-fuel price and carbon emission cost uncertainties are independent models. However, a more intricate

relationship between the two models is likely. Indeed, a period of strong growth in Europe may result in higher demand for air transportation and therefore higher prices for jet-fuel. Similarly, this higher demand for air transportation may result in more demand for carbon permits and therefore higher emission allowance prices.

The relationship between the price of jet-fuel and the price of carbon permits can be captured with a correlation matrix. This matrix is estimated by first cleaning the time series to ensure that quotes are available for both prices on the same date, and then estimating the correlation between the continuous returns of each time series. The correlation matrix is given in Eq. 62 and indicates a correlation of 19% between the two data series.

$$M_{Cor} = \begin{bmatrix} 1 & 0.199 \\ 0.199 & 1 \end{bmatrix} \quad \text{Eq. 62}$$

To include this correlation in the two stochastic models previously defined, correlated numbers need to be sampled from the standard normal distribution used in the geometric Brownian motion. This is performed using a Cholesky decomposition of the correlation matrix as shown in Eq. 63. The positive definite correlation matrix is decomposed to give a lower-triangular matrix which, when applied to a vector of uncorrelated samples, produces a sample vector with the correlation properties of the system being modeled.

$$M_{Cor} = C \cdot C^T ; C = \begin{bmatrix} 1 & 0 \\ 0.199 & 0.979 \end{bmatrix} \quad \text{Eq. 63}$$

CHAPTER 9: EXPERIMENTAL PLAN - VALIDATION

The previous chapter introduced the proof-of-concept application that is used next for the validation effort of this research. According to the Project Management Institute [201], the purpose of validation is to ensure *"that a product, service, or system meets the needs of the customer and other identified stakeholders. It often involves acceptance and suitability with external customers."* In the context of this methodology development, validation is performed to ensure that the method and its associated hypotheses actually help resolve the issues formalized in the different research questions. As a result, while the technical hypotheses were verified to ensure mathematical soundness, the method and modeling hypotheses are validated to ensure that the proposed mathematical abstractions are indeed adequate, suitable for the envisioned applications, and finally, that they properly represent all pertinent aspects of the problem.

9.1 Real options to model managerial flexibility

One top-level hypothesis in this research effort is that staggered aviation technology investments can be studied using real options analysis in order to capture the managerial flexibility offered to decision-makers. Having introduced this method hypothesis, a thorough literature review identified some gaps and led to further questioning: how can these real options methods be adapted to the problem faced by the aviation industry? This resulted in two second-level research questions, *"modeling research questions"*, and one associated second-level hypothesis, *"modeling hypothesis"*. Indeed, assuming that real options inspired methodologies present the best framework for

the analyses of long-term and highly uncertain research and development programs, the first modeling research question is related to the pertinence of these models to evaluate investments in the aerospace industry while the second modeling research question pertains to the modeling of flexibility offered to decision-makers to optimally time investments.

Research Question 1.1 — Creation of an option-thinking framework

Within the context of uncertain product and technology investment analysis, how can state-of-the-art option-based valuation methods be improved upon to ensure their domain of application is consistent with their underpinning assumptions?

Research Question 1.2 – Managerial flexibility and timing of investments

How can the flexibility offered to management to optimally time the launch of new product and technology developments be accounted for in a real options framework?

A review of the existing literature on financial options and real options points to the use of path-dependent real options for a more accurate depiction of investments in the aerospace industry. While retaining the ability to handle the managerial flexibility to fund or abandon subsequent phases in staggered developments offered by European real options, path-dependent real options enable the relaxation of the deterministic exercise time of this flexibility. In other terms, it relaxes one assumption of the more popular models by accounting for the additional flexibility to optimally time investment decisions: the flexibility is no-longer offered at a single pre-determined point in time but rather over an entire decision window. This yields the formulation of the following modeling hypothesis:

Hypothesis 1.1 – Path-dependent options to model managerial flexibility

As uncertainty unfolds, technological and market opportunities emerge and disappear. Flexible management and flexible timing of investment decisions allow the maximization of the upside potential of these opportunities. Path-dependent real options may present a means to model the flexibility offered to management in timing technology development decisions.

9.1.1 Validation process and criteria for success

This hypothesis claims that there is value created by a flexible management able to optimally time decisions regarding research and development programs, something not typically captured by traditional deterministic discounted cash flow analyses. Besides, the hypothesis states that using path-dependent real options enables the modeling of this timing flexibility and therefore enables the capture of the additional value created by active and astute management. Without loss of generality, timing flexibility refers to the ability to make an investment decision over the course of a period of time referred to as the decision window. Intuitively, this flexibility creates additional value for the business and the aim is to capture this additional value created by astute timing of investments.

To validate this hypothesis, several investigations are performed using the performance improvement package proof-of-concept study. Managerial flexibility is first introduced as the option to abandon the development program at the first decision tollgate after an initial one-year period of market research. Timing flexibility is introduced next as the possibility, if market conditions warrant it, to bring forward this decision tollgate by

reducing the amount of time spent on the market research and by starting ahead of schedule the following phase of detailed design with the aim of speeding-up the entry into service of the technology retrofit. As a consequence, there is a one year decision window during which managers have the option to either invest in the subsequent phase of the development program or to let the option expire and abandon the development program (with no salvage value). Intuitively, this timing flexibility brings additional value when compared: (1) against a *flexible* scenario with the option to invest only at the end of the initial market research (basically at the end of the decision window), and (2) against a reference *determinist* scenario with a deterministic investment decision made at the start of the initial market research with no possibility of abandonment after. The determinist, flexible, and fully flexible approaches are evaluated and compared so that a flexibility premium can be quantified in each case. For the hypothesis to be verified, this premium must be strictly positive and must be statistically significant at the 5% level of significance.

9.1.2 Preliminary testing and initial struggles

Repeated experiments with different model parameters indicate that the logit market preference model is not adequate due to excessive sensitivity. What happens is that, when the price of jet-fuel and the price of carbon emission allowances are varied, the value of the program to the manufacturer becomes extremely “volatile”, almost binary, with a roughly constant very high value and a roughly constant very low value. This is because the market preference is modeled as a steep sigmoid function that is so steep that it “saturates early” indicating either a constant 0% market preference if conditions are not favorable and almost immediately a 100% market preference whenever

the retrofit starts being profitable for airlines. With hindsight, this problem was to be expected as the PIP-Light package price does not fluctuate with demand: in the model, if the retrofit package becomes desirable for airlines, the demand increases immediately without having any form of demand-price equilibrium.

In order to bypass this problem, a different approach is undertaken in the following steps. Instead of estimating the market reaction to the retrofit, the value to the manufacturer is modeled as a fraction of the value to all its customers. This assumption is not invalid as it was reported in the literature that the NEO option for the Airbus A320 would be priced at roughly half the saving experienced by the operators [276]. This is the option that is retained for the remaining of the research: the manufacturer gets 50% of the net present value of the fuel, emission, and maintenance savings experienced by airlines.

9.1.3 Performance Improvement Package baseline evaluation

In order to prove that managerial flexibility has value and that this value is accounted for in the proposed real options analysis framework, a first step is to perform an evaluation of the technology retrofit development program using traditional capital budgeting analysis techniques. In Chapter 3, different capital budgeting techniques were reviewed and the real options approach was introduced as an extension of the discounted cash flow analysis accounting for the value of managerial flexibility. Therefore, it is natural to use the discounted cash flow technique as a reference against which the improved real options method may be benchmarked.

Baseline PIP development program assumptions

The discounted cash flow analysis requires the estimation of cash inflows and cash outflows during the life of the development program. To derive these cash flows, assumptions are made regarding the future state of the world and a jet-fuel price and carbon emission price scenario is constructed using the expectation of stochastic models for the jet-fuel and carbon emission allowances prices presented in Chapter 8. The development timeline and the development assumptions are given in Figure 117 and Table 76 respectively.




Figure 117: Baseline PIP development timeline

Table 76: Baseline PIP development assumptions

| <i>Market assumptions</i> | | <i>Development program assumptions</i> | |
|------------------------------|----------------|---|--------|
| PIP retrofit operating years | 25 | 2018 jet-fuel price (US\$/USG) | \$1.75 |
| Potential market size | 6,000 | 2018 carbon allowance price (2015-US\$/tCO ₂) | \$9.91 |
| Market share | 40% | Share of economic surplus | 40% |
| Market size shrink | 8% | Gross profit margin on sales | 50% |
| Customer types of operations | 11 | Manufacturer WACC | 9.0% |
| Customer WACC | 7.2% and 10.5% | Initial market research (M US\$) | 5.0 |
| | | Development cost (M US\$) | 95.0 |

Baseline PIP operating benefits for customers

The operating benefits for prospective customers are computed on a market by market basis using the twenty two market segments defined in Chapter 8 using two different costs of capital and eleven different types of operations represented by eleven different average flight lengths. For each of these markets, the fuel-burn reductions, the

carbon emission reductions, and the maintenance expenditure reductions are estimated over the projected number of years of operations. The projected time in operations of the PIP retrofit package is given in numbers of engine shop visits since the technology retrofit is installed during one shop visit and the engine on which it is installed is usually retired right before one subsequent shop visit (to save the maintenance expenditure). Using a maximum of 25 years in operations, the expected time-on-wing is computed. The main outputs of this analysis are given in Figure 118 and more details can be found in APPENDIX K.

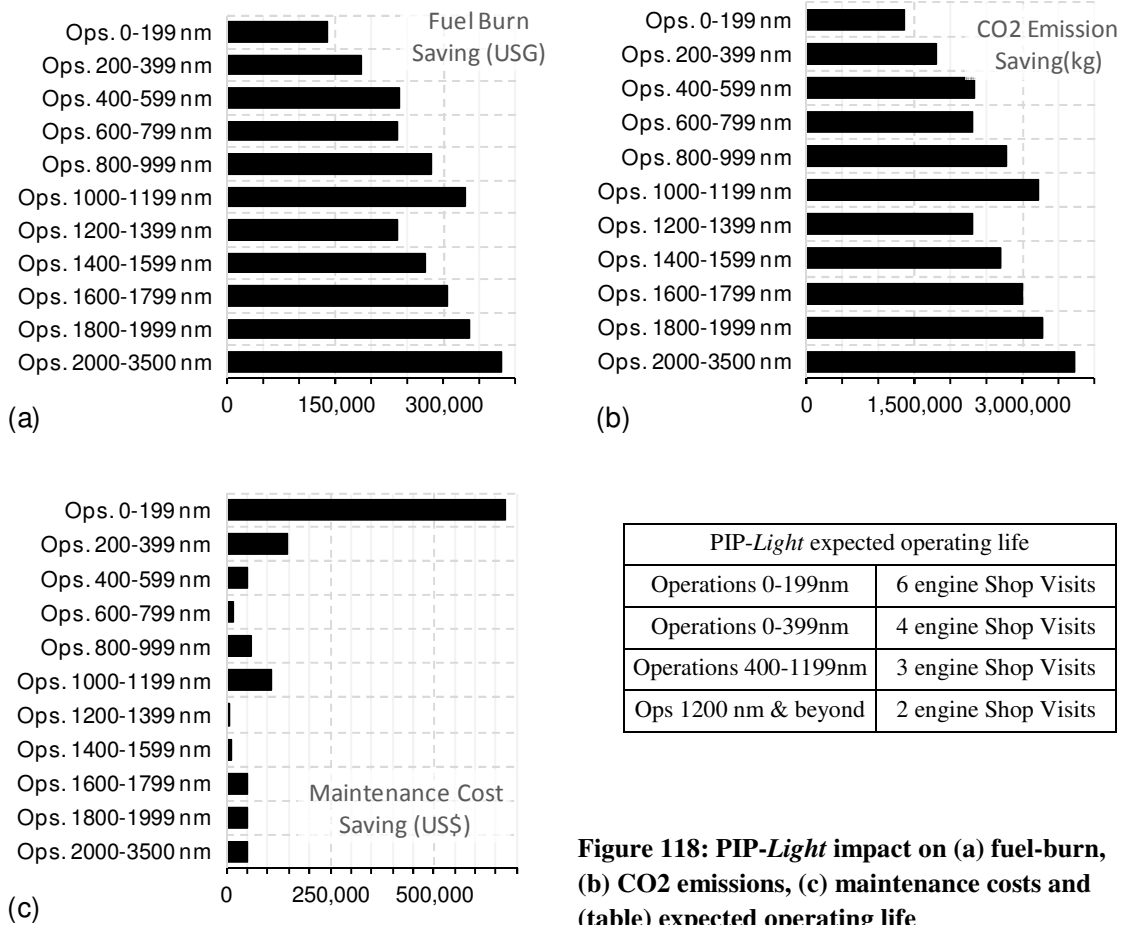


Figure 118: PIP-Light impact on (a) fuel-burn, (b) CO2 emissions, (c) maintenance costs and (table) expected operating life

Table 77 represents the net present value of the PIP retrofit for customers in each market segment as well as the size of each market segment.

Table 77: baseline PIP statistics with customer NPV and market size

| Market Segment | Customer NPV (in 2018) | | Market Size (Engines #) | |
|-------------------------|---------------------------|-------------------|----------------------------|-------------------|
| | High CAPEX markets | Low CAPEX markets | High CAPEX markets | Low CAPEX markets |
| Operations 0-199 nm | \$371,934 | \$273,514 | 54 | 54 |
| Operations 200-399 nm | \$194,438 | \$147,179 | 334 | 334 |
| Operations 400-599 nm | \$191,928 | \$147,134 | 460 | 460 |
| Operations 600-799 nm | \$182,340 | \$145,965 | 539 | 539 |
| Operations 800-999 nm | \$227,777 | \$179,505 | 469 | 469 |
| Operations 1000-1199 nm | \$279,204 | \$220,802 | 331 | 331 |
| Operations 1200-1399 nm | \$201,242 | \$168,292 | 284 | 284 |
| Operations 1400-1599 nm | \$229,429 | \$190,451 | 227 | 227 |
| Operations 1600-1799 nm | \$273,275 | \$231,374 | 148 | 148 |
| Operations 1800-1999 nm | \$301,722 | \$254,200 | 189 | 189 |
| Operations 2000-3500 nm | \$332,145 | \$277,210 | 117 | 117 |

Baseline PIP development program economic analysis

The PIP development program net present value is computed assuming no flexibility during the entire development. The entire investment is committed from the start in 2015 and the initial decision is an all-or-nothing decision: even though the program is staggered with four phases and three milestones, there is no provision in the analysis to either abandon the development at the first milestone before detailed design starts in 2016, or at the second milestone before certification and testing starts in 2017, or at the third milestone in 2018 right before production starts. The results are shown in Table 78.

Table 78: Baseline PIP development program value

| <i>Development program financial metrics</i> | |
|--|-------|
| PIP-Light revenues (2017-MUS\$) | 133.7 |
| PIP-Light net present value (2015-MUS\$) | 10.3 |

9.1.4 Limited-flexibility Performance Improvement Package development

In this section, some managerial flexibility is introduced and the decision-makers have the ability to either launch the development of the performance improvement package by funding the detailed design analysis, the certification and testing, and the production, or abandon the development altogether at the end of the initial market research.

Limited-flexibility PIP development program assumptions

In this case, the exploratory phase lasts one year between 2015 and 2016. At the end of this exploratory phase, there is flexibility to fund the following phases of development stretching until 2018 and the production coming thereafter. The alternative is to treat the investment in the exploratory phase as a sunk cost and abandon the project in 2016 if the market conditions are not favorable. As a consequence, funding the exploratory phase gives decision-makers the option to initiate the following phase of development. This is represented using a European call option on the revenues of the project with a maturity of one year and an exercise price corresponding to the entire development program cost minus the (rather low) cost of the exploratory phase. The timeline of the limited-flexibility development program is described in Figure 119.




Figure 119: Limited-flexibility PIP development timeline

Most of the other assumptions for the limited-flexibility PIP development program are identical to those presented previously in Table 76. The main difference is that two correlated stochastic processes are now used for the evolution of the jet-fuel price and for the evolution of the carbon emission allowances using the results of the calibration done in Chapter 8 in lieu of the expected 2018 prices of jet-fuel and carbon emission allowances. A summary of the main inputs is proposed in Table 79.

Table 79: Limited-flexibility PIP development assumptions

| <i>Market assumptions</i> | | <i>Development program assumptions</i> | |
|------------------------------|----------------|---|--------|
| PIP retrofit operating years | 25 | 2015 jet-fuel price (US\$/USG) | \$1.38 |
| Potential market size | 6,000 | 2015 carbon allowance price (2015-US\$/tCO ₂) | \$9.62 |
| Market share | 40% | Share of economic surplus | 40% |
| Market size shrink | 8% | Gross profit margin on sales | 50% |
| Customer types of operations | 11 | Initial market research (M US\$) | 5.0 |
| Customer WACC | 7.2% and 10.5% | Development cost (M US\$) | 95.0 |

Limited-flexibility PIP development operating benefits for customers

The operating benefits for the customers presented previously in Table 76 and Table 77 are retained for the limited-flexibility PIP development analysis.

Limited-flexibility PIP development program economic analysis

The results are provided in Table 80 and indicate a substantial increase in the program value. Because the analysis is based on Monte Carlo simulations which introduce some variability in the results, the standard error of the limited-flexibility PIP

development program value is also computed which enables the estimation of a 95% confidence interval for the program value. It is worth mentioning that the 95% confidence interval does not overlap the previously computed baseline program value. This indicates that there is a positive and statistically significant premium due to the ability to abandon the program in 2016. Mathematically speaking, this increase in value corresponds to the ability to prune unfavorable Monte Carlo trajectories at the end of the exploratory phase. Managerially speaking, this increase in value corresponds to the value created by astute management of business ventures by decision-makers. It is important to realize that this added value is not created out of thin air by the real options analysis: it always exists but it is not accounted for in typical discounted cash flow analyses.

Table 80: Limited-flexibility PIP development program value

| <i>Development program financial metrics</i> | |
|---|--------------|
| PIP-Light value (Excluding exploratory phase) (2014-MUS\$) | 18.5 |
| Exploratory phase cost (2014-MUS\$) | 5.0 |
| PIP-Light value (Total) (2014-MUS\$) | 13.5 |
| PIP-Light value standard error (2014-KUS\$) | 57.7 |
| PIP-Light value 95% confidence interval (2014-MUS\$) | 13.4 to 13.7 |

9.1.5 Flexible Performance Improvement Package evaluation

In this section, additional managerial flexibility is introduced and the decision-makers have the ability to either launch the development of the performance improvement package by funding the detailed design analysis, the certification and

testing, and the production at any point in time between 2015 and 2016. However, by 2016, a decision must be made, and if the detailed design phase is not funded, then the entire development program is abandoned.

Flexible PIP development program assumptions

In this case, the exploratory phase lasts up to one year between 2015 and 2016. However, this phase of development can be shortened in order to rush the development and start the detailed design phase early if market conditions are very favorable. However, by the end of 2016, a decision must be made whether to pursue or abandon the PIP development. If the development is abandoned, the investment in the exploratory phase is a sunk cost. If market conditions are favorable, then the detailed design, certification and testing, and production phases are funded. As a consequence, funding the exploratory phase gives decision-makers the option to initiate the following phase of development. This is represented using an American call option on the revenues of the project with a maturity of one year and an exercise price corresponding to the entire development program cost minus the (rather low) cost of the exploratory phase. The timeline of the flexible development program is described in Figure 120. Most of the other assumptions for the flexible PIP development program are identical to those presented previously in Table 79.




Figure 120: Flexible PIP development timeline

Flexible PIP development operating benefits for customers

The operating benefits for the customers presented previously in Table 76 and Table 77 are retained for the flexible PIP development analysis.

Flexible PIP development program economic analysis

The results are provided in Table 81 and indicate a significant increase in the program value. Because the analysis is based on Monte Carlo simulations which introduce some variability in the results, the standard error of the flexible PIP development program value is computed which enables the estimation of a 95% confidence interval for the program value. It is worth mentioning that the 95% confidence interval does neither overlap the previously computed baseline program value nor the previously computed limited-flexibility program value. This indicates that there is a positive and statistically significant premium due to the ability to rush the development program and start working on the detailed design early to allow an earlier entry into service and the capture of additional customers. Mathematically speaking, this increase in value corresponds to the ability to select extremely favorable Monte Carlo trajectories and bring their larger payoffs (less subject to value leakages) forward in time. Managerially speaking, this increase in value corresponds to the value created by astute management of business ventures by decision-makers and the possibility to time investment decisions early. It is important to realize that this added value is not created out of thin air by the real options analysis: it always exists but it is neither accounted for in typical discounted cash flow analyses nor in simpler real options models featuring only European types of options.

Table 81: Flexible PIP development program value

| <i>Development program investment statistics</i> | |
|--|--------------|
| Probability of program launch | 56% |
| Expected time to launch (months) | 8 |
| <i>Development program financial metrics</i> | |
| PIP-Light value (Excluding exploratory phase) (2015-MUS\$) | 19.8 |
| Exploratory phase cost (2015-MUS\$) | 5.0 |
| PIP-Light value (Total) (2015-MUS\$) | 14.8 |
| PIP-Light value standard error (2015-KUS\$) | 49.5 |
| PIP-Light value 95% confidence interval (2015-MUS\$) | 14.7 to 14.9 |

9.1.6 Fully flexible Performance Improvement Package evaluation

In this section, even more managerial flexibility is introduced and the decision-makers have the ability to launch the development of the performance improvement package by funding the detailed design analysis at any point in time between 2015 and 2016. However, by 2016, a decision must be made, and if the detailed design phase is not funded, then the entire development program is abandoned. In addition, in 2017, decision-makers have again the ability to abandon the program or proceed with the development and fund the remaining phases of certification and testing and production.

Fully flexible PIP development program assumptions

In this case, the exploratory phase lasts up to one year between 2015 and 2016. However, this phase of development can be shortened in order to rush the development and start the detailed design phase early if market conditions are very favorable.

However, by the end of 2016, a decision must be made whether to pursue or abandon the PIP development. If the development is abandoned, the investment in the exploratory phase is a sunk cost. If market conditions are favorable, then the detailed design phase is funded at the end of which decision-makers are again offered the option to fund the certification and testing and production phases. As a consequence, funding the exploratory phase gives decision-makers the compound option to initiate the detailed phase of development which in turn gives the option to initiate the following phases of development. This is represented using an American call option with a maturity of one year on a European call option on the revenues of the project with a maturity of one year and an exercise price corresponding to the entire development program cost minus the costs of the exploratory and detailed design phases. The timeline of the fully flexible development program is described in Figure 121. Most of the other assumptions for the fully flexible PIP development program are identical to those presented previously presented in Table 79.




Figure 121: Fully flexible PIP development timeline

Fully flexible PIP development operating benefits for customers

The operating benefits for the customers presented previously in Table 76 and Table 77 are retained for the fully flexible PIP development analysis.

Fully flexible PIP development program economic analysis

The results are provided in Table 82 and indicate a small increase in the program value. Because the analysis is based on Monte Carlo simulations which introduce some variability in the results, the standard error of the fully flexible PIP development program value is computed which enables the estimation of a 95% confidence interval for the program value. However, even though the 95% confidence interval does overlap neither the previously computed baseline program value nor the limited-flexibility program value, it does overlap the 95% confidence interval of the flexible program value. This indicates that there is probably a positive premium for the American-European nested option (it was just calculated), but the variability in the results prevents it from being statistically significant.

Table 82: Fully flexible PIP development program value

| <i>Development program investment statistics</i> | |
|--|--------------|
| Probability of program launch | 61% |
| Expected time to launch (months) | 7 |
| <i>Development program financial metrics</i> | |
| PIP-Light value (Excluding exploratory phase) (2015-MUS\$) | 19.9 |
| Exploratory phase cost (2015-MUS\$) | 5.0 |
| PIP-Light value (Total) (2015-MUS\$) | 14.9 |
| PIP-Light value standard error (2015-KUS\$) | 62.7 |
| PIP-Light value 95% confidence interval (2015-MUS\$) | 14.8 to 15.1 |

What is more interesting though is that, everything else staying constant, the probability of program launch increased from 56% to 61% and the expected time before

committing to the launch decreased from 8 months to 7 months. This result makes sense because the program does provide an abandonment option further into the future if conditions become unfavorable. As a result, the conditions do not need to be as favorable as in the simpler American option case to rush the investment and to start investing early since the development program is less risky.

9.1.7 Comparisons

The PIP-Light program value and the corresponding trigger boundary are given in Figure 122. These values are normalized by the development cost for the PIP-Light program (US\$100,000,000).




Figure 122: Trigger boundary for PIP-Light development program

As can be observed, the value of the development program featuring timing flexibility is substantially higher than the value of the development program without timing flexibility (about 10% higher). In the first case, the PIP-Light program launch should be delayed as current conditions do not warrant an immediate investment. In addition, the current market specification suggests that the PIP-Light development will be launched with a 61% probability and that the average time to the optimal launch is about

seven months. This yields an indication of the time available to refine or finish maturing technologies that are to be used in the PIP-Light package. In the second case, the flexibility that decision-makers have to potentially launch the program early is not accounted for despite the fact that it could be optimal. As a result, some managerial flexibility value is not accounted for and the program value is less.

Since no other parameter is changed between the first and the second analysis, the additional timing flexibility must be what is driving the additional 10% of PIP-Light program value. Finally, the 95% confidence interval for the value of the PIP-Light development program with timing flexibility (Case 1) does not contain the value of the PIP-Light development program with no timing flexibility (Case 2). Therefore, the premium for the early-exercise flexibility is statistically significant and the use of a real options approach coupled with the use of path-dependent options enables analysts to capture the value created by astute management and the value created by timing optimally the launch of research and development programs.

Between the second and the third case, the only change in the modeling corresponds to the flexibility to launch the second phase of the development program if market conditions are favorable. As a consequence, the resulting change in program value corresponds to the flexibility to abandon unprofitable ventures. Thus, the use of a real options approach enables analysts to capture the flexibility offered to managers to actively steer research and development programs into profitable directions.


As a conclusion, the implementation and use of the proposed methodology enables the capture of the managerial flexibility to abandon unprofitable ventures and the managerial flexibility to optimally time investments in development programs. The

hypothesis is validated and the additional timing flexibility is accounted for in path-dependent options.

Additional observations

It is worth mentioning that in Figure 122, the trigger boundary crosses the vertical axis at a value of 1.52. This means that the present value of the PIP-Light program must exceed the investment outlay by 52% in order to trigger an immediate launch of the development program. This is to be contrasted with the analysis of the third case using a traditional net present value analysis where the positive net present value of the PIP-Light development program suggests an immediate investment.

Sensitivity studies can also be performed to estimate the impact of model misspecification on the outcome of the analysis. Let's assume for instance that the stochastic process driving the value of the jet-fuel price features some discontinuities and jumps. The same geometric Brownian motion is used but some jumps are added. The presence of jumps is in line with one of the first observation of this research where the frequency of large jumps (1990, 2006, and 2008) was found to be abnormally high using a geometric Brownian motion specification. The additional jumps occur on average once per year with a jump size volatility of 40%. The resulting changes in the program value and in the early-exercise boundary are displayed in Figure 123.




| PIP-Light development program value (2015-MUS\$) | | |
|--|---------------------------|-------------------------------|
| Fuel Price Model | Geometric Brownian Motion | Merton Jump Diffusion Process |
| Limited Flexibility | 18.2 | 22.8 |
| Flexible | 19.8 | 23.9 |

Figure 123: Trigger boundaries for the PIP-Light with jet-fuel price following a geometric Brownian motion and a jump-diffusion process

With a significant increase in the volatility of the jet-fuel time series, the value of the embedded managerial flexibility becomes higher and results in a significant increase in the program development value. The optimal trigger boundary is also shifted up due to the increased variability introduced by the additional jumps: indeed, with increased variability comes increased likelihood that a currently profitable investment becomes loss making within the one year time-window. A simple explanation for this behavior takes the opposite scenario: if there were no variability at all, development programs would only need to be profitable by a single dollar to be in the immediate exercise area since there is no additional information to be gained by waiting, especially if the target market for the retrofit shrinks at a rate of 8% per year.

Another interesting analysis consists in comparing the current results with the value of the PIP-Light development program without accounting for the price of carbon emission allowances. One expects the value of the development program to decrease slightly owing to the decreased incentives airlines have to purchase the PIP-Light retrofit. Case 1 and Case 2 are investigated and the values of the program as well as the trigger boundary are reported in Figure 124. It seems that the taxation of carbon emissions, as it

is currently implemented, has little effects on the overall value of the development program.



| | PIP-Light (Carbon Trading) | PIP-Light (No Carbon Trading) |
|---------------------|-------------------------------|-------------------------------------|
| Limited Flexibility | 18.2 MUS\$ | 17.8 MUS\$ |
| Flexible | 19.8 MUS\$ | 19.2 MUS\$ |

Figure 124: Trigger boundaries for the PIP-Light retrofit with (black) and without (red) carbon emission allowances

Investment timing flexibility: comparisons using PIP-Involved

The analysis proceeds with the second performance improvement package. The PIP-Involved is bringing more benefits in terms of operating cost reductions but it is also offered to the market several years later when the market has already shrunk significantly and at a point in time when airlines may become reluctant to invest in an aging fleet. The parameters for this development program are given in Table 83.

Table 83: Input parameters for PIP-Involved evaluation

| | |
|--|----------------|
| Market size | 4,000 engines |
| Number of years of operations for PIP-Involved | 12 |
| Number of market segment analyzed | 22 |
| Customer WACC | 7.2% and 10.5% |
| Market shrink over time | 8% per year |
| Manufacturer gross profit margin on sales | 50% |
| PIP-Involved development cost | 500M US\$ |
| Risk free rate of return | 2% |

The resulting option price and trigger boundary are given in Figure 125. There are several striking results: the early-exercise boundary for the PIP-Involved retrofit is significantly above the trigger boundary for the PIP-Light project. There are several

reasons for this. First, PIP-Involved is a more risky project owing to the late entry into service. Consequently, launching the program early requires that the market conditions be very favorable. Next, the option value for the PIP-Involved is much lower than the option value for the PIP-Light. This is due to the shrinking market between the introduction of the PIP-Light (in 2018) and the introduction of the PIP-Involved (in 2022). In fact, about 30% of the market has vanished due to aircraft retirements. This makes the more involved and better performing PIP-Involved retrofit package an unprofitable venture for the engine manufacturer.




Figure 125: Comparison of trigger boundaries for PIP-Light(1) and PIP-Involved(2) and comparison of path-dependant option prices for the PIP-Light and PIP-Involved

9.2 Sequential moves for competitive scenarios

Having introduced a method hypothesis advocating the use of game theoretic analyses, a thorough literature review led to some further questioning: how can these game theoretic methods be adapted to the problem under review? Like in the previous section, this led to a second-level research question, “*modeling research question*” and an associated second-level hypothesis, “*modeling hypothesis*”. Indeed, assuming that game theoretic methodologies present a good framework for the analysis of competitive

research and development programs, the first modeling research question is related to the choice of models to evaluate research and development strategies in the aerospace industry.

Research Question 1.3 – Competitive scenario modeling

How can game theoretic analyses be used to adequately model competition in the aerospace industry and how can they be used to identify profitable product and technology development strategies?

Reviewing the literature and observing the nature of the competition, oligopolies have been observed owing to the significant barriers to entry for prospective competitors. Having few competitors reduces the dimension of the problem and makes it easier to list the possible strategies of the competitors. In turn, the identification of these strategies enables the formulation of sequential competitive games. While research and developments are usually made somewhat simultaneously, actual decisions to launch new research and development programs are usually made in a sequential fashion with limited room for “big surprises”. This means that one competitor is a leader while the others wait to see what happens before making a move. This leads to the following modeling hypothesis.

Hypothesis 2.1 – Equilibrium in Sequential moves for competitive scenarios

Equilibrium-types of solutions in sequential competitive scenarios provide means to quantitatively measure the impact of competing designs on profitability and to identify robust strategies.

Validation process and criteria for success

This hypothesis assumes that equilibrium-types solutions in a sequential competitive scenario lead to the most robust strategy. In order to validate this hypothesis, the strategy resulting from an equilibrium solution is compared to all other possible solutions and the performance improvement package proof-of-concept study is used again. Several strategies are defined and investigated for both the performance improvement package retrofit of the original equipment manufacturer as well as for its competitor. For this hypothesis to be verified, perturbations to the strategies defined by the equilibrium solution are first investigated. If there is no incentive for any of the competitors to deviate from the equilibrium solution, then the equilibrium solution is deemed robust and the hypothesis is validated.

Evaluation of strategies using the PIP-Light and PIP-Involved developments

The first step in the assessment of the various strategies offered to the manufacturer and to its competitor is to evaluate the payoffs of these strategies in order to later compare them. To do so, an estimate of the revenues is first computed by making assumptions regarding the market share in the different scenarios. Starting from the present state of the business where the original equipment manufacturer gets 30% of the market share with its replacement parts (life-limited parts and other pieces of the engine), the introduction of the PIP-Light with its improved aging characteristics and better fuel efficiency leads to a gain in market share which reaches 40%. The introduction of the PIP-Involved, with its significantly improved performance leading to better fuel efficiency and drastically reduced maintenance expenditures, attracts more customers and results in a market share of 50%. When the competition introduced its new PMA

package, the market shares in all cases are reduced by 10%. The summary of these hypotheses is given in Table 84.

Table 84: Market share assumptions in different competitive scenarios

| | | <i>Competition</i> | |
|--|--------------|--------------------|-----------|
| | | No Move | New PMA |
| <i>Original Equipment Manufacturer</i> | No Move | 30% - 70% | 20% - 80% |
| | PIP-Light | 40% - 60% | 30% - 70% |
| | PIP-Involved | 50% - 50% | 40% - 60% |

With the assumed market shares for the different scenarios, it becomes possible to estimate the value to the manufacturer of the different strategies. These values are summarized in Table 85 and described in a more appealing extensive tree representation in Figure 126.

Table 85: PIP development value in different competitive scenarios

| | | <i>Competition</i> | |
|--|--------------|--------------------|-------------|
| | | No Move | New PMA |
| <i>Original Equipment Manufacturer</i> | No Move | 0 / 0 | -315 / 19.7 |
| | PIP-Light | 19.8 / -88 | 4.8 / 10.8 |
| | PIP-Involved | 21.2 / -176 | 4.5 / 4.5 |

Once the profitability has been estimated for all scenarios in the sequential game, the process of finding the Nash equilibrium can start. Scenario 2 is the Nash equilibrium when the engine manufacturer is the first mover in the game and benefits from a head-start during which little competition is impacting profits. Indeed, in this particular scenario, no competitor has any incentive to deviate, and choosing a different course of action would not result in an equilibrium. Indeed, if the competitor were to change course of action, its profitability would decrease substantially (becoming negative). If the PIP manufacturer were to change strategy, the payoff of this new strategy would be

substantially less. Hence, it is more profitable for the PIP manufacturer to start investing in the PIP-Light early and forfeit the development of the more advanced PIP-Involved.

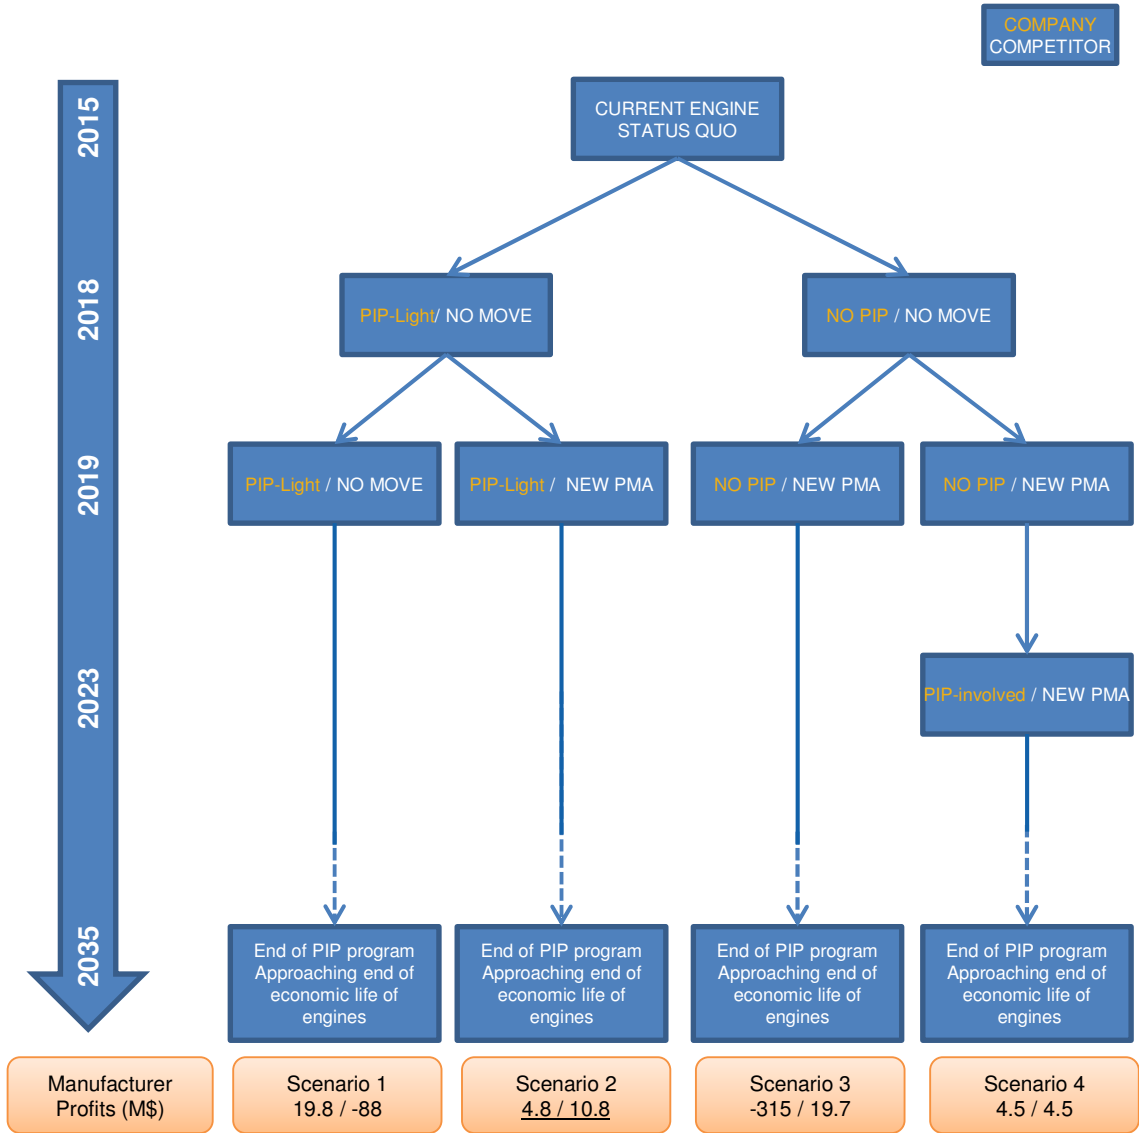


Figure 126: Sequential scenarios under investigation and selection of the Nash equilibrium

9.3 Back to overarching research question

his research question concerns the improvement of current state-of-the-art economic assessment methods required to address specific challenges associated with long-term and uncertainty-plagued aircraft and engine development programs evolving in a competitive environment.

Overarching Research Question – Improvement of value-based design methods

Within the context of aerospace research and development optimization, how can value-based design methodologies be improved to identify precursors of programmatic, technological, and market opportunities while reflecting the specific challenges associated with long-term and uncertainty-plagued aircraft and engine developments, and while accounting for the competitive nature of the business?

This overarching research question is linked to three hypotheses. The first hypothesis deals with the flexibility offered to decision-makers to take advantage of investment opportunities. The second hypothesis deals with a proposed improvement to current viability assessments by the introduction of competitive aspects early-on during the economic analysis of future concepts. The third hypothesis proposes a concurrent use of these two improvements to yield better evaluation of long-term and uncertain research and development programs with staggered investments.

Hypothesis 1 — Real options for valuation with flexibility and uncertainty

Within the context of aerospace research and development programs, real options methods enable the development of value-based design frameworks accounting for the

staggered nature of investments and the value created by managerial flexibility in uncertain environments.

Hypothesis 2 — Game theory for investigation of economic robustness with competition

Within the context of aerospace research and development programs, game theoretic methodologies enable transparent and traceable analyses that allow decision-makers to better investigate the economic robustness of selected technology and product development streams in a competitive environment characterized by uncertain moves by competitors.

Hypothesis 3 — Combined real options and game theoretic analyses

Real options methodologies combined with game theoretic methodologies allow the identification of windows of opportunities and yield analyses superior in term of robustness to either of these two analyses performed independently.

The first hypothesis claims that there is value created by active and flexible management of research and development programs. In addition, it claims that unlike traditional methods, a real option-based method is able to capture the value of this flexibility and therefore does not systematically undervalue research and development programs. This hypothesis has already been validated when hypothesis 1.1 and all lower-level hypotheses were verified and validated.

The second hypothesis claims that strategies identified with the use of a game theoretic approach are robust. To validate this hypothesis, it is sufficient to prove that competitors have no incentive to deviate from the equilibrium found with the game theoretic approach. This has already been verified and validated when hypothesis 2.1 was

validated. In fact, the construction of the equilibrium solution ensures that the strategy is robust with respect to moves by the competition.

The third hypothesis pertains to the combination of a real options methodology and a game theoretic methodology. The validation is done by comparing the research and development program value and the variability of this value in a pure game theoretic methodology, in a pure real options methodology, and in a combined methodology. The hypothesis has already been validated when the payoffs of the research and development program were estimated using real options analysis and used in the sequential competitive game.

CHAPTER 10: CONCLUSIONS AND CONTRIBUTIONS

The focus of this research endeavor is the analysis of staggered technology development programs in the aerospace industry. An extensive literature review indicated that the development of new technologies for commercial aviation involves significant risk for technologists as these programs are often driven by fixed assumptions regarding future airline needs, while being subject to an abundance of uncertainties at the technical and market levels. During the literature review, several characteristics typical of these developments were identified: the existence of milestones over the course of the development program at which the investment may be abandoned or delayed, the significant amount of uncertainties surrounding these developments both at the technical and market levels, and the long-term nature of these developments which compounds the effect of uncertainties over time.

10.1 Overarching research question

With the characteristics of technology developments programs stated above, three observations were made: (1) integrating competitive aspects early in the design ensures that development programs are robust with respect to moves by the competition and this helps mitigate some of the competitive uncertainty; (2) disregarding managerial flexibility undervalues many long-term and uncertain research and development programs; and (3) windows of opportunities emerge and disappear, and manufacturers could derive significant value by exploiting their upside potential.

In an environment with limited financial resources, decision-makers need to prioritize development programs so as to down-select only the most promising ones. This is usually done by first constructing a business case for each technology development stream, then by assessing their economic viability, and finally by down-selecting the most promising one. In order to perform this exercise, decision-makers need to equip themselves with the tools necessary to analyze these developments. Several techniques suitable for economic analyses and viability assessments have been reviewed, but most of them are found to be lacking and unable to capture the specificities of long-term developments in the aerospace industry. Consequently, the main objective of this research is to answer the following overarching research question: *“Within the context of aerospace research and development optimization, how can value-based design methodologies be improved to identify precursors of programmatic, technological, and market opportunities while reflecting the specific challenges associated with long-term and uncertainty-plagued aircraft and engine developments, and while accounting for the competitive nature of the business?”*

10.2 Method research questions and hypotheses

Drawing on the literature review, a set of three method hypotheses was formulated to answer the overarching research question. The first hypothesis relates to the improvement of value-based methodologies to handle long-term uncertain developments and to identify precursors of opportunities. The second hypothesis relates to the improvement of current viability assessments by the introduction of competitive aspects early-on during the economic analysis of future concepts. The third hypothesis

relates to the concurrent use of both of these methods to better evaluate the research and development programs.

Using the development of a performance improvement package for turbofan aircraft engines as a proof-of-concept study, it was shown that real options analyses do help improve the current state-of-the-art value-based methodologies by being able to model the usually complex and intertwined uncertainties surrounding technological developments, while accounting for the flexibility offered to management to react to the realization of these uncertainties. In particular, the real options methodology was able to quantify the value created by astute management in the presence of uncertainty by evaluating the “flexibility premium” gained by the performance improvement package manufacturer through proper use of the flexibility offered to decision-makers. This is one aspect that is traditionally absent from typical capital budgeting analyses.

Still using the performance improvement package proof-of-concept study, it was shown that using game theoretic analyses early-on during the technology development enables the formulation of robust strategies from which no competitor has any incentive to deviate. This was shown by proving that any other strategy would lead to sub-optimal profitability.

10.3 Modeling research questions and hypotheses

While performing these investigations, it soon became apparent that popular real option-based methods were insufficient to handle a complex reality. Their traditional domain of application was too narrow for realistic analyses while their implementation was too constrained for meaningful analyses. Besides, these popular real option-based

methods were not making use of the latest tools and techniques developed for financial engineering applications that turn out to be valuable for real options applications.

The first modeling hypothesis stated that Monte Carlo-based option pricing techniques were the most promising framework for the analysis of technology investments facing multiple and potentially correlated sources of uncertainty, the evolution of which may follow complex stochastic processes. Using the proof-of-concept study, this hypothesis was validated by subjecting the retrofit development to the jet-fuel price and carbon emission allowance price uncertainties, which are both correlated and which may be best described using jump diffusion processes.

The second modeling hypothesis stated that the use of path-dependent American options enables the capture of the flexibility offered to decision-makers to optimally time investments. Using the proof-of-concept study, this hypothesis was validated by first deriving the early-trigger boundary which helps decision-makers identify the proper set of conditions for which early-exercise is optimal, and then, by deriving the early-trigger premium which reflects the additional value gained by the manufacturer when allowing decision-makers to time their investment decisions optimally.

The third modeling hypothesis stated that the use of sequential competitive scenarios provided a means to measure the impact of competing designs and to identify robust strategies. Using the proof-of-concept study which featured three different technology development streams, the game theoretic analysis enabled the pruning of the strategy space to yield a single robust strategy from which none of the competitors had any incentive to deviate.

10.4 Technical research questions and hypotheses

While performing the real options investigations, it became apparent once again that popular Monte Carlo-based techniques required some modifications in order to be used efficiently in an environment riddled with uncertainties and where the evolution of the development program value over time is unobservable and therefore cannot be calibrated.

The first technical hypothesis claimed that using the non-parametric Esscher transform enables the change of probability measure required when the stochastic process representing the evolution of the development program is unknown and unobservable. This hypothesis was verified using canonical tests by comparing the transformed distributions induced by simulations under the physical probability measure and subsequent Esscher transformations to the equivalent martingale measure for some specific cases where the latter are known.

The second hypothesis claimed that a bootstrap procedure using the resampling wheel algorithm enables a resampling of the weighted distributions obtained from the Esscher transformation with the aim of generating non-weighted trajectories representing the evolution of the development program value under the equivalent martingale measure. This hypothesis was also verified using canonical tests by comparing the distribution of returns induced by these new resampled trajectories to the known equivalent martingale measure for some specific cases where the latter is known.

The third hypothesis suggests the use of the least-squares Monte Carlo algorithm to perform the valuation of real options featuring early-exercise possibilities and the construction of the trigger boundary. This hypothesis was partially verified using

canonical tests by comparing the price of real options as well as the location and shape of the trigger boundaries for specific cases where the value of the real option and the location of the trigger boundary can be estimated by other means.

The verification of this hypothesis highlighted the fact that, although the valuation was correct, the position and shape of the trigger boundary were not consistent across repeated trials. Another hypothesis was setup which claimed that the implementation of variance reduction techniques would improve the generation of trigger boundaries. It was verified that control variates, multi-start Monte Carlo simulations, and a restriction of the domain of regression of the continuation value do improve the generation of the early-exercise boundary. It was however proved that the use of quasi-Monte Carlo simulations using Sobol's low-discrepancy sequences was not yielding any benefit over regular Monte Carlo simulations.

All in all, a novel methodology, cross-fertilizing different techniques borrowed from the fields of quantitative finance, actuarial sciences, and statistics is proposed to answer the overarching research question and all derived sub-research questions. The novel method was applied to a PIP retrofit proof-of-concept study about the timing of staggered investments under uncertainty.

10.5 List of contributions

There are several contributions stemming from this research endeavor. Most of them fall under the general theme of the evaluation of staggered investments under uncertainty. A couple of contributions are related to the aircraft and engine evaluation method proposed as part of the proof-of-concept implementation. Finally, two

mathematical derivations dealing with properties of geometric Brownian motions are novel to the author's knowledge and are detailed in the appendix.

10.5.1 Contributions related to the analyses of staggered investments

On the real options side, the contributions are as follows:

- Use of the Esscher transform and its non-parametric approximation to perform a change of measure and to obtain risk-neutral distributions. To the author's knowledge, this has never been applied for the evaluation of real options.
- Use of a resampling wheel algorithm (bootstrap) to generate risk-neutral trajectories using a risk-neutral terminal distribution. To the author's knowledge, this technique has never been used to generate trajectories for the underlying process in real options applications.
- Use of least-squares Monte Carlo simulations to generate an early-investment boundary for real options applications. The least-squares Monte Carlo algorithm proposed by Longstaff and Schwartz has seen only few applications in the real options literature to the author's knowledge.
- Use of multi-start Monte Carlo simulations to improve the generation of the trigger boundary by facilitating the search for the critical price. To the author's knowledge, this technique has never been applied to improve the continuation value regression.
- Use of control variates sampled at exercise of path-dependent options to significantly improve the rate of convergence of the least-squares Monte Carlo algorithm proposed by Longstaff and Schwartz. To the author's knowledge, this technique has never been applied in real options applications.

- Use of the early-investment boundary to identify trigger events and precursors of successful research and development programs. By investigating the sensitivity of the early-investment boundary to the level of some uncertainties, events triggering the need for immediate investments can be detected.

10.5.2 Contributions related to the proof-of-concept

On the aircraft and engine evaluation side, the main contribution is as follows:

- Development of a maintenance model for turbofans powering short to medium range narrowbody aircraft. This maintenance model, constructed exclusively using public information, is able to estimate maintenance reserves depending on the type of operations and the type of environment the turbofan is operated in.

10.5.3 Other miscellaneous contributions

Two other contributions stem from the derivation of some properties related to geometric Brownian motions:

- Derivation of the probability that a geometric Brownian motion process hits a threshold during a given time-interval. This formula is then applied to estimate the probability that the process modeling the evolution of the price of jet-fuel hits a certain price-level by a certain date.
- Derivation of the expected time required for a geometric Brownian motion process to hit a threshold. This formula is then applied to estimate the average time required for the process modeling the evolution of the price of jet-fuel to hit a given price level.

10.6 Proposed extensions

This thesis dissertation develops a novel method to evaluate the economic viability of new technology and product developments facing significant market and competitive uncertainties and to detect and identify trigger events of successful developments. Two main avenues have been identified as promising for future research: more complete and thorough competitive analysis, and further extension of the domain of applicability of real options techniques.

Indeed, the competitive aspect has been only brushed over the course of this research and more advanced game theoretic and marketing analyses could be performed to enhance the forecasting power of the method. In the present research, only sequential and non-cooperative games with perfect information have been investigated. A natural way forward is the study of simultaneous games where the competitors have no knowledge of the development streams of their competitors. In this sense, this would be able to capture the surprise-effect due to closely-guarded development programs such as the one of the Boeing 787 at the turn of the century. Another natural way forward is the study of imperfect games where the different players have only partial information about the profitability of their competitors.

The domain of applicability of real options techniques could be further extended by looking at stochastic processes featuring non-stationary and non-independent increments. Indeed, the current methodology makes use of the Esscher transformation which requires that the logarithm of the underlying stochastic process features stationary and independent increments. This includes a wide variety of stochastic processes such as Wiener processes, Poisson processes, gamma processes, and inverse Gaussian processes.

However, this precludes two types of processes. The first is the set of mean-reverting processes which, by virtue of the mean reversion feature, do not have independent increments. These processes could prove useful to model the evolution over time of commodities for which a long-term return to an equilibrium level makes sense. The second is the set of heteroscedastic processes for which the volatility is changing over time. These processes could prove useful to model the periods of high volatility following significant perturbations in the economies.

APPENDIX A: GEOMETRIC BROWNIAN MOTION

Introduction and some properties

Brownian motion is named after botanist Robert Brown who, in 1827, observed the random movement of particles of pollen suspended in a fluid. It is only decades later that the transport phenomenon received the attention of Albert Einstein who worked on the modeling aspect as part of the *Annus Mirabilis* papers [285]. Brownian motion is a diffusion process and can be described in mathematics using the Wiener process which is characterized by the following properties [185]:

Definition A.1

Let W be a Wiener process. It is characterized by these four properties:

- (i) $W(0) = 0$
- (ii) $W(t)$ is almost surely continuous
- (iii) $W(t)$ has independent increment
- (iv) $W(t) - W(s) \sim N(0, t - s)$

In plain terms, a Wiener process is a stochastic process that starts at zero, with continuous sample paths, and with stationary and independent increments. A Wiener process is therefore a Markov process, which means that the probability distribution for all future values of the process depends only on its current value. It is thus unaffected by past values of the process or any other past information. A Wiener process has independent increments, which means that the probability distribution for the change in the process over any time interval is independent of any other non-overlapping time intervals.

Finally, the changes in the process over any finite interval of time are normally distributed with a variance that increases linearly with the time interval [286].

An arithmetic Brownian motion S' with parameter (S'_0, μ', σ') is defined in this dissertation as a Wiener process W with a drift term that measures the trend and a volatility term that measures the variability of the process over time as shown in Eq. 64:

$$S'(t) = S'_0 + \mu't + \sigma'W(t) \quad \text{Eq. 64}$$

Looking at this expression, an arithmetic Brownian motion can be understood to be the accumulation of independent and identically normally distributed increments over time. In infinitesimal terms, the arithmetic Brownian motion has infinitesimal random increments dS' over the infinitesimal time dt with mean $\mu'dt$ and variance σ'^2dt . This leads to the following stochastic differential equation shown in Eq. 65:

$$dS' = \mu'dt + \sigma'dW \quad \text{Eq. 65}$$

Such a model specification does not preclude S' from taking on negative values which is not consistent with the limited liability concept for asset prices. The always positive geometric Brownian motion has been introduced as the exponential of the arithmetic Brownian motion and was found particularly helpful in economics. It provides a first approximation of the dynamics of exchange rates, natural resources, and more generally many asset prices [60]. Using Ito's lemma, a stochastic differential equation for the exponential of an arithmetic Brownian motion is found and displayed in Eq. 66:

$$\text{Let: } S = e^{S'} \text{ then } dS = \left(\mu'e^{S'} + \frac{1}{2}e^{S'}\sigma'^2 \right) dt + e^{S'}\sigma'dW \quad \text{Eq. 66}$$

The terms may be reorganized by separating the deterministic (drift) and the stochastic parts. This yields Eq. 67. Performing a change of variables yields Eq. 68 which is the stochastic differential equation for the geometric Brownian motion with parameters (μ, σ) :

$$\frac{dS}{S} = \left(\mu' + \frac{1}{2} \sigma'^2 \right) dt + \sigma' dW \quad \text{Eq. 67}$$

$$\frac{dS}{S} = \mu dt + \sigma dW \text{ with change of variable } \mu = \mu' + \frac{1}{2} \sigma'^2 \text{ and } \sigma = \sigma' \quad \text{Eq. 68}$$

Investigations of this formula reveal several interesting properties. First and by construction, the geometric Brownian motion is never negative. Then, the magnitude of increments is directly related to the current realization of the process itself. This means that the innovation is relative, a property which is useful when modeling the dynamics of price as well as many other natural processes.

Estimation of parameters

In this paragraph, the estimation of parameters for the geometric Brownian motion is discussed. Even though the geometric Brownian motion is one of the most fundamental stochastic processes used in finance, estimating its drift and volatility is in fact not trivial. Drift estimation is presented for the sake of completeness as this is not usually helpful for derivative pricing, while the estimation of volatility plays a much more important role.

Drift estimation

Estimating the drift of a geometric Brownian motion is in fact quite difficult and this will be shown in this paragraph. In fact, this assertion makes perfect sense when the geometric Brownian motion is used to model prices. Indeed, if economic agents were to know

exactly the future expected price, the current price would be adjusted accordingly and the drift rate would change [287]. The definition of the geometric Brownian motion with constant drift and constant volatility is shown in Eq. 69:

$$S(t) = S_0 e^{\left(\mu - \frac{\sigma^2}{2}\right)t + \sigma W(t)} \text{ with } W(t) \sim N(0, \sqrt{t}) \quad \text{Eq. 69}$$

Let's now take a sample interval between T_1 and T_2 . Let's subdivide this interval into m different subintervals $[t_i, t_{i+1}]$ forming therefore a partition of the original interval and representing each time at which the asset price is sampled. The observed continuous return is given in Eq. 70:

$$\ln\left(\frac{S(t_{i+1})}{S(t_i)}\right) = \left(\mu - \frac{\sigma^2}{2}\right)(t_{i+1} - t_i) + \sigma(W(t_{i+1}) - W(t_i)) \quad \text{Eq. 70}$$

Summing these returns over the entire interval yields Eq. 71:

$$\sum_{i=0}^{m-1} \ln\left(\frac{S(t_{i+1})}{S(t_i)}\right) = \left(\mu - \frac{\sigma^2}{2}\right) \sum_{i=0}^{m-1} (t_{i+1} - t_i) + \sigma \sum_{i=0}^{m-1} W(t_{i+1}) - W(t_i) \quad \text{Eq. 71}$$

Rearranging this expression by observing that the first sum is simply the overall continuous return while the second sum is the interval length, and finally taking both the expectation and variance yields Eq. 72 and Eq. 73.

$$E\left[\ln\left(\frac{S(T_2)}{S(T_1)}\right)\right] = \left(\mu - \frac{\sigma^2}{2}\right)(T_2 - T_1) + \sigma E[W(T_2) - W(T_1)] \quad \text{Eq. 72}$$

$$\text{Var}\left[\ln\left(\frac{S(T_2)}{S(T_1)}\right)\right] = \left(\mu - \frac{\sigma^2}{2}\right)(T_2 - T_1) + \sigma^2 \text{Var}[W(T_2) - W(T_1)] \quad \text{Eq. 73}$$

The last term $W(T_2) - W(T_1)$ is a Wiener increment. By definition, it is normally distributed $N(0, T_2 - T_1)$ and its expectation is null. A non-biased estimator of the drift and the variance of this estimator are given in Eq. 74:

$$\hat{\mu} = \frac{1}{T_2 - T_1} \ln \left(\frac{S(T_2)}{S(T_1)} \right) + \frac{\sigma^2}{2} \text{ with } Var(\hat{\mu}) \approx \frac{\sigma^2}{T_2 - T_1} \quad \text{Eq. 74}$$

As may be seen in the previous equation, the drift term can be estimated by computing the continuous return between the first and final observation in the sample. The variance of the estimator is only a function of the interval length. There is therefore no benefit in having many observations spaced closely together as only increasing the interval length can improve the accuracy of the drift term estimator. This is the main difficulty in assessing the drift term: shorter intervals yield bad estimates while longer intervals are more accurate but may no longer be relevant and may exhibit changes in the supposedly constant drift term.

Volatility estimation

With constant drift and constant volatility, the geometric Brownian motion may be written as:

$$S(t) = S_0 e^{\left(\mu - \frac{\sigma^2}{2}\right)t + \sigma W(t)} \text{ with } W(t) \sim N(0, \sqrt{t}) \quad \text{Eq. 75}$$

Let's now take a sample interval between T_1 and T_2 and subdivide this interval in m different subintervals $[t_i, t_{i+1}]$ forming a partition of the interval and representing each time at which the asset price is sampled. The observed return of the asset price is therefore:

$$\ln\left(\frac{S(t_{i+1})}{S(t_i)}\right) = \left(\mu - \frac{\sigma^2}{2}\right)(t_{i+1} - t_i) + \sigma(W(t_{i+1}) - W(t_i)) \quad \text{Eq. 76}$$

Summing the square of these returns over the entire interval yields the following formula:

$$\begin{aligned} \sum_{i=0}^{m-1} \left[\ln\left(\frac{S(t_{i+1})}{S(t_i)}\right) \right]^2 &= \left(\mu - \frac{\sigma^2}{2}\right)^2 \sum_{i=0}^{m-1} (t_{i+1} - t_i)^2 \\ &+ \sigma^2 \sum_{i=0}^{m-1} (W(t_{i+1}) - W(t_i))^2 \\ &+ 2\sigma \left(\mu - \frac{\sigma^2}{2}\right) \sum_{i=0}^{m-1} (W(t_{i+1}) - W(t_i))(t_{i+1} - t_i) \end{aligned} \quad \text{Eq. 77}$$

Looking at the right-hand side of Eq. 77, there are three major terms. With a step size sufficiently small, the first term converges to its limit of zero. Using the properties of the quadratic variation of the Brownian motion, the second term has a limit equal to $\sigma^2(T_2 - T_1)$. Finally, the last term converges to zero as the step size gets sufficiently small. Consequently, an estimator of the volatility of the stochastic process can be computed from a time series using the formula in Eq. 78.

$$\widehat{\sigma^2} = \frac{1}{T_2 - T_1} \sum_{i=0}^{m-1} \left[\ln\left(\frac{S(t_{i+1})}{S(t_i)}\right) \right]^2 \quad \text{Eq. 78}$$

Drift and volatility estimations using maximum likelihood estimators

Another method to estimate both the drift and the volatility of a geometric Brownian motion at the same time is to use the maximum likelihood estimators of these two

quantities. Campbell et al. [280] developed the formula for the maximum likelihood function shown in Eq. 79:

$$\mathcal{L}(\mu, \sigma) = -\frac{m}{2} \ln(2\pi\sigma^2 h) - \frac{1}{2\sigma^2 h} \sum_{i=0}^{m-1} \left(\ln \left(\frac{S(t_{i+1})}{S(t_i)} \right) - \left(\mu - \frac{\sigma^2}{2} \right) h \right)^2 \quad \text{Eq. 79}$$

Differentiating this expression with regards to μ and σ^2 yields the following estimators of the drift and volatility:

$$\hat{\mu} = \frac{\hat{\sigma}^2}{2} + \frac{1}{T_2 - T_1} \sum_{i=0}^{m-1} \ln \left(\frac{S(t_{i+1})}{S(t_i)} \right) \quad \text{with} \quad \text{Var}(\hat{\mu}) \approx \frac{\sigma^2}{T_2 - T_1} \quad \text{Eq. 80}$$

$$\hat{\sigma}^2 = \frac{1}{T_2 - T_1} \sum_{i=0}^{m-1} \left(\ln \left(\frac{S(t_{i+1})}{S(t_i)} \right) - \left(\hat{\mu} - \frac{\hat{\sigma}^2}{2} \right) h \right)^2 \quad \text{with} \quad \text{Var}(\hat{\sigma}^2) \approx \frac{2\sigma^4}{m} \quad \text{Eq. 81}$$

Numerical application

For the example described in Chapter 2.1, the jet-fuel price time series from the United States Energy Information Administration [34] is used to calibrate a geometric Brownian motion. The data ranges from June 1994 to November 2012 but it is truncated after May 1997 to reflect the data available to the aircraft manufacturer at launch. The parameterization of the geometric Brownian motion is done using the maximum likelihood estimators and the results are provided in Table 86.

Table 86: U.S. Gulf Coast Kerosene-type jet-fuel spot price – June 1994 - May 1997

| U.S. Gulf Coast Kerosene-type jet fuel | |
|---|-------|
| Yearly Drift | 9.88% |
| Yearly Volatility | 33.3% |

Probability of exceeding a threshold at one point in time

The objective of this paragraph is to demonstrate how to estimate the probability that a process following a geometric Brownian motion and having a current value S_0 exceeds a given value S_t at a certain time t . By definition, the geometric Brownian motion representing the process S may be written as in Eq. 75 and the ratio expressed in Eq. 82 is distributed as shown:

$$\ln\left(\frac{S(t)}{S_0}\right) = \left(\mu - \frac{\sigma^2}{2}\right)t + \sigma W(t) \sim N\left(\left(\mu - \frac{\sigma^2}{2}\right)t, \sigma\sqrt{t}\right) \quad \text{Eq. 82}$$

This immediately yields the probability that the stock price exceeds a certain threshold at a given time expressed in Eq. 83:

$$\Pr(S(t) > S_t) = 1 - \Phi\left(\frac{\ln\left(\frac{S_t}{S_0}\right) - \left(\mu - \frac{\sigma^2}{2}\right)t}{\sigma\sqrt{t}}\right) \quad \text{Eq. 83}$$

Numerical application

Using only information available at the time of the commercial launch of the A340-500 and A340-600 aircraft (in 1997), the probability of the jet-fuel price reaching the level it did at these aircraft entry into service (in 2002), and the probability of the jet-fuel price reaching the level it did by the time of the last delivery (in 2012) are given in Table 87. Looking at these results, it was not unlikely for the spot price of jet-fuel to surge drastically as it did during the development and production of the aircraft.

Table 87: Probability of jet-fuel exceeding certain thresholds at certain dates

| | Jet-fuel reaching US\$0.67 at aircraft entry into service (2002) | Jet-fuel reaching US\$2.96 at aircraft last delivery (2012) |
|--------------------|---|--|
| Probability | 48% | 21% |

Probability of hitting a threshold during a time interval

The objective of this paragraph is to demonstrate how to estimate the probability that a process following a geometric Brownian motion exceeds a given price before a given time. This is an interesting notion to study because in many engineering problems, there are thresholds for some metrics (jet-fuel price for instance) above which some design choices may be altered (requirement to infuse fuel-saving technologies in an aircraft design for instance). Even if temporary, a stochastic process hitting a specific threshold at one point in time may influence the behavior of economic agents. Estimating the probability of these occurrences is consequently helpful for designers to perform robustness analyses.

Proposition A.1

Let S be a geometric Brownian motion with initial state S_0 , drift μ and volatility σ . Let $T \in \mathbb{R}^{+}$ be a strictly positive real number defining the interval $[0, T]$. Let M_T be the maximum of S over this interval. Let $S_1 \in \mathbb{R}^{+*}$ be strictly positive real number defining a barrier above S_0 . The probability that S exceeds S_1 at one point by time T is given by:*

$$F(T, S_0, S_1) = Pr(M_t \geq S_1)$$

$$F(T, S_0, S_1) = \phi \left(\frac{\ln \left(\frac{S_0}{S_1} \right) + \left(\mu - \frac{\sigma^2}{2} \right) T}{\sigma \sqrt{T}} \right) \\ + \exp \left(\frac{2 \left(\mu - \frac{\sigma^2}{2} \right) \ln \left(\frac{S_1}{S_0} \right)}{\sigma^2} \right) \phi \left(\frac{\ln \left(\frac{S_0}{S_1} \right) - \left(\mu - \frac{\sigma^2}{2} \right) T}{\sigma \sqrt{T}} \right)$$

Proof: In order to estimate this probability, a small detour in the world of arithmetic Brownian motion is warranted to simplify computations. By definition, the arithmetic Brownian motion follows the stochastic differential equation given in Eq. 84:

$$dS'_t = \mu' dt + \sigma' W(t) \quad \text{Eq. 84}$$

Using this arithmetic Brownian motion, Harrison [185] states that the cumulative density function for an arithmetic Brownian motion starting at zero, having a maximum M'_t below a barrier S'_3 and ending at time T below another barrier S'_2 is given in Eq. 85:

$$F_T(S'_2, S'_3) = P(S'_t \leq S'_2, M'_t \leq S'_3) \quad \text{with} \quad M'_t = \sup(S'_t, 0 \leq t \leq T) \\ F_T(S'_2, S'_3) = \phi \left(\frac{S'_2 - \mu' t}{\sigma' \sqrt{t}} \right) - \exp \left(\frac{2\mu' S'_3}{\sigma'^2} \right) \phi \left(\frac{S'_2 - 2S'_3 - \mu' t}{\sigma' \sqrt{t}} \right) \quad \text{Eq. 85}$$

This result is the cumulative density function for an arithmetic Brownian motion not hitting two distinct barriers: one barrier S'_2 concerns the final value of the process at end-time, while the other barrier S'_3 is related to the maximum of the process and is thus related to the entire process between the initial time and the end time. To yield results concerning only the second barrier S'_3 related to the maximum of the process, the first barrier S'_2 is set equal to S'_3 . Similarly, to account for a Brownian motion that does not

start at zero but rather at an initial value S'_0 , a change of variable $S'_1 = S'_3 + S'_0$ is performed and yields Eq. 86:

$$F(T, S'_0, S'_1) = F_T(S'_3, S'_3) = F_T(S'_1 - S'_0, S'_1 - S'_0)$$

$$F(T, S'_0, S'_1) = \phi\left(\frac{S'_1 - S'_0 - \mu't}{\sigma'\sqrt{t}}\right) - \exp\left(\frac{2\mu'(S'_1 - S'_0)}{\sigma'^2}\right) \phi\left(\frac{S'_0 - S'_1 - \mu't}{\sigma'\sqrt{t}}\right) \quad \text{Eq. 86}$$

Finally, of interest is the probability of hitting the barrier S_1 which is exactly the complement of the probability of the maximum being below this barrier. This yields the formula for F_{BM} , the first hitting time of an arithmetic Brownian motion, in Eq. 87:

$$F_{BM}(T, S'_0, S'_1) = \phi\left(\frac{S'_0 - S'_1 + \mu't}{\sigma'\sqrt{t}}\right) + \exp\left(\frac{2\mu'(S'_1 - S'_0)}{\sigma'^2}\right) \phi\left(\frac{S'_0 - S'_1 - \mu't}{\sigma'\sqrt{t}}\right) \quad \text{Eq. 87}$$

Having this cumulative density function for the arithmetic Brownian motion, the next step is to establish the formula for the geometric Brownian motion. This is done by means of another change of variable: if S follows a geometric Brownian motion with mean μ and standard deviation σ , then $S' = \ln(S)$ follows an arithmetic Brownian motion with mean $\mu - \sigma^2/2$ and standard deviation σ . Using Ito's lemma, the stochastic differential equation for the logarithm of the original process is indeed given by Eq. 88:

$$S' = \ln(S) \text{ then } dS' = d(\ln(S)) = \left(\mu - \frac{\sigma^2}{2}\right)dt + \sigma dW = \mu' dt + \sigma' dW \quad \text{Eq. 88}$$

Using this change of variable, plugging the new value of the drift term shown in Eq. 88, and rearranging the cumulative distribution function given in Eq. 87 yield the expression for the first hitting time of a geometric Brownian motion in Eq. 89:

$$F_{GBM}(T, S_0, S_1) = \phi \left(\frac{\ln \left(\frac{S_0}{S_1} \right) + \left(\mu - \frac{\sigma^2}{2} \right) t}{\sigma \sqrt{t}} \right) + \exp \left(\frac{2 \left(\mu - \frac{\sigma^2}{2} \right) \ln \left(\frac{S_1}{S_0} \right)}{\sigma^2} \right) \phi \left(\frac{\ln \left(\frac{S_0}{S_1} \right) - \left(\mu - \frac{\sigma^2}{2} \right) t}{\sigma \sqrt{t}} \right) \quad \text{Eq. 89}$$

Numerical application

For the example described in Chapter 2.1, the jet-fuel price time series is modeled as a geometric Brownian motion with the parameterization given in Table 86. The probability of the jet-fuel price exceeding the price it actually had at the A340-500/600 entry into service (in 2002) by the entry into service time, as well as the probability of the jet-fuel price exceeding the price it actually had at the A340-500/600 last delivery (in 2012) by the time the last aircraft was delivered are given in Table 88:

Table 88: Probability of jet-fuel price hitting certain thresholds by certain dates

| | Jet-fuel hitting US\$0.67 by aircraft entry into service in 2002 | Jet-fuel hitting US\$2.96 by aircraft last delivery in 2012 |
|--------------------|---|--|
| Probability | 72% | 34% |

Expected time to hit a threshold during a time-interval

On a related subject, another metric of interest could be the expected time for a geometric Brownian motion to first hit a given threshold. This yields a time-estimate that may be used by economic agents in order to get prepared for some specific scenarios. Dixit [288] provides a formula in Eq. 90 for the expected time T' it takes for an arithmetic Brownian motion S' with drift μ' to first hit the barrier at S'_1 from below:

$$E(T_{S'=S'_1}) = \begin{cases} \frac{S'_1 - S'_0}{\mu'} & \text{if } \mu' > 0 \\ +\infty & \text{if } \mu' \leq 0 \end{cases} \quad \text{Eq. 90}$$

Using the usual logarithmic transformation from arithmetic Brownian motion to geometric Brownian motion, the drift is adjusted accordingly, and this yields Eq. 91:

$$E(T_{S=S_1}) = \begin{cases} \frac{\ln\left(\frac{S_1}{S_0}\right)}{\mu - \frac{\sigma^2}{2}} & \text{if } \mu > \frac{\sigma^2}{2} \\ +\infty & \text{if } \mu \leq \frac{\sigma^2}{2} \end{cases} \quad \text{Eq. 91}$$

Numerical application

For the example described in Chapter 2.1 and the jet-fuel price time series parameterized as in Table 86, some estimates for the expected times for the fuel price to reach the level it did at the A340-500/600 entry into service and final delivery are given in Table 89.

Table 89: Expected time for the jet-fuel price hitting certain thresholds

| | Jet-fuel hitting US\$0.67 by aircraft entry into service | Jet-fuel hitting US\$2.96 by aircraft last delivery |
|----------------------|---|--|
| Expected Time | 5.8 years | 40 years |

APPENDIX B: JUMP DIFFUSION PROCESSES

Introduction

Geometric Brownian motion is without doubt the most widely studied stochastic process used to model the evolution of prices. Two properties of the geometric Brownian motion are discussed in this paragraph: time-scale invariance and continuity. Time-scale invariance means that the statistical properties of the process are similar at a yearly scale, a monthly scale, a daily scale, or an hourly scale. In reality, evidence tends to support the contrary: price evolutions are jaggy but continuous at the longer-horizon time-scales whereas they become highly discontinuous at shorter-horizon time-scales. Similarly, diffusion processes cannot generate sudden discontinuities in prices such as those observed when market shocks occur. In a diffusion process, large moves are the result of the accumulation of many small moves over time, and this cannot be used to easily model the sudden moves that are observed in the market [289]. Indeed, fine-tuning the calibration to account for these events results in a volatility that is too large during quiet periods in between sudden moves, while disregarding these sudden moves results in a volatility that underestimates the frequency of these events. In fact, recent events that led to large and sudden shocks include the 1979 oil crisis, the 1987 “Black Monday”, and the 2010 “Flash Crash”, and these are more frequent than the fast decaying tails of normal distributions imply.

Furthermore, the analysis of traded options has led researchers to some striking observations regarding the implied volatility and its shape. In the Black-Scholes model, the assumptions of stationary process and normally distributed returns should ensure that

the theoretical value of a vanilla option is a strictly monotonic and increasing function of volatility of the underlying asset. Therefore, a bijection exists between volatility and option price meaning that it is possible to derive a single value for the volatility – the implied volatility – from the observed price of options. Plotting this implied volatility for options with different strike prices and different maturities result in a surface called the implied volatility surface. A maturity cross-section of this implied volatility surface should be a straight horizontal line. Empirical findings are however different and exhibit a curved sloping line called the “volatility smile” [290] [202]. This means that deep out-of-the-money and deep in-the-money options are priced higher than the Black-Scholes model suggests which, in turn, implies possible misspecifications in the model.

Properties of jump-diffusion processes

One solution to capture the excess kurtosis often observed in the log-price densities is to extend the geometric Brownian motion by adding a compound Poisson¹ jump process to the diffusion process. The resulting process, called a jump diffusion process, belongs to the class of Levy processes and can be described as shown in Eq. 92, where μ is the drift of the diffusion process, σ is the volatility of the diffusion process, W_t is the Wiener process, $(N_t)_{t \geq 0}$ is the Poisson process counting the jumps, λ is the arrival intensity of the Poisson process, and finally, Y_i is the series representing the independent and identically distributed jump sizes.

$$S(t) = S_0 e^{X_t}, \text{ with } X_t = \mu t + \sigma W_t + \sum_{i=1}^{N_t} Y_i \quad \text{Eq. 92}$$

¹ A Poisson process $(N_t)_{t \geq 0}$ is a stochastic process that counts the number of events T_n in the time interval between 0 and t. The time between these events, $(T_n - T_{n-1})_{n \geq 1}$, is an independent and identically distributed sequence of exponential variables.

To fully specify the model, the jump size distribution denoted v , must also be defined. Various models have been proposed including jumps of deterministic size and jumps of random size [134][135][291]. A popular model is the 1976 jump-diffusion model pioneered by Merton [134], where the jumps in the log-price are assumed to be normally distributed with mean γ and volatility δ . This results in a particularly tractable model which is represented in Eq. 93:

$$S(t) = S_0 e^{X_t}, \text{ with } X_t = \mu t + \sigma W_t + \sum_{i=1}^{N_t} Y_i \quad \text{Eq. 93}$$

$$N_t \sim P(\lambda t) \text{ and } Y_i \sim N(\gamma, \delta)$$

Using Merton's model, some stylized results can be derived. Indeed, because the jump sizes are normally distributed, their sum is also normally distributed. Consequently, the sum in Eq. 93 (conditional on the occurrence of n jumps) can be written as a single normal distribution with properties shown in Eq. 94:

$$\sum_{i=1}^n Y_i \sim N(n\gamma, \delta\sqrt{n}) \quad \text{Eq. 94}$$

With the conditional distribution of the log-price expressed as the sum of two normal distributions, it becomes possible to further simplify and use a single normal distribution as highlighted in Eq. 95.

$$\ln\left(\frac{S(t)}{S_0}\right) \sim N\left(\left(\mu - \frac{\sigma^2}{2}\right)t, \sigma\sqrt{t}\right) + N(N_t\gamma, \delta\sqrt{N_t}) \quad \text{Eq. 95}$$

$$\ln\left(\frac{S(t)}{S_0}\right) \sim N\left(\left(\mu - \frac{\sigma^2}{2}\right)t + N_t\gamma, \sqrt{\sigma^2 t + \delta^2 N_t}\right)$$

Now, weighting this conditional distribution by the probability of having one, two etc. jumps lead to the following unconditional distribution in Eq. 96:

$$P(S(t) \leq S) = \sum_{n=0}^{+\infty} e^{-\lambda t} \frac{(\lambda t)^n}{n!} F_{n,t}(S), \quad \text{with } F_{n,t} \text{ the lognormal cdf} \quad \text{Eq. 96}$$

In turn, this can be used to express the price of a vanilla option as a Poisson-weighted infinite sum of Black-Scholes option prices, denoted V^{BS} . For instance, the formula for a call option on an asset following the Merton jump-diffusion process, V^{JD} , is provided in Eq. 97:

$$V^{JD}(S(t), t) = e^{-r \cdot t} \sum_{n=0}^{+\infty} e^{-\lambda(T-t)} \frac{(\lambda(T-t))^n}{n!} V^{BS}(S_n, t, \sigma_n)$$

$$\text{With } \sigma_n = \sqrt{\sigma^2 + n \frac{\delta^2}{T-t}} \text{ and } S_n = S_0 e^{\left(n\gamma + \frac{n\delta^2}{2} - \lambda(T-t) e^{\left(\gamma + \frac{\delta^2}{2} \right)} + \lambda(T-t) \right)} \quad \text{Eq. 97}$$

Detection of jumps estimation of Merton jump diffusion parameters

The detection of jumps as well as the calibration of jump diffusion processes is a notoriously difficult task. The main reason is the confounding of the continuous-time part and the jump part of stochastic processes reported in discrete observations [292]. The time-smoothing effect of less frequently sampled observations exacerbates the difficulty of detecting sudden jumps because jumps then tend to get averaged out [135]. Consequently, a significant amount of research is presently being undertaken in the field of empirical finance to provide analysts with methods and techniques to detect jumps and calibrate jump diffusion models [289][293][294][295]. Ait-Sahalia [135] uses maximum likelihood estimations to study the effect of the presence of jumps on the ability to

properly identify the volatility of log-return processes. Lee and Mykland [292] use a nonparametric test to detect the presence of jumps in the log-return time series and to estimate jump intensity and jump size distribution. Cont and Tankov [296] propose an entropy-minimization scheme to calibrate jump diffusion models using observed option prices. More recently, Tankov and Voltchkova [290] use observed option prices to calibrate jump diffusion processes by minimizing the squared norm of the difference between market and model prices. Going further into details is beyond the scope of this research but a few techniques and key points are highlighted in the following paragraphs. Interested readers are referred to the research work cited for further information.

Detection of jumps

Several methods have been proposed to detect the presence of jumps. The procedure proposed by Lee and Mykland [292] is retained in this research for two reasons. First, this is a non-parametric jump test and therefore the test is robust with respect to model specification. Second, the authors demonstrate that their test outperforms competing non-parametric jump tests such as that of Barndorff-Nielsen and Shephard [293] and that of Jiang and Oomen [297]. The intuition behind the test of Lee and Mykland is that if the volatility is high, then the occurrence of a jump in the market (leading to an abnormally high or low return) may not be distinguishable because the returns stemming from usual continuous innovations may be just as high when observed in discrete times. Therefore, looking at the magnitude of returns is not sufficient. Instead, to enable the detection of jumps, the returns are standardized by the instantaneous volatility which is a measure explaining the local variation due to the continuous part of the process. To provide a test suitable for non-stationary processes, the *instantaneous* volatility is used. It is estimated

using a rolling window, immediately preceding the time at which a jump detection test is performed.

The instantaneous volatility estimator is not computed using the realized quadratic power variation (traditional method to estimate the realized volatility of a geometric Brownian motion) since this estimator is inconsistent in the presence of jumps. Instead, the realized bipower variation defined in Eq. 98 as the sum of products of consecutive absolute returns is used as a nonparametric estimator of the volatility.

$$\widehat{\sigma(t_i)}^2 = \lim_{n \rightarrow \infty} \sum_{i=3}^n \left| \ln \left(\frac{S(t_i)}{S(t_{i-1})} \right) \right| \left| \ln \left(\frac{S(t_{i-1})}{S(t_{i-2})} \right) \right| \quad \text{Eq. 98}$$

With the definition of the instantaneous volatility, the test statistic $\mathcal{L}(i)$ which tests at time t_i whether a jump occurred from t_{i-1} to t_i is given in Eq. 99. It uses a rolling window of $K-1$ observations immediately preceding t_i to estimate the instantaneous volatility.

$$\mathcal{L}(i) = \frac{\ln \left(\frac{S(t_i)}{S(t_{i-1})} \right)}{\widehat{\sigma(t_i)}} \quad \text{Eq. 99}$$

$$\widehat{\sigma(t_i)}^2 = \frac{1}{K-2} \lim_{n \rightarrow \infty} \sum_{j=i-K+2}^{i-1} \left| \ln \left(\frac{S(t_j)}{S(t_{j-1})} \right) \right| \left| \ln \left(\frac{S(t_{j-1})}{S(t_{j-2})} \right) \right|$$

The asymptotic behavior of the jump detection statistic $\mathcal{L}(i)$ is then given by Eq. 100 (without jumps) and Eq. 111 (with jumps) under some mild conditions concerning the size of the rolling window ($K = O(\Delta t^\alpha)$ where $-1 < \alpha < -0.5$):

$$\left\{ \begin{array}{l} \text{No jump: } \mathcal{L}(i) \sim N \left(0, \frac{\pi}{2} \right) \quad \text{Eq. 100} \\ \text{One jump } Y \text{ at time } \tau: \mathcal{L}(i) \sim N \left(\frac{\sqrt{\pi} \cdot Y(\tau)}{\sqrt{2} \cdot \sigma \cdot \sqrt{\Delta t}}, \frac{\pi}{2} \right) \text{ and } \lim_{\Delta t \rightarrow 0} \mathcal{L}(i) = \infty \quad \text{Eq. 101} \end{array} \right.$$

On one hand, Eq. 100 indicates that without the presence of jumps, the test statistic follows a normal distribution with zero mean and a variance equals to $\pi/2$. On the other hand, Eq. 101 indicates that whenever a jump of amplitude Y occurs at a time τ between t_{i-1} and t_i , then the test statistic follows a normal distribution with non-zero mean and variance $\pi/2$. Of interest is the fact that the jump detection statistic exhibits a very different behavior when the frequency of observations increases: $\mathcal{L}(i) \rightarrow \infty$ as $\Delta t \rightarrow 0$.

The final step consists in selecting the rejection region for the test statistic. Because the test statistic exhibits very different behavior in the absence and presence of jumps, it may be possible to reject the absence of jump hypothesis. Indeed, in the absence of jump, the test statistic follows a normal distribution centered on zero, while in the presence of jumps it follows a non-centered normal distribution with a mean that may become very large. Consequently, the question becomes how large the test statistic can be without any jump. Under the absence of jump hypothesis, the asymptotic distribution of maximums of the test statistic is given by Eq. 102, where n is the number of observations:

$$\frac{\max|\mathcal{L}(i)| - C_n}{S_n} \rightarrow \xi \text{ with } P(\xi \leq x) = \exp(-e^{-x})$$

$$\text{With: } C_n = \sqrt{\pi \cdot \ln n} - \frac{\ln \pi + \ln(\ln n)}{\sqrt{\frac{16 \ln n}{\pi}}}, \quad S_n = \frac{\sqrt{\pi}}{\sqrt{4 \cdot \ln n}} \quad \text{Eq. 102}$$

Calibration of jump intensity

In the jump diffusion process of Merton, the occurrence of jumps is governed by a Poisson counting process of intensity λ and the size of jumps is governed by a normal

distribution of mean γ and volatility δ . As previously mentioned, jumps are used to model the higher-than-expected occurrence of abnormally large returns (leptokurtic property) which cannot be explained with a pure diffusive process. Changing the drift of the return process (increase or decrease) is not the intent of using jumps. Consequently, the combination (γ, δ) is often selected such that it does not impact the drift of the return process. This leads in turn to a model with fewer free parameters as shown in Eq. 103:

$$(\gamma, \delta) \text{ such that: } \gamma = -\frac{1}{2}\delta^2 \quad \text{Eq. 103}$$

Of more interest is the jump arrival intensity λ which may be directly estimated using the non-parametric jump test presented above. Indeed, since the jump detection test enables the detection of a jump within an interval t_{i-1} to t_i , it becomes possible to estimate the mean number of jumps within any time interval and in particular the mean number of jumps $\hat{\lambda}$ within the unit interval. Lee and Mykland show that this estimator $\hat{\lambda}$ converges to the actual intensity of the Poisson counting process λ .

APPENDIX C: PROBABILITY MEASURE

The concepts of probability measure, equivalent probability measure, and change of probability measure are used extensively throughout this thesis and more generally for asset pricing in mathematical finance. In order not to clutter the main text with unnecessary mathematical definitions, these concepts are described in the following paragraph in greater details. This enables the reader to have access to all relevant information to understand the work presented while not being distracted by abstract concepts.

Introduction

Before digging any further, let's first introduce some basic building blocks that are required to understand some of the aforementioned concepts. For many financial applications, the universe of all possible outcomes denoted by Ω refers to the sequence of asset prices over time. This space is both infinite since assets may take an infinite number of different prices, but also uncountable as it is not possible to enumerate all the prices and sequences of prices that may take place. Infinite and uncountable spaces present many challenges. The first is the inability to simply sum the probabilities of all elements within a subset to yield the probability of the entire subset. Another challenge is the non-tractability, as it is impossible to describe each and every subset to later assign probabilities.

To reduce this complexity, mathematicians introduced the notion of sigma algebra or sigma field (often denoted as σ -algebra) as a collection of subsets of interest for which a probability measure is defined. This collection of subsets is designed to ensure that only

events of interest are included and it is constructed to ensure that the usual probability operations (probability of a union and probability of a complement) can be defined in a consistent way for any subset within this σ -algebra. It is therefore defined to ensure measurability of each subset contained within the collection.

Definition C.1

Let Ω be a non-empty set and let \mathcal{F} be a collection of subsets of Ω . \mathcal{F} is a σ -algebra provided that:

- (i) The empty set \emptyset belongs to \mathcal{F} ,
- (ii) Whenever a set A belongs to \mathcal{F} , its complement A^c also belongs to \mathcal{F}
- (iii) Whenever a countable and possibly infinite sequence of sets A_1, A_2, \dots belongs to \mathcal{F} , then their union $\bigcup_{i=1}^{\infty} A_i$ also belongs to \mathcal{F}

In a less abstract way, if S denotes a continuous-time stochastic process (such as an asset price), then the information set which contains the observations of S up until time t constitutes a σ -algebra denoted by \mathcal{F}_t . It is called the sigma algebra generated from the observations of S and it represents the state of information about which paths are possible. In fact, if the possible paths are known, then the impossible paths are also known as well as any union of them. Consequently, this interpretation of the \mathcal{F}_t σ -algebra as containing the information learned by all the observations up to time t is consistent with the definition given above. Besides, as time marches on, more and more information becomes available while past information is still retained. As a result, the sequence of σ -algebras \mathcal{F}_t is increasing with time and represents the evolution of

information as it becomes available over time. This collection of increasing σ -algebras \mathcal{F}_t is called a natural filtration.

Measure and probability measure

Having defined a collection of subsets that is closed under both the complement and countable unions ensures that the space is measurable. A measure is a function which assigns a non-negative real number to any subset within the σ -algebra. It can be defined as follows:

Definition C.2

Let Ω be a non-empty set and let \mathcal{F} be a σ -algebra of subsets of Ω . A measure \mathbb{P} is a function that assigns to every set A in \mathcal{F} :

- (i) $\mathbb{P}(A) \geq 0$
- (ii) $\mathbb{P}(\emptyset) = 0$
- (iii) Whenever A_1, A_2, \dots is a countable sequence of pair-wise disjoint sets in \mathcal{F} ,
then:

$$\mathbb{P}\left(\bigcup_{i=1}^{\infty} A_i\right) = \sum_{i=1}^{\infty} \mathbb{P}(A_i)$$

A special case of measure is the probability measure. A probability measure is a real-valued function defined on a set of events in a probability space that satisfies the aforementioned measure properties while assigning one to the entire probability space. In other words, it is simply a normalized measure and its definition is given below.

Definition C.3

Let Ω be a non-empty set and let \mathcal{F} be a σ -algebra of subsets of Ω . A probability measure \mathbb{P} is a function that assigns to every set A in \mathcal{F} a real number in $[0, 1]$ called the probability of A and written $\mathbb{P}(A)$ such that:

- (i) $\mathbb{P}(\Omega) = 1$
- (ii) Whenever A_1, A_2, \dots is a sequence of disjoint sets in \mathcal{F} , then:

$$\mathbb{P}\left(\bigcup_{i=1}^{\infty} A_i\right) = \sum_{i=1}^{\infty} \mathbb{P}(A_i)$$

The triple $(\Omega, \mathcal{F}, \mathbb{P})$ is called a probability space.

Humans are usually very familiar with the historical probability measure (or observable probability measure or statistical probability measure). It is directly related to the observations that human-beings make and the likelihood they observe. Generally speaking, when models are constructed to simulate the experience of the real world, they are calibrated using the historical probability measure; hence the other name: statistical probability measure. It is however not the most convenient probability measure to work with for derivative pricing.

Equivalent probability measure

The definition of a probability measure given earlier does not guarantee that it is unique. Of interest in finance is the ability to change from one probability measure \mathbb{P} to another probability measure \mathbb{Q} without losing any information. For instance, if one trajectory of an asset price is possible under one probability measure, it must be possible under an equivalent probability measure, even if the likelihood of happening is minute, in order

not to lose information about that path. More generally, if one probability measure allows something that another probability measure doesn't, then some information is lost while switching from one measure to the other. To avoid this issue, the concept of equivalent probability measure is introduced.

Definition C.4

Two measures \mathbb{P} and \mathbb{Q} are equivalent if they operate on the same sample space and agree on what is possible and what is impossible. Formally, \mathbb{P} and \mathbb{Q} are equivalent if for every event A in Ω :

$$\mathbb{P}(A) = 0 \Leftrightarrow \mathbb{Q}(A) = 0$$

Having both probability measures agree on what is possible and impossible allows the definition of likelihood ratios and Radon-Nikodym derivatives which are used to perform changes of probability measures.

Change of probability measure

A popular technique to price derivatives consists in converting discounted asset price processes into martingales which are basically driftless processes. Mathematically, there are two venues to do this: either tweaking the process by operating on the values of the process, or operating on the probabilities associated with the process. The first technique is not used because it requires the knowledge of the drift of the process which is usually not available. The second technique is used extensively in finance and it aims at shifting the mean of a process by transforming the probability measure. For the purpose of derivative pricing, the transformation is done by switching from the historical probability measure, under which models describing asset price dynamics are calibrated, to a

synthetic probability measure called the risk-neutral measure, under which discounted prices are martingales (the expected discounted price is the current price).

Let's now look at the mechanics of creating a new probability measure equivalent to the original historical probability measure. This is done through an almost surely positive random variable Z with expectation equal to one under the original probability measure as shown below:

Theorem C.5

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space and let Z be an almost surely non-negative random variable such that $E(Z) = 1$. For $A \in \mathcal{F}$ let's now define $\mathbb{Q}(A) = \int_A Z(\omega) d\mathbb{P}(\omega)$. Then:

- (i) \mathbb{Q} is a probability measure
- (ii) If S is a non-negative random variable, then $E_{\mathbb{Q}}(S) = E_{\mathbb{P}}(S \cdot Z)$
- (iii) If Z is almost surely strictly positive, then $E_{\mathbb{P}}(S) = E_{\mathbb{Q}}\left(\frac{S}{Z}\right)$ for every non-negative random variable S

To use the proper terminology, the almost surely positive random variable Z with unit expectation used to link the two probability measures is called the *Radon-Nikodym derivative* of \mathbb{Q} with respect to \mathbb{P} . It is usually denoted as follows:

$$Z = \frac{d\mathbb{Q}}{d\mathbb{P}}$$

More intuitively, the Radon-Nikodym derivative may be interpreted using the likelihood ratio between two functions $f_{\mathbb{P}}^n$ and $f_{\mathbb{Q}}^n$, with each function describing the likelihood of the asset price following a particular trajectory described respectively under probability

measures \mathbb{P} and \mathbb{Q} . In this simplified setting, the Radon-Nikodym derivative $\frac{d\mathbb{Q}}{d\mathbb{P}}$ is

defined as the limit of the likelihood ratio as the “sampling” along trajectories goes to infinity. As the ratio must always be defined, the requirement for equivalent probability measures also becomes apparent.

Definition C.6

\mathbb{P} and \mathbb{Q} are two equivalent probability measures. Given a path for a sequence of n ordered time increments $\{t_1, t_2 \dots t_n\}$, let’s define $\{S_1, S_2 \dots S_n\}$ the realization of the process S at each of these increments. The Radon-Nikodym derivative $\frac{d\mathbb{Q}}{d\mathbb{P}}$ is defined as

the limit as n goes to infinity of the likelihood ratio:

$$\frac{d\mathbb{Q}}{d\mathbb{P}} = \lim_{n \rightarrow \infty} \frac{f_{\mathbb{P}}^n(S_1, S_2 \dots S_n)}{f_{\mathbb{Q}}^n(S_1, S_2 \dots S_n)}$$

Cameron-Martin-Girsanov theorem for change of measure

In the previous paragraph, the change of probability measure technique has been formalized using the Radon-Nikodym derivative. It was presented in a rather general setting and did not provide an appealing way to actually perform the change of measure. In some specific cases, the transformation can nonetheless be quite elegant. Fortunately, this is the case for Brownian motions which happen to be ubiquitous in finance.

The mapping between a stochastic process expressed under the \mathbb{P} measure and expressed under the \mathbb{Q} measure is exactly what the Cameron-Martin-Girsanov theorem provides. It ensures that, under some mild constraints, one can distort a probability measure to change the mean of a process. It is expressed as follows:

Definition C.7

Let W_t be a Brownian motion under probability measure \mathbb{P} and let γ_t be an \mathcal{F} -previsible process satisfying the condition that $E_{\mathbb{P}} \left(\exp \left(\frac{1}{2} \int_0^T \left(\frac{\mu - r_f}{\sigma} \right)^2 dt \right) \right)$ is finite. Then, there exists a measure \mathbb{Q} such that:

- (i) \mathbb{Q} is equivalent to \mathbb{P}
- (ii) $\frac{d\mathbb{Q}}{d\mathbb{P}} = \exp \left(- \int_0^T \gamma_t dW_t - \frac{1}{2} \int_0^T \gamma_t^2 dt \right)$
- (iii) $W_t^{\mathbb{Q}} = W_t + \int_0^T \gamma_s dW_s$ is a standard Brownian motion under \mathbb{Q}

In other words, the newly constructed Brownian motion $W_t^{\mathbb{Q}}$ is a driftless Brownian motion under the new probability measure \mathbb{Q} , while the original Brownian motion W_t now has drift $-\gamma_t$ under the probability measure \mathbb{Q} . Practically, what this means is that by choosing an appropriate process γ_t , it is possible to eliminate the drift of any Brownian motion so as to make it driftless. For financial applications, a proper choice of γ_t defines a new probability measure \mathbb{Q} under which discounted asset price processes are martingales.

Application

Let's assume that the asset price is governed by the following stochastic differential equation with constant drift μ and constant volatility σ . This model can be calibrated using real data obtained from historical time series. Therefore, under the historical probability measure \mathbb{P} , it is given by:

$$dS = \mu S dt + \sigma S dW$$

Let's now define γ as the following constant:

$$\gamma = \frac{\mu - r_f}{\sigma}$$

Then, using the Cameron-Martin-Girsanov theorem, a new equivalent probability measure \mathbb{Q} is defined, and under this new probability measure, the asset price process can be written as:

$$dS = \mu S dt + \sigma S dW_t^{\mathbb{Q}} - \gamma \sigma S dt = \mu S dt + \sigma S dW_t^{\mathbb{Q}} - (\mu - r_f) S dt$$

Rearranging the terms leads to the following familiar expression:

$$dS = r_f S dt + \sigma S dW$$

Under this measure, the drift of stock prices is identical to the risk-free rate of return. Better said, the discounted price is driftless. This is why this measure is called the risk-neutral probability measure. Incidentally, the constant γ defined by $\gamma = \frac{\mu - r_f}{\sigma}$ is usually called the risk premium, as it measures the extra return demanded by investors to hold one unit of risk.

Change of measure for some popular processes

In this paragraph, a change of measure is applied to different popular stochastic processes. For illustration purposes, the aim is to change the probability measure so that these processes either become risk-neutral or have their drift equal the risk-free rate of return. The Radon-Nikodym derivative denoted by ξ_t is also provided.

- Arithmetic Brownian motion (standard Brownian motion)

Let's start with the density function $f_{\mathbb{P}}$ given in Eq. 104 for a standard Brownian motion with constant drift μ' and constant volatility σ' described as $dS'_t = \mu' dt + \sigma' dW_t$. For the sake of simplicity, the process is assumed to start at zero at the time-origin.

$$f_{\mathbb{P}}(S'_t, t) = \frac{1}{\sigma' \sqrt{2\pi t}} \exp\left(\frac{-(S'_t - \mu' t)^2}{2\sigma'^2 t}\right) \quad \text{Eq. 104}$$

Let's now introduce the Radon-Nikodym derivative as shown in Eq. 105:

$$\xi_t = \frac{d\mathbb{Q}}{d\mathbb{P}} = \exp\left(-\frac{1}{2}\left(\frac{\mu' - r_f}{\sigma'}\right)^2 t - \frac{\mu' - r_f}{\sigma'} W_t\right) \quad \text{Eq. 105}$$

The Novikov condition stated in the Cameron-Martin-Girsanov theorem, expressed as

$E_{\mathbb{P}}\left(\exp\left(\frac{1}{2}\int_0^T \left(\frac{\mu' - r_f}{\sigma'}\right)^2 dt\right)\right) < \infty$, is certainly satisfied. The change of measure is done

using the transformation $f_{\mathbb{Q}}(S'_t, t) = f_{\mathbb{P}}(S'_t, t) \frac{d\mathbb{Q}}{d\mathbb{P}_t}$. This leads to the new density $f_{\mathbb{Q}}$ under probability measure \mathbb{Q} shown in Eq. 106:

$$f_{\mathbb{Q}}(S'_t, t) = \frac{1}{\sigma' \sqrt{2\pi t}} \exp\left(\frac{-(S'_t - \mu' t)^2}{2\sigma'^2 t}\right) \cdot \exp\left(-\frac{1}{2}\left(\frac{\mu' - r_f}{\sigma'}\right)^2 t - \frac{\mu' - r_f}{\sigma'} W_t\right) \quad \text{Eq. 106}$$

Expanding the previous expression leads to Eq. 107:

$$f_{\mathbb{Q}}(S'_t, t) = \frac{1}{\sigma' \sqrt{2\pi t}} \exp\left(\frac{-S_t'^2 + 2S_t' \mu' t - \mu'^2 t^2}{2\sigma'^2 t} + \frac{-t\mu'^2 + 2tr_f \mu' - tr_f^2}{2\sigma'^2} - \frac{\mu' S_t' - r_f S_t' - \mu'^2 t + r_f \mu' t}{\sigma'^2}\right) \quad \text{Eq. 107}$$

Simplifying and rearranging the terms in Eq. 107 yield Eq. 108 which is the density of a standard Brownian motion with constant drift r_f and constant volatility σ' . The change of measure therefore produced another standard Brownian motion but with a different drift r_f .

$$f_{\mathbb{Q}}(S'_t, t) = \frac{1}{\sigma' \sqrt{2\pi t}} \exp\left(-\frac{(S'_t - r_f t)^2}{2\sigma'^2 t}\right). \quad \text{Eq. 108}$$

- Geometric Brownian motion

Let's start again with the density function $f_{\mathbb{P}}$ given in Eq. 109 for a geometric Brownian motion with constant drift μ and constant volatility σ described as $dS_t = \mu S_t dt + \sigma S_t dW_t$. For the sake of simplicity, the process is assumed to start at the value S_0 at the time-origin.

$$f_{\mathbb{P}}(S_t, t) = \frac{1}{\sigma S_t \sqrt{2\pi t}} \exp\left(\frac{-(\ln(S_t) - \ln(S_0) - \hat{\mu}t)^2}{2\sigma^2 t}\right) \text{ with } \hat{\mu} = \mu - \frac{\sigma^2}{2} \quad \text{Eq. 109}$$

Let's now introduce the Radon-Nikodym derivative as shown in Eq. 110:

$$\xi_t = \frac{d\mathbb{Q}}{d\mathbb{P}} = \exp\left(-\frac{1}{2}\left(\frac{\hat{\mu} - r_f}{\sigma}\right)^2 t - \frac{\hat{\mu} - r_f}{\sigma} W_t\right) \quad \text{Eq. 110}$$

The Novikov condition stated in the Cameron-Martin-Girsanov theorem, expressed as $E_{\mathbb{P}}\left(\exp\left(\frac{1}{2}\int_0^T \left(\frac{\mu - r_f}{\sigma}\right)^2 dt\right)\right) < \infty$, is certainly satisfied. The change of measure is done

using the transformation $f_{\mathbb{Q}}(S_t, t) = f_{\mathbb{P}}(S_t, t) \frac{d\mathbb{Q}}{d\mathbb{P}_t}$ and this leads to the new density $f_{\mathbb{Q}}$ under probability measure \mathbb{Q} shown in Eq. 111:

$$f_{\mathbb{Q}}(S_t, t) = \frac{1}{\sigma S_t \sqrt{2\pi t}} \exp\left(\frac{-(\ln(S_t) - \ln(S_0) - \hat{\mu}t)^2}{2\sigma^2 t}\right) \cdot \exp\left(-\frac{1}{2}\left(\frac{\hat{\mu} - r_f}{\sigma}\right)^2 t - \frac{\hat{\mu} - r_f}{\sigma} W_t\right) \quad \text{Eq. 111}$$

Expanding the previous expression yields Eq. 112:

$$f_{\mathbb{Q}}(S_t, t) = \frac{1}{\sigma \sqrt{2\pi t}} \exp\left[\frac{-\ln(S_t)^2 - \ln(S_0)^2 - \hat{\mu}^2 t^2}{2\sigma^2 t} + \frac{\ln(S_t)(\ln(S_0) + \hat{\mu}t)}{2\sigma^2 t} + \frac{\ln(S_0)(\ln(S_t) - \hat{\mu}t)}{2\sigma^2 t} + \frac{\hat{\mu}t(\ln(S_t) - \ln(S_0))}{2\sigma^2 t} + \frac{-t\hat{\mu}^2 + 2tr_f\hat{\mu} - tr_f^2}{2\sigma^2} + \frac{-\hat{\mu}\ln(S_t) + \hat{\mu}\ln(S_0) + \hat{\mu}^2 t}{\sigma^2} + \frac{r_f \ln(S_t) - r_f \ln(S_0) - r_f \hat{\mu}t}{\sigma^2}\right] \quad \text{Eq. 112}$$

Simplifying and then rearranging the terms as previously done yield Eq. 113 which is the density of a geometric Brownian motion with constant drift r_f and constant volatility σ . The change of measure therefore produced another geometric Brownian motion but with a different drift r_f .

$$f_{\mathbb{Q}}(S_t, t) = \frac{1}{\sigma S_t \sqrt{2\pi t}} \exp\left(\frac{-(\ln(S_t) - \ln(S_0) - r_f t)^2}{2\sigma^2 t}\right). \quad \text{Eq. 113}$$

- Jump-diffusion process

The jump-diffusion process of Merton [134] is used. This model features a diffusion part, similar to a geometric Brownian motion, as well as a discontinuous part. It is assumed that the diffusion part and the discontinuous part of the jump-diffusion process are independent of each other. The discontinuous part is made of a finite number of jumps, the occurrence of which follows a Poisson counting process $N(t)$ of intensity λ . This means that the number of jump occurrences during a given time interval is directly related to the intensity λ as shown in Eq. 114.

$$\mathbb{P}(N(t) = k) = \frac{(\lambda t)^k}{k!} e^{-\lambda t} \quad \text{and} \quad \mathbb{P}(N(dt) = 1) \approx \lambda dt \quad \text{Eq. 114}$$

Provided that a jump occurred, its amplitude is random and denoted by the random variable J_i . The random variables corresponding to jump sizes are assumed to be independent and identically distributed. Such a jump process with random occurrences and random sizes is called a compound Poisson process and its mathematical expression is given in Eq. 115.

$$Q(t) = \sum_{i=1}^{N(t)} J_i, \quad t \geq 0 \quad \text{Eq. 115}$$

In Merton's model, the jump sizes are not directly modeled. Instead, it is the absolute price jump size defined as $J_t = S_{t^+}/S_{t^-}$ that is assumed to be log-normally distributed. The stochastic differential equation corresponding to this process is shown in Eq. 116. The fact that the absolute price jump size is log-normally distributed ensures that a jump to ruin can happen, but that an asset initially with positive price cannot end-up having a

negative price after a downward jump. In this stochastic differential equation, the term $J_t - 1$ is the relative price jump size.

$$dS_t = \mu S_t dt + \sigma S_t dW_t + (J_t - 1)S_t dN_t$$

$$dN_t = \begin{cases} 1 & \text{with probability } \lambda dt \\ 0 & \text{with probability } 1 - \lambda dt \end{cases} \quad \text{Eq. 116}$$

In this setting, the Radon-Nikodym derivative of the jump-diffusion process can be written as the product of a component related to the diffusion part and a component related to the jump part. For the latter, N_t denotes the Poisson process (number of jumps) while φ denotes the density for size of jumps. These two components are expressed in Eq. 117 and Eq. 118 respectively. With a single asset, the market is incomplete and there may be several risk-neutral measures characterized by different values of $\bar{\lambda}$ and $\bar{\varphi}$. Extra stocks are required to determine a unique risk-neutral measure.

$$\xi_t^{Diff} = \exp\left(-\frac{1}{2}\left(\frac{\mu - r_f}{\sigma}\right)^2 t - \frac{\mu - r_f}{\sigma} W_t\right) \quad \text{Eq. 117}$$

$$\xi_t^{Jump} = \exp\left((\lambda - \bar{\lambda})t\right) \prod_{k=1}^{N_t} \frac{\bar{\lambda}\bar{\varphi}(J_k)}{\lambda\varphi(J_k)} \quad \text{Eq. 118}$$

The step-by-step derivation of the risk-neutral transformation is beyond the scope of this thesis. However, the end-result is formulated in Eq. 119 and an interested reader is referred to Cont and Tankov [298] and Shreve [138] for a proper introduction to jump processes. Heuristically, the extra expected return due to jumps is removed from the risk-free drift of the process.

$$dS_t = (r_f - \lambda E(J_t - 1))dt + \sigma S_t dW_t^{\mathbb{Q}} + (J_t - 1)S_t dN_t \quad \text{Eq. 119}$$

- Mean-reverting process

Mean-reverting processes are another class of stochastic processes used extensively in finance. They may be used to model either returns that eventually move back to an average long-term return or commodity prices that may be disturbed but return towards a long-term equilibrium price. Mean-reversion for commodities is usually explained by the convenience yield and the cost of carry which affect the value of actually owning commodities. A popular mean-reverting process is the Ornstein-Uhlenbeck stochastic process. Mathematically, it can be expressed by Eq. 120 where the constant η represents the speed of adjustment or speed of mean-reversion, while \bar{S} represents the long-term mean.

$$dS_t = \eta(\bar{S} - S_t)dt + \sigma dW_t \quad \text{Eq. 120}$$

The step-by-step derivation of the risk-neutral transformation is beyond the scope of this thesis. However, to get to the risk-neutral form, the value of the market-price of risk for a diffusive process is introduced and defined as λ . This metric is estimated by observing the market and finding another purely diffusive asset for which the market-price of risk can be estimated as was presented previously for the geometric Brownian motion. Using this result, the risk-neutral form for the Ornstein-Uhlenbeck stochastic process is formulated in Eq. 121. An interested reader is referred to Bjerksund and Ekern [299] for the actual derivation.

$$dS_t = \eta\left(\bar{S} - \frac{\sigma\lambda}{\eta} - S_t\right)dt + \sigma dW_t^Q \quad \text{Eq. 121}$$

APPENDIX D: ESSCHER TRANSFORM

The Esscher transform is at the core of this research effort. It is consequently appropriate to provide a bit more information about this time-honored technique commonly used in actuarial finance and recently adapted for financial engineering applications by Gerber and Shiu [140]. The Esscher approximation was initially introduced by Esscher [141] to approximate the (upper) tails of total claim distributions by shifting the mean of the aggregate claims to a point of interest. In the context of collective risk theory, this practice is useful to estimate insurance premiums for stop-loss insurance policies. For financial applications, the Esscher transform is applied to the return distribution of one or many assets, and the transformation induces a parametric change of probability measure. The parameter h (possibly a vector) is chosen to ensure that the discounted price of each asset becomes a martingale under the new probability measure.

For a single probability density function f and a real number h , the Esscher transform f_{Ess} with parameter h is expressed using the moment generating function M of f as shown in Eq. 122. Looking at this definition, the Esscher transform is the product of an exponential function and a density function, normalized by a moment generating function. As a result, this transformation induces an equivalent probability measure as both distributions agree on sets with probability zero.

$$f_{Ess}(x, h) = \frac{e^{hx} f(x)}{M(h)}, \text{ with } h \in \mathbb{R} \text{ and } M(h) = \int_{-\infty}^{\infty} e^{hx} f(x) dx \quad \text{Eq. 122}$$

Now, let's adapt this definition for a process with stationary and independent increments.

For positive time indices, let's denote by S such a process and let's denote by X its

corresponding continuously compounded rate of return process with null value at the time-origin, as shown in Eq. 123. X has an infinitely divisible distribution which means that it can be described as the sum of independent and identically distributed random variables. For asset pricing purposes, this is equivalent to saying that a price observation can be done infinitely many times during a given time-interval and that the return distributions of the price observations are independent and identically distributed. Besides, the return over the entire time interval is, in the limit, the sum of the returns for each of the observation.

$$S(t) = S(0) \cdot e^{X(t)}, \quad \text{with } X(0) = 0 \quad \text{Eq. 123}$$

Let's introduce F as the cumulative density function for these continuously compounded returns and let's call f its density counterpart. Both functions are dependent on two variables: the return value x and the time index t as shown in Eq. 124 and Eq. 125:

$$F(x, t) = P(X(t) \leq x) \quad \text{Eq. 124}$$

$$f(x, t) = \frac{d}{dx}(F(x, t)) \quad \text{Eq. 125}$$

Let's now introduce the moment generating function for the process X as shown in Eq. 126. Again, it is a function of two variables: one is the usual parameter h of moment generating functions while the other is the time index t .

$$M(h, t) = \int_{-\infty}^{\infty} e^{hx} f(x, t) dx \quad \text{Eq. 126}$$

X is an infinitely divisible distribution and each increment is independent. By assuming that M is continuous at $t=0$, the properties of moment generating functions lead to the following identity proven in Breiman [300]:

$$M(h, t) = M(h, 1)^t \quad \text{Eq. 127}$$

Having defined the moment generating function M , all the brick-elements are available to define the new Esscher transform, denoted by f_{ESS} , of parameter h for a stochastic process X . The definition is given for any positive t in Eq. 128:

$$f_{ESS}(x, t, h) = \frac{e^{hx} f(x, t)}{M(h, t)}, \quad \text{for } x \in \mathbb{R} \text{ and } h \in \mathbb{R} \text{ and } t \in \mathbb{R}^+ \quad \text{Eq. 128}$$

Looking at this definition, the Esscher transform is the product of an exponential function and a density function, normalized by a moment generating function. Consequently, the expectation or, better said, the integral with regards to x over the entire real space is one. In addition, both the original density function f and the Esscher transformed function f_{ESS} have the same support. As a result, this transformation induces an equivalent probability measure as both distributions agree on zero probability sets. The function f_{ESS} is the new distribution function of X under the new probability measure and it is called the Esscher transform of the original distribution. The moment generating function M_{ESS} for this new function has three parameters in total: h and t stemming from the Esscher transformed distribution itself, as well as the usual moment generating function parameter. Its expression is provided in Eq. 129.

$$M_{ESS}(z, t; h) = E_{ESS}^h(e^{X(t)}) = \int_{-\infty}^{\infty} e^{zx} f(x, t, h) dx = \int_{-\infty}^{\infty} \frac{e^{h+zx} f(x, t)}{M(h, t)} dx$$

$$M_{ESS}(z, t; h) = \frac{M(z + h, t)}{M(h, t)}$$

Eq. 129

The essence of the Esscher technique is to use the parameter h to ensure that the new probability measure is an equivalent martingale measure. Let's call h^* this specific value of h . Under the equivalent martingale measure, the discounted price of assets is a martingale which means that the current price of these assets is exactly their expectation. As a result, h^* is determined such that $S(0) = E_{ESS}^{h^*}(e^{-r_f t} S(t))$. Using also the fact that $S(t) = S(0) \cdot e^{X(t)}$ leads to solving Eq. 130:

$$S(0) = E_{ESS}^{h^*}(e^{-r_f t} \cdot S(0) \cdot e^{X(t)}) = e^{-r_f t} \cdot S(0) \cdot E_{ESS}^{h^*}(e^{X(t)})$$

Eq. 130

After some simplifications, one gets Eq. 131:

$$e^{r_f t} = E_{ESS}^h(e^{X(t)}) = M_{ESS}(z, t, h)$$

Eq. 131

Now, using the identity expressed in Eq. 127 yields the identity shown in Eq. 132 for the moment generating function of the Esscher transformed distribution:

$$M_{ESS}(z, t; h) = (M_{ESS}(z, 1; h))^t$$

Eq. 132

Using both the logarithm of Eq. 131 as well as Eq. 132, h^* is determined such that it solves the equation displayed in Eq. 133:

$$r_f = \ln (M_{ESS}(1,1; h^*))$$

Eq. 133

There is no closed-form expression for h^* but any numerical solver or numerical method should be able to handle this problem. Gerber and Shiu [151] have shown that the parameter h^* is unique. The corresponding transformation is called the risk-neutral Esscher transform and the corresponding probability measure is called the risk-neutral probability measure or equivalent martingale measure. The Esscher transformation has been applied to various pricing problems in finance. It presents the advantage of being both a rather straightforward and versatile technique. Indeed, it can handle many different types of processes, including some of the most commonly used stochastic processes in finance, such as diffusion processes and diffusion processes with jumps.

Non-parametric Esscher transform

The motivation behind the non-parametric Esscher transform methodology is to be able to risk-neutralize a distribution of asset prices or cash flows that may be obtained in the first place either through observations, bootstrapping techniques, or thanks to simulations. The ability to generate a risk-neutral distribution directly from a real, observable distribution without the need to create and calibrate a stochastic model can prove extremely handy for real options analysis. Indeed, when performing a probabilistic design analysis, there is usually a need to evaluate many different concepts and therefore many different design points. Calibrating a stochastic model for each and every of these data points would be very time-consuming and inefficient.

A technique that directly fits the data was first proposed for financial applications by Pereira et al. [154]. It is a data-driven technique and starts with a distribution of asset prices from which the log-return is estimated. To follow usual notations, all empirical estimates (empirical rate of return and empirical moment generating function) are

denoted with a hat. Using a sample of size n representing the distribution of asset prices, a vector of empirical continuously compounded rates of return is constructed and denoted by \widehat{X}_t as shown in Eq. 134. Each index i represents a different scenario and a different state of the economy under the historical probability measure.

$$\widehat{X}_t = [x_t^1, x_t^2, x_t^i \dots x_t^n] = \left[\ln\left(\frac{S_t^1}{S_{t-1}^1}\right), \ln\left(\frac{S_t^2}{S_{t-1}^2}\right), \ln\left(\frac{S_t^i}{S_{t-1}^i}\right) \dots \ln\left(\frac{S_t^n}{S_{t-1}^n}\right) \right] \quad \text{Eq. 134}$$

This vector is then used to derive the empirical moment generating function, analogous to the one defined earlier in Eq. 126. This empirical moment generating function is denoted by \widehat{M}_t and is given by Eq. 135.

$$\widehat{M}(h, t) = \frac{1}{n} \sum_{i=1}^n e^{hx_t^i} \quad \text{Eq. 135}$$

Similarly, the empirical moment generating function of the yet to be generated Esscher transformed distribution is given by Eq. 136.

$$\widehat{M}_{ESS}(z, t, h) = \frac{\widehat{M}(z + h, t)}{\widehat{M}(h, t)} \quad \text{Eq. 136}$$

Continuing to draw the parallel with the “original” Esscher transform technique previously presented, the empirically-driven value \widehat{h}^* that will make the Esscher transform a risk-neutral Esscher transform was originally given by Eq. 133, which in the non-parametric case, yields Eq. 137.

$$r_f = \ln\left(\widehat{M}_{ESS}(1, 1; \widehat{h}^*)\right) = \ln\left(\frac{\widehat{M}(\widehat{h}^* + 1, 1)}{\widehat{M}(\widehat{h}^*, 1)}\right) \quad \text{Eq. 137}$$

Rearranging the terms in Eq. 137 and using the vector of observed returns yield the equation for \widehat{h}^* described in Eq. 138.

$$\widehat{h}^* = \arg \left\{ \ln \left(\frac{\sum_{i=1}^n e^{(h+1)x_t^i}}{\sum_{i=1}^n e^{hx_t^i}} \right) = r_f \right\} \quad \text{Eq. 138}$$

Having solved for \widehat{h}^* using numerical methods (such as Newton-Raphson, bisection, etc.), the next step is to actually perform the change of measure. How to do that? Since the idea is to change the probability measure, one technique is to assign a different weight to each and every observation of the original distribution. In the original sample of size n representing the original distribution, each observation carries exactly the same weight or the same probability, which is exactly $1/n$. The non-parametric Esscher transform is going to change these probabilities by assigning a specific weight to each of these observations in order to tilt and distort the original probability measure. The resulting sample, still of size n , is risk-neutral and represents a drawing from the risk-neutral Esscher transform.

To make the sample risk-neutral, a risk-neutral probability vector $\mathbb{Q}_t^{\widehat{h}^*}$ is introduced. Its expression is exactly the “empirical” counterpart of the Esscher transform initially shown in Eq. 128. It is however normalized to ensure that the sum of its elements is equal to one as given in Eq. 139. This set of probabilities is risk-neutral, and it is this set that is used for the computation of expectations necessary to perform options pricing.

$$\mathbb{Q}_t^{h^*} = \left[\frac{e^{\widehat{h}^* x_t^1}}{\sum_{i=1}^n e^{\widehat{h}^* x_t^i}}, \frac{e^{\widehat{h}^* x_t^2}}{\sum_{i=1}^n e^{\widehat{h}^* x_t^i}}, \dots, \frac{e^{\widehat{h}^* x_t^n}}{\sum_{i=1}^n e^{\widehat{h}^* x_t^i}} \right] \quad \text{Eq. 139}$$

APPENDIX E: MODERN PORTFOLIO THEORY

Modern portfolio theory (MPT) aims at formulating an optimization framework for the design of financial portfolios. In essence, modern portfolio theory recognizes that investing is a trade-off between expected return and risk, and that investors will therefore seek the highest possible reward (greed) while avoiding risk as much as possible. In this setting, the reward is expressed as the expected rate of return of an asset while the risk is expressed as the uncertainty or volatility of the rate of return. The theoretical justification for this analysis is Samuelson's *Fundamental Approximation Theorem of Portfolio Analysis* which proves that under some compactness¹ assumptions:

- The importance of all moments of the return distribution beyond the variance is much smaller than the expected value and the variance itself
- The variance of the return distribution is as important as the expected value of the return for the investor welfare

Consequently, the modern portfolio theory relies on a tradeoff analysis between mean and volatility (square root of variance). To illustrate the theory, several examples featuring increasingly more complex situations are described and analyzed next.

In an economy featuring two perfectly correlated assets A and B , a portfolio could be any combination of these two assets. In this case, both the expected return of the portfolio and its volatility would be linear combinations of respectively the expected return and the

¹ The distribution of the rate of return on a portfolio is said to be compact if the risk can be controlled by an investor at any time. Practically, this means that sudden jumps in stock prices are absent and that as a position in a risky portfolio is held for shorter and shorter times, the risk to the investor decreases and approaches zero.

volatility of assets A and B as shown in Eq. 140 and Eq. 141 below. The blue line in the exhibit (a) of Figure 127 represents all the different possible combinations of A and B .

$$E(r_p) = w_A E(r_A) + (1 - w_A) E(r_B) \quad \text{Eq. 140}$$

$$\sigma_p^2 = w_A^2 \sigma_A^2 + (1 - w_A)^2 \sigma_B^2 + 2w_A(1 - w_A)\sigma_A\sigma_B = (w_A\sigma_A + (1 - w_A)\sigma_B)^2 \quad \text{Eq. 141}$$

In an economy featuring two negatively correlated assets A and B , a portfolio combining any of these two assets would have an expected return being linear in the expected return of assets A and B as shown in Eq. 142. What is interesting in this case is that there exists a portfolio with a specific value of w_A for which the risk of the portfolio is neutralized as shown in Eq. 143 below and in exhibit (b) of Figure 127.

$$E(r_p) = w_A E(r_A) + (1 - w_A) E(r_B) \quad \text{Eq. 142}$$

$$\sigma_p^2 = w_A^2 \sigma_A^2 + (1 - w_A)^2 \sigma_B^2 - 2w_A(1 - w_A)\sigma_A\sigma_B = (w_A\sigma_A - (1 - w_A)\sigma_B)^2 \quad \text{Eq. 143}$$

In particular, there exists a portfolio such that: $\sigma_p = 0 \Leftrightarrow w_A = \frac{\sigma_B}{\sigma_B - \sigma_A}$




Figure 127: Expected return of a portfolio with two perfectly and negatively correlated assets

In an economy featuring two uncorrelated assets A and B , a portfolio combining any of these two assets would have an expected return being linear in the expected return of

assets A and B and would feature a volatility in between the maximum and minimum of the two previous examples. This is mathematically described by Eq. 144 and Eq. 145.

$$E(r_p) = w_A E(r_A) + (1 - w_A) E(r_B) \quad \text{Eq. 144}$$

$$\sigma_p^2 = w_A^2 \sigma_A^2 + (1 - w_A)^2 \sigma_B^2 \quad \text{Eq. 145}$$

A more generic setting is an economy featuring two correlated assets A and B with a correlation strictly in between plus and minus one. In this case, a portfolio combining any of these two assets would have an expected return being linear in the expected return of assets A and B , while its volatility would be subject to the correlation between the two. This is shown in Eq. 146 and Eq. 147 below. Again, the volatility of the constructed portfolios would be between the maximum and minimum volatilities of the first two examples. This is represented by the blue line that curves to the left in the exhibit (a) of Figure 128.

$$E(r_p) = w_A E(r_A) + (1 - w_A) E(r_B) \quad \text{Eq. 146}$$

$$\sigma_p^2 = w_A^2 \sigma_A^2 + (1 - w_A)^2 \sigma_B^2 + 2w_A(1 - w_A)\sigma_A\sigma_B\rho_{AB} \quad \text{Eq. 147}$$

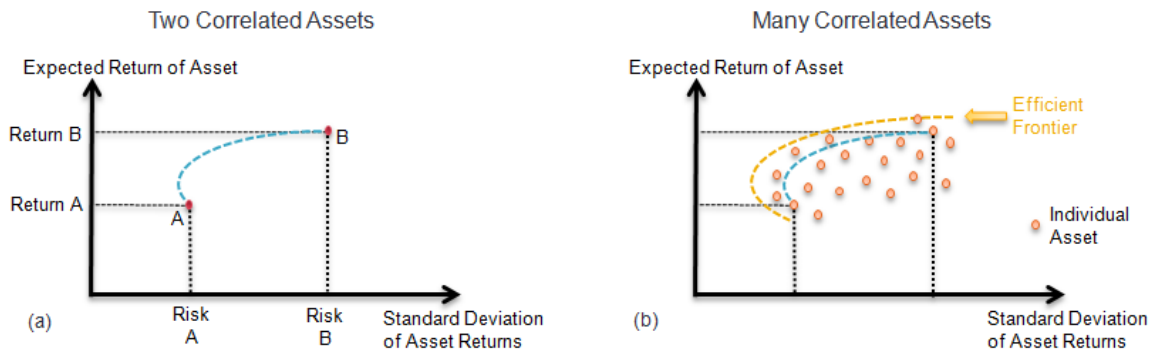


Figure 128: Expected return of a portfolio with two and more correlated assets

Let's now move to an economy featuring more than two correlated assets. There are an infinite number of asset combinations: some of these combinations result in portfolios with similar rates of return but different volatilities, some others result in portfolios with similar volatilities but different rates of return. Of interest is the locus of portfolios that, for a given rate of return, have minimum volatilities. These portfolios form a Pareto frontier of *efficient portfolios* and are represented by the yellow contour line in exhibit (b) of Figure 128. A risk-averse investor would select an efficient portfolio featuring lower volatility whereas a risk-neutral investor may be more inclined to take risks and select an efficient portfolio with higher return.

In the last example, a risk-free asset such as a Treasury bill is added to the previous economy. Since the risk-free asset has zero volatility, combinations of the risk-free asset and one efficient portfolio from the Pareto frontier may be graphically represented by the red line joining these two assets in exhibit (a) of Figure 129. Of interest is the green line tangent to the Pareto frontier in exhibit (b) of Figure 129. Since this line has the highest slope possible, portfolios along this *capital allocation line (CAL)* joining the risk-free asset and the tangency portfolio dominate all other portfolios from a mean-variance standpoint. In this case, the useful part of the efficient frontier collapses to a single efficient portfolio called the tangency portfolio.




Figure 129: Expected returns and volatilities for portfolios featuring a risk-free asset and many correlated assets

Along the capital allocation line, the expected return and risk tradeoff can be mathematically represented by equation Eq. 148.

$$E(r_p) = r_f + \left(\frac{E(r_t) - r_f}{\sigma_t} \right) \sigma_p \quad \text{Eq. 148}$$

APPENDIX F: CAPITAL ASSET PRICING MODEL

Modern portfolio theory provides the foundation on which the Capital Asset Pricing Model (CAPM) is built. The model developed by Sharpe [61], Lintner [62], and Mossin [63] aims at providing a prediction of the relationship that should be observed between the expected rate of return of an asset and its risk. In this model, the authors argue that only systematic risk should be rewarded since exposure to idiosyncratic risk can be mitigated by holding a well-diversified portfolio. This relationship between risk and return is of great importance in finance as it helps benchmark the rate of returns of investments.

One of the main assumptions underlying the model is that investors constitute a group of rational, risk-averse individuals with homogeneous expectations. This means that despite their different initial wealth and different attitudes towards risk, all investors analyze securities in the same way and share the same economic view of the world. Having access to all the information at the same time, they will all end up with the same expected returns, the same correlation matrix, and therefore the same set of efficient portfolios and the same tangency portfolio.

Of interest is the fact that investors will hold various combinations (according to their risk-aversion) of the same assets, namely the risk-free asset and the efficient tangency portfolio. Since all investors choose to hold the same portfolio of risky assets, this portfolio must represent the overall economy and must include all available assets. This is why this efficient tangency portfolio is referred to as the *market portfolio* and the

corresponding capital allocation line is referred to as the *capital market line* as shown in exhibit (a) of Figure 130.

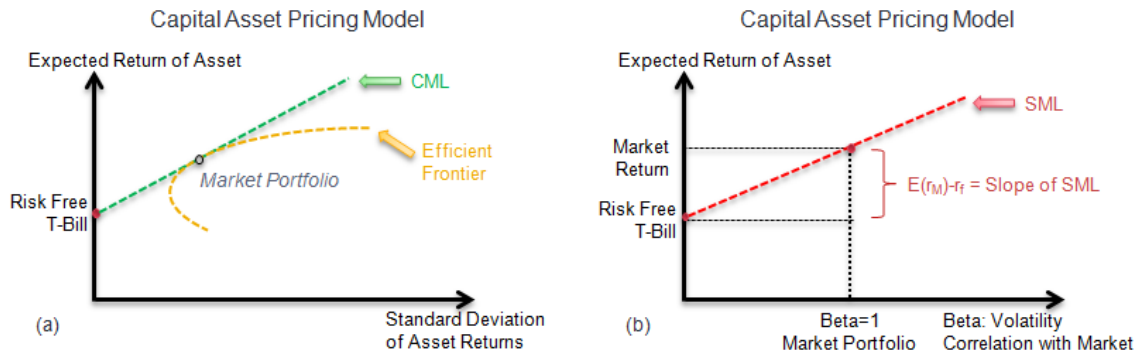


Figure 130: Capital Market Line (CML) and Security Market Line (SML)

Another interesting aspect of the capital asset pricing model is that it recognizes that investors should not be rewarded for idiosyncratic risk. Indeed, as portfolios get diversified, some sources of risks that are asset-specific will eventually see their impact diminish (law of averages), while some other sources of risks might even cancel out (negative correlations between the returns of some assets). However, as diversified as the portfolio might be, there will still be some residual risk because of non-controllable and non-predictable macroeconomics factors that affect the overall economy and therefore the return of each individual asset. This is the systematic or market risk. As much as there is no reason to reward investors exposing themselves to unnecessary risks by not diversifying their portfolios, the CAPM is built upon the insight that the risk premium on an asset is determined by its contribution to the risk of the overall market portfolio. Let's now translate mathematically the implications of these assumptions by first showing in Eq. 149 the contributions of asset i to the return and variance of the market portfolio.

$$w_i \cdot \sum_{k=1}^n w_k \cdot Cov(r_i, r_k) = w_i \cdot [E(r_i) - r_f] \quad \text{Eq. 149}$$

Now, let's express in Eq. 150 the covariance of asset i with the market portfolio using the expanded definition of the market return. It becomes obvious that the contribution of asset i to the overall variance of the market portfolio can be re-written in a single covariance term as shown in Eq. 151.

$$Cov(r_i, r_M) = Cov\left(r_i, \sum_{k=1}^n w_k r_k\right) = \sum_{k=1}^n w_k \cdot Cov(r_i, r_k) \quad \text{Eq. 150}$$

$$w_i \cdot \sum_{k=1}^n w_k \cdot Cov(r_i, r_k) = w_i \cdot Cov(r_i, r_M) \quad \text{Eq. 151}$$

The reward-to-risk ratio for holding some asset i in the market portfolio is therefore given in Eq. 152 by the ratio of the asset contribution to the risk premium over its contribution to the variance. It was also proven earlier that the market portfolio is the tangency portfolio, and therefore its reward-to-risk ratio (risk premium) is known and given by Eq. 153.

$$\text{Reward to risk for Asset } i = \frac{w_i [E(r_i) - r_f]}{w_i \cdot Cov(r_i, r_M)} = \frac{E(r_i) - r_f}{Cov(r_i, r_M)} \quad \text{Eq. 152}$$

$$\text{Reward to risk for Market Portfolio} = \frac{E(r_M) - r_f}{Cov(r_M, r_M)} = \frac{E(r_M) - r_f}{\sigma^2} \quad \text{Eq. 153}$$

There is no reason for the reward-to-risk ratio of asset i to be different from the reward-to-risk ratio of the market portfolio. Otherwise, the economy would no longer be in

equilibrium as demand for one of the most performing asset would drive its return down. As a result, these two ratios must be equal. Rearranging the equality between Eq. 152 and Eq. 153 leads to Eq. 154 and Eq. 155. The new parameter *beta* denoted by β is introduced as the ratio of the covariance over the market portfolio variance. It measures the contribution of asset *i* to the variance of the market portfolio as a fraction of the total variance of the market portfolio.

$$\frac{E(r_M) - r_f}{\sigma^2} = \frac{E(r_i) - r_f}{Cov(r_i, r_M)} \Leftrightarrow E(r_i) - r_f = \frac{Cov(r_i, r_M)}{\sigma^2} [E(r_M) - r_f] \quad \text{Eq. 154}$$

$$E(r_i) = r_f + \beta [E(r_M) - r_f] \quad \text{Eq. 155}$$

This last equation is the well-known expected return-beta relationship of the CAPM, which defines the theoretical risk premium demanded by investors for each asset in the economy. This relationship can be graphically portrayed as the *security market line* (SML) shown in exhibit (b) of Figure 130. Despite the similarity between the capital market line and the security market line, there is a fundamental difference between the two. Contrary to the capital market line which links reward and risk for an entire and well-diversified portfolio (hence standard deviation the measure of risk), the security market line links reward and risk for individual assets (hence beta the measure of risk).

APPENDIX G: BLACK-SCHOLES-MERTON OPTION PRICING MODEL

The Black-Scholes-Merton model is a mathematical model of a financial market containing certain derivative instruments. With this model, Fischer Black and Myron Scholes were able to derive analytical formulas for the market price of the simplest financial options (European call and put options also called plain vanilla options). The derivation of this formula in 1973 was immediately embraced by practitioners and led to a boom in the trading of derivatives worldwide.

Setting up the Black-Scholes-Merton model

The idea introduced by Black and Scholes in their seminal paper [66] is to design a self-financing portfolio consisting of one share of the underlying stock as well as some quantity of the option such that the entire portfolio is riskless. The riskless portfolio means that it is insensitive to changes in the price of the underlying, and the rate of return of the portfolio is exactly the risk-free rate of return. To simplify the analysis, Black and Scholes make several assumptions regarding the market model:

- No arbitrage opportunity which precludes the possibility of riskless profit.
- Underlying stock $S(t)$ does not pay any dividend or any other distribution.
- Possibility to borrow or lend cash at the same constant short-term interest rate r .
- No transaction cost on buying or selling the stock $S(t)$ or the option $V(t)$.

- Ability to buy or sell any amount of the underlying stock $S(t)$ of the associated option $V(t)$ or a riskless bond. This includes positive amounts (long position), negative amounts (short position), as well as fractional amounts.

- Underlying stock $S(t)$ follows a geometric Brownian motion with a constant trend μ and a constant volatility σ . A geometric Brownian motion is a continuous-time stochastic process which satisfies the following stochastic differential equation in Eq. 156 where W is a Wiener¹ process:

$$dS = \mu S dt + \sigma S dW \quad \text{Eq. 156}$$

The value of the riskless portfolio made of one share of the underlying as well as θ shares of the options at time t is given by the following formula in Eq. 157:

$$\Pi(t) = S(t) + \theta V(t) \quad \text{Eq. 157}$$

The self-financing assumption means that the change in value of the portfolio is due only to the change in value of the assets within the portfolio as shown in Eq. 158:

$$d\Pi = dS + \theta dV \quad \text{Eq. 158}$$

Using Ito's lemma, we have Eq. 159:

$$dV = \left(\frac{\partial V}{\partial t} + \mu S \frac{\partial V}{\partial S} + \frac{1}{2} \sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} \right) dt + \sigma S \frac{\partial V}{\partial S} dW \quad \text{Eq. 159}$$

¹ Wiener process: A Wiener process is a continuous-time stochastic process named for Norbert Wiener. It is characterized by the following properties:

- $W_0 = 0$
- The function $t \rightarrow W_t$ is almost surely continuous everywhere
- W_t has independent increments with $W_t - W_s$ normally distributed $N(0, t-s)$

And therefore Eq. 160:

$$d\Pi = \left(\mu S + \theta \frac{\partial V}{\partial t} + \mu S \theta \frac{\partial V}{\partial S} + \frac{1}{2} \theta \sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} \right) dt + \left(\sigma S + \theta \sigma S \frac{\partial V}{\partial S} \right) dW \quad \text{Eq. 160}$$

Under this set of assumptions, the value of the option depends only on the price of the stock as well as on the time. Therefore, it is possible to create a hedged position which is absolutely risk-free. This means that the rate of return is certain and is exactly the risk-free rate of return. Mathematically, this means that the deterministic part of Eq. 160 is equal to the risk-free interest rate and the stochastic part is null.

For the stochastic part (dW):

$$\sigma S + \theta \sigma S \frac{\partial V}{\partial S} = 0 \Leftrightarrow \theta = -\frac{\partial S}{\partial V} \quad \text{Eq. 161}$$

For the deterministic part (dt):

$$\mu S + \theta \frac{\partial V}{\partial t} + \mu S \theta \frac{\partial V}{\partial S} + \frac{1}{2} \theta \sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} = r\Pi = r(S + \theta V) \quad \text{Eq. 162}$$

Dividing by θ (which is never null since the ratio $1/\theta$ is never infinite) and then simplifying using the left-hand side of Eq. 161 yields Eq. 163:

$$\frac{\partial V}{\partial t} + \frac{1}{2} \sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} = r \left(\frac{S}{\theta} + V \right) \quad \text{Eq. 163}$$

Rearranging the terms in Eq. 163 using the right-hand side of Eq. 161 leads to the Black-Scholes-Merton partial differential equation displayed in Eq. 164:

$$\frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} + rS \frac{\partial V}{\partial S} - rV = 0 \quad \text{Eq. 164}$$

Solving for the derivative price

The Black-Scholes-Merton partial differential equation is called backward parabolic. To solve such an equation, initial and terminal conditions (at maturity) are required. The terminal condition is given by the type of derivative under investigation and therefore by the payoff at maturity. The discounted Feynman-Kac theorem briefly stated below may be applied to get the solution to the partial differential equation.

Discounted Feynman-Kac theorem

Suppose that S_t follows the stochastic process

$$dS_t = \mu(S_t, t)dt + \sigma(S_t, t)dW_t^Q$$

Where W_t^Q is a Brownian motion under the measure \mathbb{Q} .

Let $V(S_t, t)$ be a differentiable function of S_t and t , and suppose that $V(S_t, t)$ follows the partial differential equation given by:

$$\frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2(S_t, t) \frac{\partial^2 V}{\partial S^2} + \mu(S_t, t) \frac{\partial V}{\partial S} - r(S_t, t)V(S_t, t) = 0$$

With boundary conditions $V(x_T, T)$, then the solution $V(S_t, t)$ can be expressed as a conditional expectation:

$$V(S_t, t) = E^{\mathbb{Q}} \left[e^{-\int_t^T r(S_u, u) du} \cdot V(S_T, T) | \mathcal{F}_t \right]$$

Let's now apply the discounted Feynman-Kac theorem to the Black-Scholes-Merton partial differential equation by recognizing the drift term $\mu(S_t, t) = \mu S$, the volatility term $\sigma(S_t, t) = \sigma S$ and the constant risk-free rate $r(S_u, u) = r$. This leads to the following equation expressed in Eq. 165:

$$V(S_t, t) = E^{\mathbb{Q}} \left[e^{-\int_t^T r du} \cdot V(S_T, T) | \mathcal{F}_t \right] = e^{-r(T-t)} E^{\mathbb{Q}} [V(S_T, T) | \mathcal{F}_t] \quad \text{Eq. 165}$$

Now, to go further and obtain a closed-form solution to this mathematical expectation, the boundary conditions must be specified. The boundary conditions are dependent on the type of derivative that is to be priced. For a European call option with strike price K , the boundary condition may be expressed as $V(S_t, T) = \max(S_t - K, 0)$. In this case, the European call option price is given by Eq. 166:

$$V(S_t, t) = e^{-r(T-t)} E^{\mathbb{Q}} [\max(S_T - K, 0) | \mathcal{F}_t] \quad \text{Eq. 166}$$

The stock price follows a geometric Brownian motion and therefore it can be written as in Eq. 167:

$$S_T = S_t e^{\left(r - \frac{\sigma^2}{2}\right)(T-t) + \sigma(W(T) - W(t))} \quad \text{Eq. 167}$$

Therefore, Eq. 168 represents the European call option expression:

$$V(S_t, t) = e^{-r(T-t)} E^{\mathbb{Q}} \left[\max \left(S_t e^{\left(r - \frac{\sigma^2}{2}\right)(T-t) + \sigma(W(T) - W(t))} - K, 0 \right) | \mathcal{F}_t \right] \quad \text{Eq. 168}$$

$W(T) - W(t)$ is a normally distributed random variable with mean zero and variance $T - t$. The expectation can be replaced by an integral featuring the normal probability density function as shown in Eq. 169:

$$V(S_t, t) = \frac{e^{-r(T-t)}}{\sqrt{2\pi(T-t)}} \int_{x=-\infty}^{+\infty} \max\left(S_t e^{\left(r-\frac{\sigma^2}{2}\right)(T-t)+\sigma \cdot x} - K, 0\right) \cdot e^{\frac{-x^2}{2(T-t)}} dx \quad \text{Eq. 169}$$

Letting: $y = x/\sqrt{T-t}$, and performing the resulting change of variable yields the expectation in Eq. 170:

$$V(S_t, t) = \frac{e^{-r(T-t)}}{\sqrt{2\pi}} \int_{y=-\infty}^{+\infty} \max\left(S_t e^{\left(r-\frac{\sigma^2}{2}\right)(T-t)+\sigma y\sqrt{T-t}} - K, 0\right) \cdot e^{\frac{-y^2}{2}} dy \quad \text{Eq. 170}$$

The *max* function may be removed by choosing an appropriate lower bound for the integral. This yields Eq. 171 below:

$$V(S_t, t) = \frac{e^{-r(T-t)}}{\sqrt{2\pi}} \int_{y=\frac{\ln\left(\frac{K}{S_t}\right)-\left(r-\frac{\sigma^2}{2}\right)(T-t)}{\sigma\sqrt{T-t}}}^{+\infty} \left(S_t e^{\left(r-\frac{\sigma^2}{2}\right)(T-t)+\sigma y\sqrt{T-t}} - K\right) \cdot e^{\frac{-y^2}{2}} dy \quad \text{Eq. 171}$$

This expression can then be split into two parts. Letting: $\alpha = \frac{\ln\left(\frac{K}{S_t}\right)-\left(r-\frac{\sigma^2}{2}\right)(T-t)}{\sigma\sqrt{T-t}}$, this leads

to Eq. 172 and Eq. 173:

$$I(S_t, t) = \frac{e^{-r(T-t)}}{\sqrt{2\pi}} \int_{y=\alpha}^{+\infty} S_t e^{\left(r-\frac{\sigma^2}{2}\right)(T-t)+\sigma y\sqrt{T-t}} \cdot e^{\frac{-y^2}{2}} dy \quad \text{Eq. 172}$$

$$J(S_t, t) = \frac{e^{-r(T-t)}}{\sqrt{2\pi}} \int_{y=\alpha}^{+\infty} K \cdot e^{-\frac{y^2}{2}} dy \quad \text{Eq. 173}$$

With S_t being a constant, rearranging the first integral I yields Eq. 174:

$$I(S_t, t) = \frac{S_t}{\sqrt{2\pi}} \int_{y=\alpha}^{+\infty} e^{-\frac{1}{2}(y^2 - 2\sigma y\sqrt{T-t} + \sigma^2(T-t))} dy = \frac{S_t}{\sqrt{2\pi}} \int_{y=\alpha}^{+\infty} e^{-\frac{1}{2}(y - \sigma\sqrt{T-t})^2} dy \quad \text{Eq. 174}$$

Finally, letting $z = y - \sigma\sqrt{T-t}$ and using the symmetry of the normal distribution leads to Eq. 175:

$$I(S_t, t) = \frac{S_t}{\sqrt{2\pi}} \int_{z=\alpha - \sigma\sqrt{T-t}}^{+\infty} e^{-\frac{z^2}{2}} dz = \frac{S_t}{\sqrt{2\pi}} \int_{z=-\infty}^{-\alpha + \sigma\sqrt{T-t}} e^{-\frac{z^2}{2}} dz \quad \text{Eq. 175}$$

Plugging the expression for α and using the normal cumulative distribution function N gives Eq. 176:

$$I(S_t, t) = \frac{S_t}{\sqrt{2\pi}} \int_{z=-\infty}^{\frac{\ln\left(\frac{S_t}{K}\right) + \left(r + \frac{\sigma^2}{2}\right)(T-t)}{\sigma\sqrt{T-t}}} e^{-\frac{z^2}{2}} dz = S_t N\left(\frac{\ln\left(\frac{S_t}{K}\right) + \left(r + \frac{\sigma^2}{2}\right)(T-t)}{\sigma\sqrt{T-t}}\right) \quad \text{Eq. 176}$$

With K being constant, rearranging the second integral J and using the symmetry of the normal distribution yields Eq. 177 which, after simplification, leads to Eq. 178.

$$J(S_t, t) = \frac{Ke^{-r(T-t)}}{\sqrt{2\pi}} \int_{y=\alpha}^{+\infty} e^{-\frac{y^2}{2}} dy = \frac{Ke^{-r(T-t)}}{\sqrt{2\pi}} \int_{y=-\infty}^{\frac{\ln\left(\frac{S_t}{K}\right) + \left(r - \frac{\sigma^2}{2}\right)(T-t)}{\sigma\sqrt{T-t}}} e^{-\frac{y^2}{2}} dy \quad \text{Eq. 177}$$

$$J(S_t, t) = Ke^{-r(T-t)} \cdot N\left(\frac{\ln\left(\frac{S_t}{K}\right) + \left(r - \frac{\sigma^2}{2}\right)(T-t)}{\sigma\sqrt{T-t}}\right) \quad \text{Eq. 178}$$

This finally leads to the Black-Scholes pricing formula for a European call option shown in Eq. 179:

$$V(S_t, t) = S_t \cdot N(d_1(t)) - Ke^{-r(T-t)}N(d_2(t))$$

With:

$$d_1(t) = \frac{\ln\left(\frac{S_t}{K}\right) + \left(r + \frac{\sigma^2}{2}\right)(T-t)}{\sigma\sqrt{T-t}} \quad \text{and} \quad d_2(t) = \frac{\ln\left(\frac{S_t}{K}\right) + \left(r - \frac{\sigma^2}{2}\right)(T-t)}{\sigma\sqrt{T-t}} \quad \text{Eq. 179}$$

APPENDIX H: IMPLEMENTATION OF REAL OPTIONS METHOD

The purpose of the implementation is to develop an environment to verify, validate, and finally apply the proposed methodology. In this section, several aspects of the implementation are described. The first aspect consists in presenting the architecture retained for the implementation of the proposed methodology while substantiating choices made. The second aspect consists in presenting the set of tools that are used to verify the different steps of the proposed methodology and in discussing the structure retained and the choices made.

Implementation Environment

Terminology

Due to the analogies between real options and financial options, the program used to assess the value of real options is called a *pricer* or an *option pricer* in the following paragraphs.

Language selection

There are several venues for the implementation of a real option-based program evaluation calculator. It is customary, in the financial engineering industry, to use C, C++ or Java to code object-oriented option pricing algorithms [301] [302]. This enables the use of pre-existing libraries and ensures a fast execution time which can be important, especially for algorithmic “nanosecond” trading. Extremely fast runtimes and the availability of pre-existing financial engineering libraries are nevertheless not primary

requirements for this research. Indeed, most analyses will only be run a limited number of times and the fundamentally new methodology proposed in this research will only minimally leverage the availability of pre-existing libraries.

However, one significant goal of this research is to show that a real options approach can be used and implemented with some relative ease by people with various backgrounds (corporate finance, management, engineering, etc.). In this context, a widely-used and ubiquitous programming language found at most workstations within companies would be ideal. Consequently, a spreadsheet type of environment augmented by object-oriented Visual Basic for Application (VBA) programming is appropriate. Microsoft Excel¹ has been retained to perform most computations in this research. If the need for faster execution speed were to arise, some computationally intensive routines could always be translated and coded as Excel add-ins in the C++ language.

Developing the Real Options Toolbox

Architecting the implementation

The implementation is articulated around an Excel Spreadsheet interface where all the interactions between the user and the FLAVIA program take place. The inputs include the description of the development program to investigate, the description of the real option to investigate, the description of the stochastic processes driving the uncertainties, and finally the description of some technical parameters used by the real options pricing program. Two screenshots of the user interface are displayed in Figure 131 and Figure 132: the first one shows the interface used for the specification of technical parameters

¹ Excel is developed by the Microsoft Corporation, Redmond, WA

(such as the number of Monte Carlo replications and time steps and the convergence criteria for solvers) while the second one shows the interface used to define the stochastic processes modeling the evolution of uncertainties and to describe the staggered development timeline. The outputs include an estimate of the real option price as well as an approximation of the early-exercise boundary.

| Modeling Parameters | |
|--|--------|
| Time Step Nbr | 180 |
| Nested Time Step Nbr | 1 |
| Monte Carlo Run | 30000 |
| Nested Monte Carlo Run | 30000 |
| Resampled Monte Carlo Run | 30000 |
| Nested Resampled Monte Carlo Run | 30000 |
| Trigger Time Monte Carlo Run | 4095 |
| Nested Trigger Time Monte Carlo Run | 4095 |
| Resampling Pool Size | 1 |
| Nested Resampling Pool Size | 1 |
| Esscher Solver Convergence Criteria | 1E-11 |
| Esscher Solver Max iterations | 100 |
| LSMC Interpolation Basis Size | 5 |
| Critical Price Solver Convergence Criteria | 1E-04 |
| Critical Price Solver Max Iteration | 25 |
| Refined Simulation Starting Points Nbr | 50 |
| Boundary Range Factor (Bounded Side) | 0.85 |
| Boundary Range Factor (Unbounded Side) | 4 |
| Natural Boundary Convergence Criteria | 0.0001 |
| Exercise Boundary Regression Basis Size | 2 |
| Quasi Monte Carlo Analysis | FALSE |
| Control Variate and Moment Correction | TRUE |
| Multi Start Simulation | TRUE |
| Critical Price Outlier Removal | TRUE |
| Compute Trigger Time Expectation | TRUE |
| Display Intermediate Graph | TRUE |

| Development Program Parameters | |
|--------------------------------|--------|
| Uncertainty Number | 1 |
| Milestone Number | 2 |
| Program Length (year) | 1 |
| Yearly Market Size | 200 |
| Program Shrink per Year | 4.00% |
| Profit per Sale | 0.000 |
| Risk Free Rate | 5.00% |
| WACC | 13.50% |

Figure 131: FLAVIA interface for technical parameters specification

| Development Program Timeline | | | | |
|-------------------------------------|-----------------------|---------------|-------------------------|------------|
| Decision Milestone | Pre-Conceptual Design | Detail Design | Testing & Certification | Production |
| Phase Decision Window | 1.00 | 1.00 | | |
| Phase Time Length | 1.00 | 1.00 | | |
| Phase Investment | 1.00 | 1.20 | | |
| Investment Decision (Invest / Sell) | INVEST | INVEST | SELL | SELL |
| Type of Flexibility | AMERICAN | EUROPEAN | | |

| Stochastic Processes | 2 | 2 | 2 | |
|--|----------------|------------------------|----|--|
| Uncertainty | Fuel Price | CO2 Permit Price | 2 | |
| Stochastic Process | JUMP DIFFUSION | JUMP DIFF / COR. JUMPS | BM | |
| Current Price (US\$) | 2 | 100 | | |
| Long-Term Drift (yr) | 5.00% | 20.00% | | |
| Dividend Rate (yr) | 4.00% | 0.00% | | |
| Std Deviation (yr) | 20.00% | 44% | | |
| Jump Rate | 100.00% | 100.00% | | |
| Jump Size Average | -8.00% | -2.00% | | |
| Jump Size Volatility | 40.00% | 20.00% | | |
| Diffusion Correlation First Uncertainty | 1.00 | 0.60 | | |
| Diffusion Correlation Second Uncertainty | 0.60 | 1.00 | | |
| Diffusion Correlation Third Uncertainty | -0.20 | -0.50 | | |
| Diffusion Correlation Fourth Uncertainty | | | | |
| Diffusion Correlation Fifth Uncertainty | | | | |
| Diffusion Correlation Sixth Uncertainty | | | | |
| Jump Size Correlation First Uncertainty | 1.00 | 0.80 | | |
| Jump Size Correlation Second Uncertainty | 0.80 | 1.00 | | |
| Jump Size Correlation Third Uncertainty | -0.50 | -0.20 | | |
| Jump Size Correlation Fourth Uncertainty | | | | |
| Jump Size Correlation Fifth Uncertainty | | | | |
| Jump Size Correlation Sixth Uncertainty | | | | |

Figure 132: FLAVIA interface for uncertainties and program timeline specifications

The program itself is entirely written in VBA using an object-oriented logic. Interaction between the spreadsheet and the VBA code is kept to a minimum to increase execution speed. The architecture of the program is articulated around seven different classes which are highlighted in Figure 133 and described in the following paragraphs.

The first class “Development Program” deals with all the computations pertaining to the research and development program such as technology modeling, aircraft operating cost estimation, market preference estimation, and profit estimation. For verification purposes, the option pricer is tested using standard financial options – not real options – so that option prices from other sources can be used to check results from the option pricer. Therefore an additional subroutine called “Basic value estimation” is added for verification purposes. Its goal is to bypass all computations pertaining to a research and development business plan.

The second class “Random Numbers” deals with the generation of random numbers according to different types of probability distributions. Sampling for uniform distributions, normal distributions, triangular distributions, as well as Poisson distributions are implemented in the class. Uniformly distributed random numbers are generated using either a uniform number generator or a low discrepancy sequence such as the Halton [187] and Sobol [189] sequences. Standard normally distributed random numbers are generated with the Box-Muller [303] transform using a source of uniformly distributed random numbers. For multivariate distributions, correlations can be accounted for using the Cholesky factorization of the correlation matrix [131]. The correlation matrix represents the correlation structure between the different dimensions.

The third class “Stochastic Process” deals with the simulation of stochastic processes. Several types of stochastic processes can be simulated including arithmetic Brownian motions, geometric Brownian motions, multi-dimensional geometric Brownian motions, as well as jump diffusion processes. The stochastic processes require “time-series” of random numbers to model the innovation terms and these are obtained using the random number class described previously. When correlation exists between different stochastic processes, correlated time-series of random numbers are used for the innovation factors.

The fourth class “Risk Neutralization” deals with the risk-neutralization of stochastic processes simulated under the physical probability measure. All computations pertaining to the Esscher transforms are performed using subroutines in this class. This includes the computation of the empirical moment generating function, the computation of the Esscher transform, as well as the estimation of the Esscher parameter using either a bisection solver or a Newton-Raphson solver.

The fifth class “Resampling” deals with the resampling of the weighted risk-neutral distribution to obtain a uniformly distributed discrete approximation of the risk-neutral distribution. The basic bootstrap subroutine and a faster, more efficient bootstrap subroutine belong to this class.

The sixth class “Option Type” deals with the algorithms used to evaluate options. The algorithms to price European options as well as path-dependent Bermudan and American options are implemented in this class. For the analysis of path-dependent options, the Longstaff-Schwartz method is used and therefore some regression routines (polynomial basis functions) are implemented to estimate conditional expectations.

The seventh class “Option Payoff” deals with the computations pertaining to the evaluation of the option payoff. Option is left undefined as this class can handle many different types of payoffs, stemming from both the analysis of real options and financial options. These include the usual vanilla put and call payoffs.

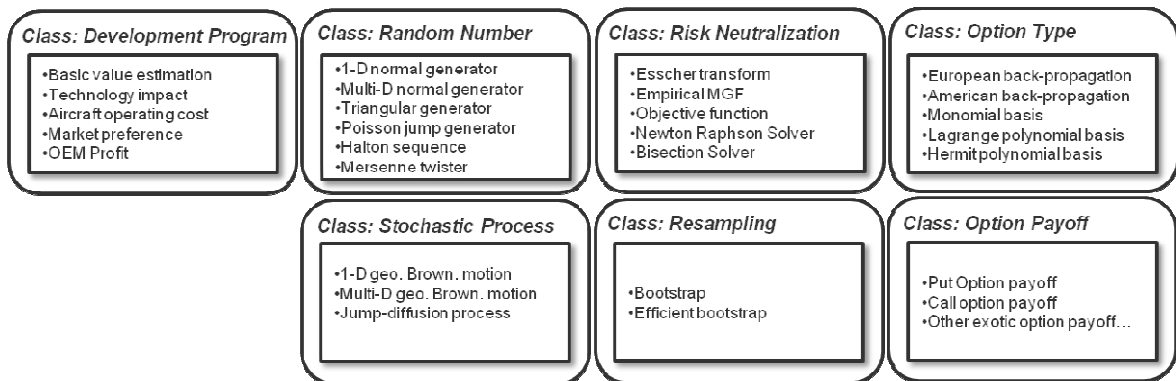


Figure 133: Classes implementation for the real options tool

APPENDIX I: SOLVING PARTIAL DIFFERENTIAL EQUATIONS WITH FINITE-DIFFERENCE SCHEMES

Finite-difference methods

Finite-difference methods are widely used numerical schemes that enable the pricing of European as well as Bermudan and American options. Finite-difference methods were first proposed by Schwartz [176] and Brennan and Schwartz [177] [178] to solve the Black-Scholes partial differential equation. Finite-difference methods present a means to obtain numerical solutions to partial differential equations by discretizing the time and asset-price space into a mesh of evenly distributed nodes and then approximating partial derivatives as finite differences at these nodes. Two popular ways to express these finite differences are the forward difference and the backward difference leading to respectively an explicit or an implicit finite-difference scheme. The difference is highlighted in Figure 134.

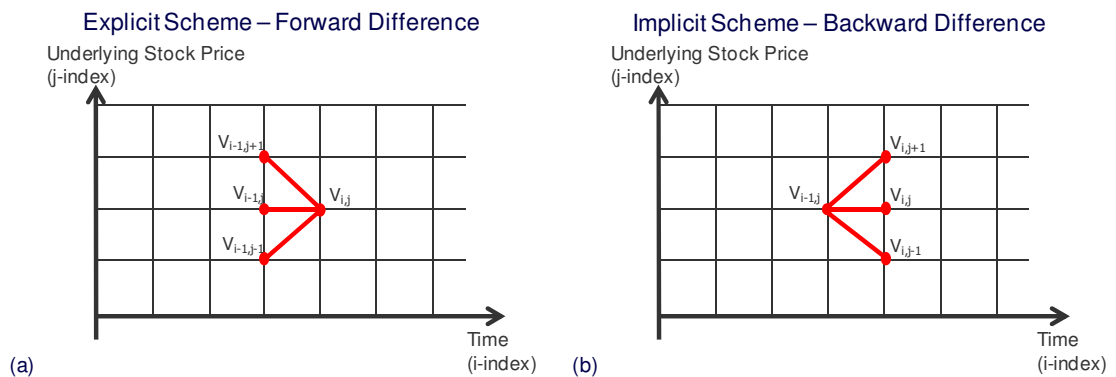


Figure 134: Forward (a) and backward (b) finite-difference for the explicit and implicit schemes

Because of potential numerical instabilities with explicit schemes, the implicit numerical scheme is used and the Black-Scholes partial differential equation approximation is expressed in Eq. 180.

$$\begin{aligned} \frac{V(i, j) - V(i - 1, j)}{\Delta t} + (r - q)j\Delta S \frac{V(i, j + 1) - V(i, j - 1)}{2\Delta S} \\ + \frac{\sigma^2}{2} j^2 (\Delta S)^2 \frac{V(i, j + 1) - 2V(i, j) + V(i, j - 1)}{(\Delta S)^2} = rV(i, j) \end{aligned} \quad \text{Eq. 180}$$

Boundary conditions

Boundary conditions at the extremities of the mesh enable the estimation of the option price which is then propagated throughout the mesh using the finite-difference approximation of the partial differential equation. The boundary conditions are usually set at the maturity of the option (because the option payoff is usually known at maturity), for extremely large value of the underlying asset (because the option payoff can be approximated for these large values), and for extremely small value of the underlying asset (again, because the option payoff can be approximated for these small values). For an American call option, these boundary conditions are expressed as shown in Eq. 181:

$$\begin{aligned} V(S_T, T) &= \max(S_T - K, 0) \\ \lim_{S \rightarrow +\infty} V(S_{t_k}, t_k) &= S_{t_k} \cdot e^{-qt_k} - K \cdot e^{-rt_k} \\ \lim_{S \rightarrow 0} V(S_{t_k}, t_k) &= 0 \end{aligned} \quad \text{Eq. 181}$$

Propagation in the mesh

Back-propagating the option value using the implicit scheme at each node in the mesh requires solving a sequence of linear systems with identical coefficients as shown in Eq. 182.

$$V(i + 1, j) = \alpha_j V(i, j - 1) + \beta_j V(i, j) + \gamma_j V(i + 1, j)$$

$$\text{With } \begin{cases} \alpha_j = \frac{1}{2} \Delta t ((r - q) \cdot j - \sigma^2 j^2) \\ \beta_j = 1 + \sigma^2 j^2 \Delta t + r \Delta t \\ \gamma_j = -\frac{1}{2} \Delta t ((r - q) \cdot j + \sigma^2 j^2) \end{cases} \quad \text{Eq. 182}$$

The sequence of linear systems with identical coefficients can be expressed in matrix format and this leads to the matrix equation of Eq. 183. Solving for the solution of these systems may be done by inverting the tri-diagonal coefficient matrix.

$$\begin{bmatrix} \beta_1 & \gamma_1 & 0 & \dots & 0 & 0 & 0 \\ \alpha_2 & \beta_2 & \gamma_2 & \dots & 0 & 0 & 0 \\ & \vdots & & \ddots & & \vdots & \\ 0 & 0 & 0 & \dots & \alpha_{N-2} & \beta_{N-2} & \gamma_{N-2} \\ 0 & 0 & 0 & \dots & 0 & \alpha_{N-1} & \beta_{N-1} \end{bmatrix} \begin{bmatrix} V(i, 1) \\ \vdots \\ V(i, N-1) \end{bmatrix} = \begin{bmatrix} V(i+1, 1) - \alpha_1 V(i, 0) \\ V(i+1, 2) \\ \vdots \\ V(i+1, N-2) \\ V(i+1, N-1) - \gamma_{N-1} V(i, N) \end{bmatrix} \quad \text{Eq. 183}$$

Nevertheless, this is accomplished most efficiently by factoring the tri-diagonal matrix into a lower and upper triangular parts and then solving each individual systems (L-U factorization [304]). Discussing in details the implementation of the implicit scheme for finite-difference methods is beyond the scope of this discussion and an interested reader is referred to the textbook of Wilmott, Howison, and Dewynne [305].

Early-exercise boundary

Another interesting aspect of solving a partial differential equation using a finite-difference scheme is the ability to directly generate the early-investment boundary.

Generating the early-investment boundary is done by checking if the early-exercise privilege is exercised at each and every node in the time and asset-price mesh. The boundary is approximated at each time cross-section by looking at neighboring nodes that have different exercise policies (i.e. the critical stock price is defined by two neighboring nodes, one with the option exercised early and one with the option kept open). This approximation is used to check the accuracy of the early-exercise boundary generated by the proposed approach.

APPENDIX J: IMPLEMENTATION VERIFICATION


Finite-difference scheme implementation: choosing the appropriate discretization

For real options featured in staggered research and development investments, a reasonable range for the maturity is probably from half a year up to two years while a reasonable range for the volatility is probably from twenty percent up to forty percent. Below the lower bounds, a real options methodology is probably not warranted as there is little uncertainty and little time to learn. Above the upper bounds, investments become extremely risky and probably unrealistic. With the time and space ranges defined, the granularity of the grid is then investigated to meet a given accuracy target. The accuracy is defined using the error criterion in Eq. 184 which represents the average error between the Black-Scholes price V_{BS} and the finite-difference price V_{FD} for European options. An error threshold of 10^{-4} is retained as it leads to an *average* accuracy up to the third significant digit for the valuation of an at-the-money option with unit strike.

$$\epsilon = \frac{1}{N} \sum_{i=1}^N |V_{BS}(S_i) - V_{FD}(S_i)| \quad \text{Eq. 184}$$

The graphs in exhibits (a), (b), (c), and (d) of Table 90 show the error criterion for different combinations of the number of time and space steps. These four graphs represent the lower and upper bounds for the maturity and volatility ranges and indicate that the longest maturities and highest volatilities are the most demanding cases. In addition, going beyond 500 steps in the space dimension yields little improvement and thus a discretization of 500 steps is retained for the space dimension.

Table 90: Selection of a time-space grid for European options with unit spot to strike ratios



Using a space discretization of 500 steps, the graph in Figure 135 describes how many time steps are required to meet the target accuracy for different maturities and volatilities. It appears that 400 time steps are sufficient to meet the accuracy target for volatilities up to 40% percent and maturities up to two years.

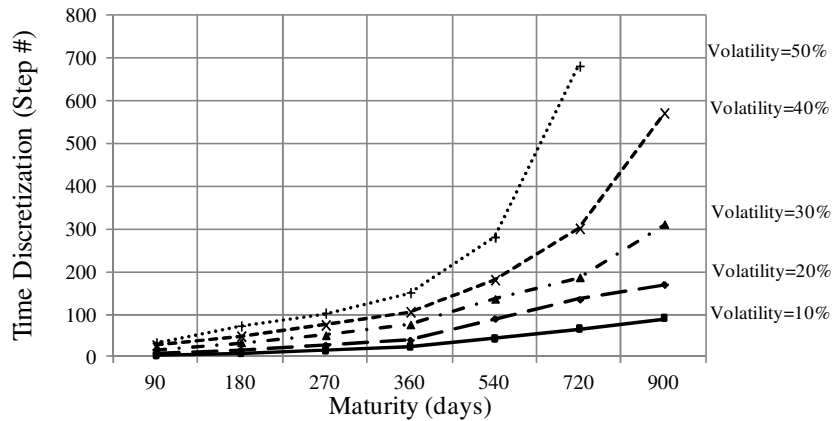


Figure 135: Required number of time steps to meet target accuracy

Finite-difference scheme implementation compared to Black-Scholes for option pricing

Comparing option prices computed using the Black-Scholes equation to those obtained with finite-difference schemes may be difficult due to the wide range of magnitudes for option prices. In particular, some option prices are very small and computing relative errors may prove misleading. Following Knuth [306], the relative difference d_r between two estimates V_{BS} and V_{FD} is retained. The relative difference is defined in Eq. 185 and used to characterize the magnitude of errors shown in Table 91 for notional European put and call options.

$$d_r = \frac{|V_{BS} - V_{FD}|}{\max(|V_{BS}|, |V_{FD}|)} \quad \text{Eq. 185}$$

Table 91: Comparison between finite-difference method and Black-Scholes for European put and call options

| Option Parameters ^(a) | | | | | European Call Option | | | European Put Option | | |
|----------------------------------|----------|-------|------|-----|----------------------|--------------------------|---------------------|---------------------|--------------------------|---------------------|
| Spot to Strike Ratio | σ | r_f | q | T | Black Scholes | Finite-Difference Method | Relative Difference | Black Scholes | Finite-Difference Method | Relative Difference |
| 0.8 | 20% | 2.0% | 0.0% | 180 | 3.62E-03 | 3.63E-03 | 0.22% | 1.94E-01 | 1.94E-01 | 0.00% |
| 0.8 | 20% | 2.0% | 0.0% | 360 | 1.43E-02 | 1.43E-02 | 0.02% | 1.94E-01 | 1.94E-01 | 0.00% |
| 0.8 | 20% | 2.0% | 0.0% | 720 | 3.85E-02 | 3.85E-02 | 0.02% | 1.99E-01 | 1.99E-01 | 0.00% |
| 0.8 | 20% | 2.0% | 4.0% | 180 | 2.58E-03 | 2.58E-03 | 0.16% | 2.08E-01 | 2.08E-01 | 0.00% |
| 0.8 | 20% | 2.0% | 4.0% | 360 | 9.40E-03 | 9.39E-03 | 0.03% | 2.21E-01 | 2.21E-01 | 0.00% |
| 0.8 | 20% | 2.0% | 4.0% | 720 | 2.25E-02 | 2.25E-02 | 0.08% | 2.45E-01 | 2.45E-01 | 0.01% |
| 0.8 | 20% | 8.0% | 0.0% | 180 | 5.70E-03 | 5.71E-03 | 0.28% | 1.66E-01 | 1.67E-01 | 0.01% |
| 0.8 | 20% | 8.0% | 0.0% | 360 | 2.38E-02 | 2.39E-02 | 0.06% | 1.47E-01 | 1.47E-01 | 0.02% |
| 0.8 | 20% | 8.0% | 0.0% | 720 | 6.91E-02 | 6.91E-02 | 0.00% | 1.21E-01 | 1.21E-01 | 0.02% |
| 0.8 | 20% | 8.0% | 4.0% | 180 | 4.14E-03 | 4.15E-03 | 0.23% | 1.81E-01 | 1.81E-01 | 0.01% |
| 0.8 | 20% | 8.0% | 4.0% | 360 | 1.64E-02 | 1.64E-02 | 0.02% | 1.71E-01 | 1.71E-01 | 0.01% |
| 0.8 | 20% | 8.0% | 4.0% | 720 | 4.37E-02 | 4.37E-02 | 0.04% | 1.57E-01 | 1.57E-01 | 0.00% |
| 0.8 | 40% | 2.0% | 0.0% | 180 | 3.26E-02 | 3.26E-02 | 0.03% | 2.23E-01 | 2.23E-01 | 0.01% |
| 0.8 | 40% | 2.0% | 0.0% | 360 | 6.85E-02 | 6.85E-02 | 0.04% | 2.49E-01 | 2.49E-01 | 0.01% |
| 0.8 | 40% | 2.0% | 0.0% | 720 | 1.26E-01 | 1.25E-01 | 0.05% | 2.86E-01 | 2.86E-01 | 0.02% |
| 0.8 | 40% | 2.0% | 4.0% | 180 | 2.85E-02 | 2.85E-02 | 0.05% | 2.34E-01 | 2.34E-01 | 0.01% |
| 0.8 | 40% | 2.0% | 4.0% | 360 | 5.72E-02 | 5.72E-02 | 0.06% | 2.69E-01 | 2.69E-01 | 0.01% |
| 0.8 | 40% | 2.0% | 4.0% | 720 | 9.75E-02 | 9.74E-02 | 0.06% | 3.20E-01 | 3.20E-01 | 0.02% |
| 0.8 | 40% | 8.0% | 0.0% | 180 | 3.85E-02 | 3.85E-02 | 0.02% | 1.99E-01 | 1.99E-01 | 0.00% |
| 0.8 | 40% | 8.0% | 0.0% | 360 | 8.35E-02 | 8.35E-02 | 0.04% | 2.07E-01 | 2.07E-01 | 0.01% |
| 0.8 | 40% | 8.0% | 0.0% | 720 | 1.59E-01 | 1.59E-01 | 0.04% | 2.11E-01 | 2.11E-01 | 0.02% |
| 0.8 | 40% | 8.0% | 4.0% | 180 | 3.38E-02 | 3.38E-02 | 0.04% | 2.10E-01 | 2.10E-01 | 0.01% |

Table 91 Continued

| | | | | | | | | | | |
|-----|-----|------|------|-----|----------|----------|-------|----------|----------|-------|
| 0.8 | 40% | 8.0% | 4.0% | 360 | 7.04E-02 | 7.04E-02 | 0.05% | 2.25E-01 | 2.25E-01 | 0.01% |
| 0.8 | 40% | 8.0% | 4.0% | 720 | 1.26E-01 | 1.26E-01 | 0.06% | 2.39E-01 | 2.39E-01 | 0.02% |
| 0.9 | 20% | 2.0% | 0.0% | 180 | 1.99E-02 | 1.99E-02 | 0.05% | 1.10E-01 | 1.10E-01 | 0.01% |
| 0.9 | 20% | 2.0% | 0.0% | 360 | 4.15E-02 | 4.15E-02 | 0.05% | 1.22E-01 | 1.22E-01 | 0.02% |
| 0.9 | 20% | 2.0% | 0.0% | 720 | 7.70E-02 | 7.70E-02 | 0.04% | 1.38E-01 | 1.38E-01 | 0.02% |
| 0.9 | 20% | 2.0% | 4.0% | 180 | 1.54E-02 | 1.54E-02 | 0.08% | 1.23E-01 | 1.23E-01 | 0.01% |
| 0.9 | 20% | 2.0% | 4.0% | 360 | 2.96E-02 | 2.96E-02 | 0.08% | 1.45E-01 | 1.45E-01 | 0.02% |
| 0.9 | 20% | 2.0% | 4.0% | 720 | 4.87E-02 | 4.86E-02 | 0.07% | 1.79E-01 | 1.79E-01 | 0.02% |
| 0.9 | 20% | 8.0% | 0.0% | 180 | 2.75E-02 | 2.75E-02 | 0.03% | 8.83E-02 | 8.83E-02 | 0.01% |
| 0.9 | 20% | 8.0% | 0.0% | 360 | 6.16E-02 | 6.16E-02 | 0.03% | 8.47E-02 | 8.47E-02 | 0.01% |
| 0.9 | 20% | 8.0% | 0.0% | 720 | 1.24E-01 | 1.24E-01 | 0.03% | 7.64E-02 | 7.64E-02 | 0.01% |
| 0.9 | 20% | 8.0% | 4.0% | 180 | 2.18E-02 | 2.18E-02 | 0.05% | 1.00E-01 | 1.00E-01 | 0.01% |
| 0.9 | 20% | 8.0% | 4.0% | 360 | 4.58E-02 | 4.57E-02 | 0.05% | 1.04E-01 | 1.04E-01 | 0.02% |
| 0.9 | 20% | 8.0% | 4.0% | 720 | 8.43E-02 | 8.43E-02 | 0.05% | 1.06E-01 | 1.06E-01 | 0.02% |
| 0.9 | 40% | 2.0% | 0.0% | 180 | 6.72E-02 | 6.72E-02 | 0.05% | 1.57E-01 | 1.57E-01 | 0.02% |
| 0.9 | 40% | 2.0% | 0.0% | 360 | 1.12E-01 | 1.12E-01 | 0.04% | 1.92E-01 | 1.92E-01 | 0.02% |
| 0.9 | 40% | 2.0% | 0.0% | 720 | 1.78E-01 | 1.78E-01 | 0.04% | 2.39E-01 | 2.39E-01 | 0.03% |
| 0.9 | 40% | 2.0% | 4.0% | 180 | 5.99E-02 | 5.99E-02 | 0.05% | 1.68E-01 | 1.68E-01 | 0.02% |
| 0.9 | 40% | 2.0% | 4.0% | 360 | 9.55E-02 | 9.55E-02 | 0.05% | 2.11E-01 | 2.11E-01 | 0.02% |
| 0.9 | 40% | 2.0% | 4.0% | 720 | 1.41E-01 | 1.41E-01 | 0.06% | 2.71E-01 | 2.71E-01 | 0.03% |
| 0.9 | 40% | 8.0% | 0.0% | 180 | 7.70E-02 | 7.70E-02 | 0.04% | 1.38E-01 | 1.38E-01 | 0.02% |
| 0.9 | 40% | 8.0% | 0.0% | 360 | 1.33E-01 | 1.33E-01 | 0.04% | 1.56E-01 | 1.56E-01 | 0.03% |
| 0.9 | 40% | 8.0% | 0.0% | 720 | 2.20E-01 | 2.20E-01 | 0.04% | 1.72E-01 | 1.72E-01 | 0.03% |
| 0.9 | 40% | 8.0% | 4.0% | 180 | 6.90E-02 | 6.89E-02 | 0.05% | 1.48E-01 | 1.48E-01 | 0.02% |
| 0.9 | 40% | 8.0% | 4.0% | 360 | 1.14E-01 | 1.14E-01 | 0.05% | 1.73E-01 | 1.73E-01 | 0.03% |
| 0.9 | 40% | 8.0% | 4.0% | 720 | 1.77E-01 | 1.77E-01 | 0.05% | 1.98E-01 | 1.98E-01 | 0.04% |
| 1 | 20% | 2.0% | 0.0% | 180 | 6.12E-02 | 6.12E-02 | 0.04% | 5.13E-02 | 5.12E-02 | 0.05% |
| 1 | 20% | 2.0% | 0.0% | 360 | 8.92E-02 | 8.91E-02 | 0.04% | 6.94E-02 | 6.93E-02 | 0.05% |
| 1 | 20% | 2.0% | 0.0% | 720 | 1.31E-01 | 1.31E-01 | 0.03% | 9.17E-02 | 9.17E-02 | 0.04% |
| 1 | 20% | 2.0% | 4.0% | 180 | 5.07E-02 | 5.07E-02 | 0.05% | 6.06E-02 | 6.06E-02 | 0.04% |
| 1 | 20% | 2.0% | 4.0% | 360 | 6.80E-02 | 6.80E-02 | 0.05% | 8.74E-02 | 8.74E-02 | 0.04% |
| 1 | 20% | 2.0% | 4.0% | 720 | 8.81E-02 | 8.81E-02 | 0.05% | 1.26E-01 | 1.26E-01 | 0.04% |
| 1 | 20% | 8.0% | 0.0% | 180 | 7.71E-02 | 7.70E-02 | 0.03% | 3.79E-02 | 3.78E-02 | 0.07% |
| 1 | 20% | 8.0% | 0.0% | 360 | 1.21E-01 | 1.21E-01 | 0.03% | 4.42E-02 | 4.41E-02 | 0.06% |
| 1 | 20% | 8.0% | 0.0% | 720 | 1.94E-01 | 1.94E-01 | 0.03% | 4.63E-02 | 4.63E-02 | 0.05% |
| 1 | 20% | 8.0% | 4.0% | 180 | 6.50E-02 | 6.49E-02 | 0.04% | 4.55E-02 | 4.55E-02 | 0.06% |
| 1 | 20% | 8.0% | 4.0% | 360 | 9.54E-02 | 9.53E-02 | 0.04% | 5.77E-02 | 5.77E-02 | 0.05% |
| 1 | 20% | 8.0% | 4.0% | 720 | 1.39E-01 | 1.39E-01 | 0.04% | 6.83E-02 | 6.82E-02 | 0.05% |
| 1 | 40% | 2.0% | 0.0% | 180 | 1.17E-01 | 1.17E-01 | 0.03% | 1.07E-01 | 1.07E-01 | 0.04% |
| 1 | 40% | 2.0% | 0.0% | 360 | 1.67E-01 | 1.67E-01 | 0.03% | 1.47E-01 | 1.47E-01 | 0.04% |
| 1 | 40% | 2.0% | 0.0% | 720 | 2.38E-01 | 2.38E-01 | 0.04% | 1.99E-01 | 1.99E-01 | 0.04% |
| 1 | 40% | 2.0% | 4.0% | 180 | 1.06E-01 | 1.06E-01 | 0.04% | 1.16E-01 | 1.16E-01 | 0.04% |
| 1 | 40% | 2.0% | 4.0% | 360 | 1.44E-01 | 1.44E-01 | 0.04% | 1.64E-01 | 1.64E-01 | 0.04% |
| 1 | 40% | 2.0% | 4.0% | 720 | 1.91E-01 | 1.91E-01 | 0.05% | 2.29E-01 | 2.29E-01 | 0.04% |
| 1 | 40% | 8.0% | 0.0% | 180 | 1.31E-01 | 1.31E-01 | 0.03% | 9.17E-02 | 9.17E-02 | 0.04% |
| 1 | 40% | 8.0% | 0.0% | 360 | 1.94E-01 | 1.94E-01 | 0.03% | 1.17E-01 | 1.17E-01 | 0.04% |
| 1 | 40% | 8.0% | 0.0% | 720 | 2.88E-01 | 2.88E-01 | 0.03% | 1.40E-01 | 1.40E-01 | 0.05% |
| 1 | 40% | 8.0% | 4.0% | 180 | 1.19E-01 | 1.19E-01 | 0.04% | 9.97E-02 | 9.97E-02 | 0.04% |
| 1 | 40% | 8.0% | 4.0% | 360 | 1.69E-01 | 1.69E-01 | 0.04% | 1.31E-01 | 1.31E-01 | 0.04% |

Table 91 Continued

| | | | | | | | | | | |
|-----|-----|------|------|-----|----------|----------|-------|----------|----------|-------|
| 1 | 40% | 8.0% | 4.0% | 720 | 2.35E-01 | 2.35E-01 | 0.05% | 1.64E-01 | 1.64E-01 | 0.05% |
| 1.1 | 20% | 2.0% | 0.0% | 180 | 1.29E-01 | 1.29E-01 | 0.01% | 1.95E-02 | 1.95E-02 | 0.08% |
| 1.1 | 20% | 2.0% | 0.0% | 360 | 1.56E-01 | 1.56E-01 | 0.02% | 3.63E-02 | 3.63E-02 | 0.07% |
| 1.1 | 20% | 2.0% | 0.0% | 720 | 1.98E-01 | 1.98E-01 | 0.02% | 5.92E-02 | 5.91E-02 | 0.06% |
| 1.1 | 20% | 2.0% | 4.0% | 180 | 1.13E-01 | 1.13E-01 | 0.01% | 2.45E-02 | 2.45E-02 | 0.06% |
| 1.1 | 20% | 2.0% | 4.0% | 360 | 1.25E-01 | 1.25E-01 | 0.02% | 4.85E-02 | 4.84E-02 | 0.05% |
| 1.1 | 20% | 2.0% | 4.0% | 720 | 1.41E-01 | 1.41E-01 | 0.02% | 8.59E-02 | 8.59E-02 | 0.05% |
| 1.1 | 20% | 8.0% | 0.0% | 180 | 1.52E-01 | 1.52E-01 | 0.01% | 1.31E-02 | 1.31E-02 | 0.12% |
| 1.1 | 20% | 8.0% | 0.0% | 360 | 1.98E-01 | 1.98E-01 | 0.02% | 2.11E-02 | 2.10E-02 | 0.11% |
| 1.1 | 20% | 8.0% | 0.0% | 720 | 2.75E-01 | 2.75E-01 | 0.02% | 2.72E-02 | 2.72E-02 | 0.10% |
| 1.1 | 20% | 8.0% | 4.0% | 180 | 1.34E-01 | 1.34E-01 | 0.01% | 1.68E-02 | 1.68E-02 | 0.09% |
| 1.1 | 20% | 8.0% | 4.0% | 360 | 1.63E-01 | 1.63E-01 | 0.02% | 2.93E-02 | 2.93E-02 | 0.09% |
| 1.1 | 20% | 8.0% | 4.0% | 720 | 2.06E-01 | 2.06E-01 | 0.03% | 4.27E-02 | 4.27E-02 | 0.09% |
| 1.1 | 40% | 2.0% | 0.0% | 180 | 1.80E-01 | 1.80E-01 | 0.02% | 7.05E-02 | 7.05E-02 | 0.05% |
| 1.1 | 40% | 2.0% | 0.0% | 360 | 2.31E-01 | 2.31E-01 | 0.02% | 1.12E-01 | 1.12E-01 | 0.05% |
| 1.1 | 40% | 2.0% | 0.0% | 720 | 3.05E-01 | 3.05E-01 | 0.03% | 1.66E-01 | 1.66E-01 | 0.06% |
| 1.1 | 40% | 2.0% | 4.0% | 180 | 1.66E-01 | 1.66E-01 | 0.02% | 7.74E-02 | 7.74E-02 | 0.05% |
| 1.1 | 40% | 2.0% | 4.0% | 360 | 2.03E-01 | 2.03E-01 | 0.03% | 1.26E-01 | 1.26E-01 | 0.04% |
| 1.1 | 40% | 2.0% | 4.0% | 720 | 2.48E-01 | 2.48E-01 | 0.04% | 1.94E-01 | 1.94E-01 | 0.05% |
| 1.1 | 40% | 8.0% | 0.0% | 180 | 1.98E-01 | 1.98E-01 | 0.02% | 5.92E-02 | 5.91E-02 | 0.06% |
| 1.1 | 40% | 8.0% | 0.0% | 360 | 2.64E-01 | 2.64E-01 | 0.02% | 8.68E-02 | 8.68E-02 | 0.06% |
| 1.1 | 40% | 8.0% | 0.0% | 720 | 3.62E-01 | 3.62E-01 | 0.03% | 1.14E-01 | 1.14E-01 | 0.07% |
| 1.1 | 40% | 8.0% | 4.0% | 180 | 1.83E-01 | 1.83E-01 | 0.02% | 6.53E-02 | 6.52E-02 | 0.06% |
| 1.1 | 40% | 8.0% | 4.0% | 360 | 2.33E-01 | 2.33E-01 | 0.03% | 9.88E-02 | 9.88E-02 | 0.06% |
| 1.1 | 40% | 8.0% | 4.0% | 720 | 2.99E-01 | 2.99E-01 | 0.04% | 1.36E-01 | 1.36E-01 | 0.07% |
| 1.2 | 20% | 2.0% | 0.0% | 180 | 2.16E-01 | 2.16E-01 | 0.00% | 6.16E-03 | 6.16E-03 | 0.03% |
| 1.2 | 20% | 2.0% | 0.0% | 360 | 2.37E-01 | 2.37E-01 | 0.00% | 1.76E-02 | 1.76E-02 | 0.05% |
| 1.2 | 20% | 2.0% | 0.0% | 720 | 2.76E-01 | 2.76E-01 | 0.01% | 3.72E-02 | 3.72E-02 | 0.07% |
| 1.2 | 20% | 2.0% | 4.0% | 180 | 1.94E-01 | 1.94E-01 | 0.00% | 8.23E-03 | 8.23E-03 | 0.07% |
| 1.2 | 20% | 2.0% | 4.0% | 360 | 1.98E-01 | 1.98E-01 | 0.00% | 2.50E-02 | 2.50E-02 | 0.02% |
| 1.2 | 20% | 2.0% | 4.0% | 720 | 2.04E-01 | 2.04E-01 | 0.01% | 5.71E-02 | 5.71E-02 | 0.04% |
| 1.2 | 20% | 8.0% | 0.0% | 180 | 2.43E-01 | 2.43E-01 | 0.00% | 3.76E-03 | 3.76E-03 | 0.05% |
| 1.2 | 20% | 8.0% | 0.0% | 360 | 2.86E-01 | 2.86E-01 | 0.01% | 9.32E-03 | 9.31E-03 | 0.13% |
| 1.2 | 20% | 8.0% | 0.0% | 720 | 3.63E-01 | 3.63E-01 | 0.01% | 1.56E-02 | 1.56E-02 | 0.14% |
| 1.2 | 20% | 8.0% | 4.0% | 180 | 2.21E-01 | 2.21E-01 | 0.00% | 5.14E-03 | 5.14E-03 | 0.00% |
| 1.2 | 20% | 8.0% | 4.0% | 360 | 2.44E-01 | 2.44E-01 | 0.01% | 1.38E-02 | 1.38E-02 | 0.09% |
| 1.2 | 20% | 8.0% | 4.0% | 720 | 2.82E-01 | 2.82E-01 | 0.02% | 2.61E-02 | 2.60E-02 | 0.11% |
| 1.2 | 40% | 2.0% | 0.0% | 180 | 2.55E-01 | 2.55E-01 | 0.01% | 4.53E-02 | 4.53E-02 | 0.05% |
| 1.2 | 40% | 2.0% | 0.0% | 360 | 3.04E-01 | 3.04E-01 | 0.01% | 8.42E-02 | 8.41E-02 | 0.05% |
| 1.2 | 40% | 2.0% | 0.0% | 720 | 3.78E-01 | 3.78E-01 | 0.03% | 1.39E-01 | 1.39E-01 | 0.07% |
| 1.2 | 40% | 2.0% | 4.0% | 180 | 2.37E-01 | 2.37E-01 | 0.01% | 5.04E-02 | 5.04E-02 | 0.04% |
| 1.2 | 40% | 2.0% | 4.0% | 360 | 2.69E-01 | 2.69E-01 | 0.02% | 9.62E-02 | 9.62E-02 | 0.05% |
| 1.2 | 40% | 2.0% | 4.0% | 720 | 3.11E-01 | 3.11E-01 | 0.03% | 1.64E-01 | 1.64E-01 | 0.07% |
| 1.2 | 40% | 8.0% | 0.0% | 180 | 2.76E-01 | 2.76E-01 | 0.01% | 3.72E-02 | 3.72E-02 | 0.07% |
| 1.2 | 40% | 8.0% | 0.0% | 360 | 3.41E-01 | 3.41E-01 | 0.02% | 6.41E-02 | 6.40E-02 | 0.07% |
| 1.2 | 40% | 8.0% | 0.0% | 720 | 4.41E-01 | 4.41E-01 | 0.03% | 9.36E-02 | 9.36E-02 | 0.10% |

Table 91 Continued

| | | | | | | | | | | |
|---|-----|------|------|-----|----------|----------|-------|----------|----------|-------|
| 1.2 | 40% | 8.0% | 4.0% | 180 | 2.57E-01 | 2.57E-01 | 0.01% | 4.16E-02 | 4.16E-02 | 0.06% |
| 1.2 | 40% | 8.0% | 4.0% | 360 | 3.04E-01 | 3.04E-01 | 0.02% | 7.40E-02 | 7.39E-02 | 0.07% |
| 1.2 | 40% | 8.0% | 4.0% | 720 | 3.68E-01 | 3.68E-01 | 0.03% | 1.13E-01 | 1.13E-01 | 0.09% |
| ^(a) Comparison is made with Black-Scholes formula with σ = standard deviation of returns; r_f = riskless rate of interest; q = dividend yield; and T = time to expiration | | | | | | | | | | |

APPENDIX K: PIP RESULTS

Impact of PIP-Light infusion on the engine operations in different market segments

| | Average Take-Off Derate | Severity Factor | Ease of Access to Capital Markets | | | Limited Access to Capital Markets | | |
|---------------------------|-------------------------|-----------------|-----------------------------------|-------------------------|-----------------------------|-----------------------------------|-------------------------|-----------------------------|
| | | | Fuel-burn Difference | CO2 Emission Difference | Maintenance Cost Difference | Fuel-Burn Difference | CO2 Emission Difference | Maintenance Cost Difference |
| Ops. Between 0-199 nm | 17% | 1.48 | 138,874 | 1,355,262 | 669,876 | 138,874 | 1,355,262 | 669,876 |
| Ops. between 200-399 nm | 15% | 1.21 | 184,987 | 1,805,269 | 144,158 | 184,987 | 1,805,269 | 144,158 |
| Ops. between 400-599 nm | 14% | 1.14 | 237,448 | 2,317,231 | 50,981 | 237,448 | 2,317,231 | 50,981 |
| Ops. between 600-799 nm | 14% | 1.08 | 234,833 | 2,291,717 | 16,935 | 234,833 | 2,291,717 | 16,935 |
| Ops. between 800-999 nm | 13% | 1.04 | 282,643 | 2,758,292 | 59,492 | 282,643 | 2,758,292 | 59,492 |
| Ops. between 1000-1199 nm | 12% | 1.04 | 330,449 | 3,224,826 | 109,158 | 330,449 | 3,224,826 | 109,158 |
| Ops. between 1200-1399 nm | 8% | 1.17 | 235,881 | 2,301,945 | 6,400 | 235,881 | 2,301,945 | 6,400 |
| Ops. between 1400-1599 nm | 5% | 1.33 | 274,755 | 2,681,307 | 10,520 | 274,755 | 2,681,307 | 10,520 |
| Ops. between 1600-1799 nm | 1% | 1.53 | 306,023 | 2,986,452 | 49,045 | 306,023 | 2,986,452 | 49,045 |
| Ops. between 1800-1999 nm | 0% | 1.57 | 335,760 | 3,276,649 | 51,830 | 335,760 | 3,276,649 | 51,830 |
| Ops. between 2000-3500 nm | 0% | 1.59 | 380,059 | 3,708,961 | 51,450 | 380,059 | 3,708,961 | 51,450 |

APPENDIX L: VERIFICATION DEFINITIONS

In this section, definitions of terms commonly used during the verification are provided. Emphasis is put on the specific circumstances in which these terms are used.

Probability distribution:

A probability distribution is a function that assigns a probability to each measurable subset of the outcomes of a random survey. In other terms, it is a statement about the frequency of outcomes. In the context of this thesis, the terms “probability distributions” and “probability measures” are used interchangeably.

Population distribution:

In the context of this thesis, population distributions or simply distributions are rarely used “as is” because they usually have infinite support. Instead, sampling from population distributions is performed. The term distribution refers to the entity (population and associated probability measure) from which a sample is drawn or from which a sample is assumed to be drawn.

Sample:

A sample is a subset of a population. In the context of this thesis, a sample is a set of independently and identically distributed experimental realizations. When performing Monte Carlo analyses, the set consisting of all Monte Carlo realizations constitutes a sample.

Empirical distribution:

The empirical probability distribution is the ratio of the number of times a specific outcomes occurred to the total number of trials performed in an experiment. In this research, empirical distributions estimate probabilities from experiences and observations and are constructed from samples usually generated via Monte Carlo simulations. The terms “empirical distribution” and “experimental distribution” are used interchangeably.

Test case:

A test case is a set of all relevant parameters used to perform a trial. In the context of this thesis, a test case is defined by a list of conditions and variables under which a trial is performed. For option pricing purposes, a test case may be defined by the type of stochastic process used as well as its parameterization. The terms “test case” and “test scenario” are used interchangeably.

Trial:

A trial or simply a test is a procedure that can be infinitely repeated and that corresponds to the examination of a single test case. In the context of this thesis, most trials are random in that they may have more than one possible outcome. The terms “trial” and “test” are used interchangeably.

Experiment:

In the context of thesis, an experiment is a collection of several independent trials performed to study the variability of their outcomes and to subject these outcomes to

statistical hypothesis testing. In this sense, an experiment is a composed experiment in which each individual repetition is a trial.

Statistical hypothesis testing:

Statistical hypothesis testing is a method of statistical inference used for testing a statistical hypothesis. Results from such tests are called statistically significant if they are predicted to be unlikely to occur by sampling error alone using a threshold probability called the significance level.

Null hypothesis:

The null hypothesis is a hypothesis being tested in an attempt to either disprove, reject or nullify it. During statistical hypothesis testing, the observed sample results are compared with the distribution under the null hypothesis and the likelihood of finding the obtained results is thereby determined. In the context of thesis, the null hypothesis is assumed to be true unless evidence indicates otherwise.

p-value:

During hypothesis testing, the p -value helps determine the significance of results: it is defined as the probability of finding the observed sample results, or more extreme results, when the null hypothesis is actually true. In the context of this thesis, p -values are computed to check how different an experimental empirical probability distribution is from an expected theoretical probability distribution.

Test statistic:

A test statistic is a measure of an attribute of a sample. In statistical hypothesis testing, the null hypothesis test is usually formulated in terms of a test statistic considered as a numerical summary of the data contained within the sample. In the context of this thesis, test statistics are constructed to perform hypothesis testing regarding the means of samples, the empirical distribution functions of samples, and the options prices.

REFERENCES

- [1] AIA, *Year end review and forecast*, 2012.
- [2] D. Raymer, *Aircraft design: a conceptual approach*, J. Przemieniecki, Ed., AIAA, 1999.
- [3] D. Schrage, "Technology for rotorcraft affordability through integrated product / process development (IPPD)," Montreal, Canada, 1999.
- [4] D. Schrage and J. E. Rogan, *The impact of concurrent engineering on aerospace systems design*, G. I. o. T. School of Aerospace Engineering, Ed., Atlanta, Georgia., 1991.
- [5] W. Marx, D. Mavris and D. Schrage, "Integrated product development for the wing structural design of the high speed civil transport," Panama City, FL, 1994.
- [6] W. Engler, P. Biltgen and D. Mavris, "Concept selection using an interactive reconfigurable matrix of alternatives (IRMA)," Reno, NV, 2007.
- [7] G. Taguchi, *Introduction to quality engineering*, Quality Resources, 1986.
- [8] D. Mavris and O. Bandte, "Economic uncertainty assessment using a combined design of experiments / Monte Carlo simulation approach with application to an," San Diego, CA, 1995.
- [9] P. Collopy and P. Hollingsworth, "Value-Driven Design," *Journal of Aircraft*, vol. 48, p. 749, May-June 2011.
- [10] P. Birtles, *Lockheed L1011 Tristar*, Motorbooks International Publishing Company, 1998.

- [11] C. Tang and J. Zimmerman, "Managing new product development and supply chain risks: the Boeing 787 case," *Supply Chain Forum: An international Journal*, vol. 10, pp. 74-85, 2009.
- [12] S. Trimble, *CFM revamps Leap assembly ahead of ramp-up*, 2015.
- [13] D. Vidalon, T. Hephher and A. Pal, *Seatmaker Zodiac faces American Airlines suit over delays*, 2015.
- [14] P. Nolte, A. Apffelstaedt and V. Gollnick, "Quantitative Assessment of Technology Impact on Aviation Fuel Efficiency," in *Air Transport and Operations*, 2012.
- [15] J. Moxon, *Airbus A380 Aircraft Profile*, 2007.
- [16] D. Gates, *Boeing celebrates 787 delivery as program's costs top \$32 billion*, 2011.
- [17] R. Marowitz, *Larger Bombardier C Series set for first flight next week in Mirabel*, 2015.
- [18] A. Rothman, *Airbus vows computers will speak same language after A380 delay*, 2006.
- [19] A. Levin and S. Ray, *Boeing 787 battery fix wins approval to resume flights*, 2013.
- [20] C. Justin and D. Mavris, "Option-based approach to value engine maintenance cost guarantees and engine maintenance contracts," ATIO Conference, Virginia Beach, 2011.
- [21] M. Liehr, A. Grossler and M. Klein, "Understanding business cycles in the airline market," International Conference of the System Dynamics Society, Wellington, New Zealand, 1999.
- [22] D. Gates, *Boeing 787's problems blamed on outsourcing, lack of oversight*, 2013.

- [23] N. Clark, *The Airbus saga: crossed wires and a multibillion-euro delay*, 2006.
- [24] A. Joyce, Volatility in the global airline industry is shaping its future evolution, vol. Air transportation in the 21st century, J. O'Connell and G. Williams, Eds., Ashgate Publishing Limited, 2011.
- [25] AIA, *Aerospace industry net profit after taxes*, 2012.
- [26] C. Justin, S. Briceño and D. Mavris, "Strategic analysis of competitive markets: a case study for the narrowbody market," ATIO Conference, Virginia Beach, 2011.
- [27] AIRBUS, *Global Market Forecast 2015-2034*, 2015.
- [28] T. Friedman, *The world is flat: a brief history of the twenty-first century*, 1 ed., Farrar, Strauss and Giroux, 2005.
- [29] Flightglobal, *BAe wins launch aid for Airbus A340-500/600*, 1998.
- [30] H. P. Ring, *Managing for long term growth - At the bottom of the aerospace cycle*, 2003.
- [31] D. Michaels, *The secret price of a jet airliner*, 2012.
- [32] Wikipedia, *Airbus A340*, 2013.
- [33] Wikipedia, *Boeing 777*, 2013.
- [34] EIA, *U.S. Gulf Coast kerosene-type jet fuel spot price*, 2013.
- [35] Flightglobal, *Airbus to offer cash back on A340 as 777 stretches lead*, 2006.
- [36] P. Samuelson, "Rational theory of warrant pricing," vol. 6, pp. 13-39, 1965.
- [37] C. Justin, S. Briceno and D. Mavris, "A competitive and real-options framework for the economic analysis of large aerospace programs," ICAS, Brisbane, QLD,

2012.

- [38] M. Babej and T. Pollack, *Boeing versus Airbus*, 2006.
- [39] AIRBUS, *Global Market Forecast 2001-2020*, 2002.
- [40] A. Doyle, *Airbus set to finalise risk sharing roster for A380*, 2002.
- [41] AIRBUS, *Airbus confirms further A380 delay and launches company restructuring plan*, 2006.
- [42] N. Forgeard, *Airbus presentation - Global Investor Forum 2003*, 2003.
- [43] D. Michaels, *Airbus on track to double profit margin by 2015*, 2013.
- [44] C. Justin, E. Garcia and D. Mavris, "Aircraft valuation: a network approach to the evaluation of aircraft for fleet planning and strategic decision making," ATIO Conference, Fort Worth, TX, 2010.
- [45] D. Shannon, *Airbus says A380F development 'interrupted' by UPS cancellation, but still on offer*, 2007.
- [46] AviationToday, *Emergence of 'A330-200 Lite' unlikely to impact existing values*, 2004.
- [47] M. Kingsley-Jones, *A cut above*, 2005.
- [48] M. Kingsley-Jones, *Airbus goes for extra width - A350 XWB special report*, 2006.
- [49] M. Kingsley-Jones, *Playing catch-up: no room for delays of the Airbus A350*, 2007.
- [50] M. Kingsley-Jones, *Airbus rolls out A350 XWB design revisions*, 2007.
- [51] A. Rothman, *Airbus A350 to be 'Hellish Ride' for Enders in bid for top job*, 2010.

- [52] M. H. Weingartner, "Some new views on the payback period and capital budgeting decisions," *Management Sciences*, vol. 15, pp. B-594, 1969.
- [53] R. Byrne, A. Charnes, W. Cooper and K. Kortanek, "A chance-constrained approach to capital budgeting with portfolio type payback and liquidity constraint and horizon posture controls," *Journal of Financial and Quantitative Analysis*, vol. 2, pp. 339-364, 1967.
- [54] I. Fisher, *The theory of interest*, Augustus, Kelley, 1930.
- [55] J. B. Williams, *The theory of investment value*, Harvard University Press, 1938.
- [56] T. Klamer, "Empirical evidence on the adoption of sophisticated capital budgeting techniques," *Journal of Business*, pp. 382-397, 1972.
- [57] J. Graham and C. Harvey, "The theory and practice of corporate finance: evidence from the field," *Journal of Financial Economics*, vol. 60, pp. 187-243, 2001.
- [58] P. A. Horvath, "Compounding and discounting in continuous time," *The Quarterly Review of Economics and Finance*, vol. 35, no. 3, pp. 315-325, 1995.
- [59] S. A. Ross, R. W. Westerfield and J. Jaffe, *Corporate finance*, Mcgraw Hill, 2003.
- [60] J. Treynor, "Towards a theory of market value of risky assets," 1962.
- [61] W. Sharpe, "Capital asset prices: a theory of market equilibrium under conditions of risk," *The Journal of Finance*, vol. 19, pp. 425-442, 1964.
- [62] J. Lintner, "The valuation of risky assets and the selection of risky investments in stock portfolios and capital budgets," *Review of Economics and Statistics*, vol. 47, pp. 13-37, 1965.
- [63] J. Mossin, "Equilibrium in a capital asset market," *Econometrica*, vol. 34, pp. 768-

783, 1966.

- [64] H. Markowitz, "Portfolio selection," *The Journal of Finance*, vol. 7, pp. 77-91, 1952.
- [65] Z. Bodie, A. Marcus and A. Kane, *Investments*, 6th ed., McGraw-Hill, 2004, pp. 146-449.
- [66] F. Black and M. Scholes, "The pricing of options and corporate liabilities," *Journal of Political Economy*, vol. 81, no. 3, pp. 637-654, May - Jun 1973.
- [67] R. C. Merton, "Theory of rational option pricing," *The Bell Journal of Economics and Management Science*, vol. 4, no. 1, pp. pp. 141-183, 1973.
- [68] E. G. Haug, *The complete guide to option pricing formulas*, 2nd ed., McGraw-Hill, 1997.
- [69] M. Porter, "How competitive forces shape strategy," *Harvard Business Review*, Vols. March-April, 1979.
- [70] A. A. Cournot, *Recherches sur les principes mathématiques de la théorie des richesses*, Hachette, 1838.
- [71] C. Darwin, *The descent of man, and selection in relation to sex*, John Murray, 1871.
- [72] J. V. Neumann and O. Morgenstern, *Theory of games and economic behavior*, Princeton University Press, 1944.
- [73] J. F. Nash, "Non-cooperative games," 1950.
- [74] D. Fudenberg and J. Tirole, *Game theory*, MIT Press, 1991.
- [75] C. Justin, S. Briceno, D. Mavris and F. Villeneuve, "A competitive market

- approach to gas turbine technology portfolio selection," in *Proceedings of the ASME Turbo Expo 2011, Volume 3*, 2011.
- [76] R. Gibbons, *Game theory for applied economists*, Princeton University Press, 1992.
- [77] M. Allais, "Le comportement de l'homme rationnel devant le risque, critique des postulats et axiomes de l'Ecole Americaine," *Econometrica*, vol. 21, pp. 503-546, 1953.
- [78] H. Simon, *Models of bounded rationality: economic analysis and public policy*, vol. Volume 1, MIT Press, 1982.
- [79] G. Gigerenzer and R. Selten, *Bounded rationality: the adaptive toolbox*, G. Gigerenzer and R. Selten, Eds., MIT Press, 2002.
- [80] D. Kahneman and A. Tversky, "Prospect theory: an analysis of decision under risk," *Econometrica*, vol. 47, pp. 263-292, 1979.
- [81] C. Krychowski, "Apport et limites des options réelles à la décision d'investissement stratégique : une étude appliquée dans le secteur des télécommunications," 2007.
- [82] S. C. Myers, "Determinants of corporate borrowing," *Journal of Financial Economics*, vol. 5, no. 2, pp. 147-175, 1977.
- [83] J. Stonier, "What is an aircraft purchase option worth? Quantifying asset flexibility created through manufacturer lead-time reductions and product commonality," *Handbook of Airline Finance*, 1999.
- [84] I. Fernandez-Martin, "Valuation of design adaptability in aerospace systems," 2008.

- [85] B. Miller and J.-P. Clarke, "Investments under uncertainty in air transportation," *Journal of the Transportation Research Forum*, vol. 44, pp. 61-74, Spring 2005.
- [86] V. Datar and S. Mathews, "European real options: an intuitive algorithm for the Black-Scholes formula," *Journal of Applied Finance*, vol. 14, p. 45, Spring 2004.
- [87] S. Mathews, "Valuing risky projects with real options," *Research-Technology Management*, vol. 52, 2009.
- [88] B. Miller and J.-P. Clarke, "Strategic guidance in the development of new aircraft programs: a practical real options approach," *IEEE Transaction on Engineering Management*, vol. 55, pp. 566-578, November 2008.
- [89] T. Luehrman, "Investment opportunities as real options: getting started on the numbers," *Harvard Business Review*, vol. Case 9, July-August 1998.
- [90] J. Shank and D. Peterson, "Strategic cost analysis for capital spending decisions," *Cost Management*, vol. 19, p. 14, 2005.
- [91] O. Pinon, "A methodology for the valuation and selection of adaptable technology portfolios and its application to small and medium airports," 2012.
- [92] J. Cox and S. Ross, "The valuation of options for alternative stochastic processes," *Journal of Financial Economics*, vol. 3, pp. 145-166, 1976.
- [93] M. Baxter and A. Rennie, *Financial calculus. An introduction to derivative pricing*, Cambridge University Press, 2003.
- [94] P. Boyle, "Options: a Monte Carlo approach," *Journal of Financial Economics*, vol. 4, pp. 323-338, 1977.
- [95] F. Jamshidian, "An exact bond option formula," *The Journal of Finance*, vol. 44,

pp. 205-209, 1989.

- [96] H. Geman, N. E. Karoui and J.-C. Rochet, "Change of numeraire, change of probability measure and option pricing," *Journal of Applied Probability*, vol. 32, pp. 443-458, 1995.
- [97] D. Duffie, *Dynamic asset pricing theory*, Princeton University Press, 2001.
- [98] P. Ouwehand and G. West, *Pricing rainbow options*, 2006, pp. 74-80.
- [99] J. Cox, S. Ross and M. Rubinstein, "Option pricing: a simplified approach," *Journal of Financial Economics*, vol. 7, pp. 229-263, 1979.
- [100] J. Andreasen, B. Jensen and R. Poulsen, "Eight valuation methods in financial mathematics: the Black-Scholes formula as example," *The Mathematical Scientist*, vol. 23, pp. 18-40, 1999.
- [101] M. Rubinstein, "Implied binomial tree," *The Journal of Finance*, vol. 69, pp. 771-818, 1994.
- [102] J. M. Harrison and D. M. Kreps, "Martingales and arbitrage in multiperiod securities markets," *Journal of Economic Theory*, vol. 20, pp. 281-408, 1979.
- [103] T. Copeland, T. Koller and J. Murrin, *Valuation. Measuring and managing the value of companies*, Third Edition ed., McKinsey, 2000.
- [104] T. Copeland and V. Antikarov, *Real options - a practitioner's guide*, Texere, 2001.
- [105] A. Borison, "Real options analysis: where are the Emperor's clothes?," *Journal of Applied Corporate Finance*, vol. 17, pp. 17-31, 2005.
- [106] T. Copeland and V. Antikarov, "Real options: Meeting the Georgetown challenge," *Journal of Applied Corporate Finance*, vol. 17, pp. 32-51, 2005.

- [107] A. Borison, "A response to "Real options: meeting the Georgetown challenge", "*Journal of Applied Corporate Finance*, vol. 17, pp. 52-54, Spring 2005.
- [108] P. Samuelson, "Proof that properly anticipated prices fluctuate randomly," *Industrial Management Review*, vol. 6, pp. 41-49, 1965.
- [109] E. Fama, "The behavior of stock-market prices," *Journal of Business*, vol. 38, pp. 34-105, 1965.
- [110] J.-C. Duan, "Conditionally fat-tailed distributions and the volatility smile in options," 1999.
- [111] A. Lo and C. MacKinlay, "Stock market prices do not follow random walks: evidence from simple specifications," *Review of Financial Studies*, vol. 1, pp. 41-66, 1988.
- [112] T. Bollerslev, "A conditionally heteroskedastic time series model for speculative prices and rates of return," *The Review of Economics and Statistics*, vol. 69, pp. 542-547, 1986.
- [113] A. Kemna, "Case studies on real options," *Financial Management*, vol. 22, pp. 259-270, 1993.
- [114] H. Weeds, "Strategic delays in a real options model of R&D competition," *The Review of Economic Studies*, vol. 69, pp. 729-747, 2002.
- [115] O. Pinon, E. Garcia and D. Mavris, *Evaluating flexibility in airport capacity-enhancing technology investments*, 2012.
- [116] R. McDonald and D. Siegel, "The value of waiting to invest," *The Quarterly Review of Economics and Finance*, pp. 707-727, 1986.

- [117] S. Grenadier, "Valuing lease contracts: a real-options approach," *Journal of Financial Economics*, vol. 38, pp. 297-331, 1995.
- [118] S. Majd and R. Pindyck, "Time to build, option value and investment decisions," *Journal of Financial Economics*, vol. 18, pp. 7-27, 1987.
- [119] A. Bernardo and B. Chowdry, "Resources, real options and corporate strategy," *Journal of Financial Economics*, vol. 63, pp. 211-234, 2002.
- [120] M. Dias and k. Rocha, "Petroleum concessions with extensible options using mean reversion with jumps to model oil prices," 1999.
- [121] L. Trigeorgis, "Evaluating leases with complex operating options," *European Journal of Operational Research*, vol. 91, pp. 315-329, 1996.
- [122] J. Stonier, "Airline fleet planning, financing, and hedging decisions under conditions of uncertainty," in *Handbook of Airline Strategy*, New York, McGraw-Hill, 2001.
- [123] N. Lewis, D. Enke and D. Spurlock, "Valuation for the strategic management of research and development projects: the deferral option," *Engineering Management Journal*, vol. 16, pp. 36-48, 2004.
- [124] W. Hahn and J. Dyer, "Discrete time modeling of mean-reverting stochastic processes for real option valuation," *European Journal of Operational Research*, vol. 184, pp. 534-548, 2008.
- [125] C. Bastian-Pinto, L. Brendao and W. Hahn, "Flexibility as a source of value in the production of alternative fuels: the ethanol case," *Energy Economics*, vol. 31, pp. 411-422, 2009.

- [126] N. Metropolis and S. Ulam, "The Monte Carlo method," *Journal of the American Statistical Association*, vol. 44, pp. 335-341, 1949.
- [127] M. Broadie and P. Glasserman, "Estimating security price derivatives using simulation," *Management Sciences*, vol. 42, pp. 269-285, 1996.
- [128] S. Rose, "Valuation of interacting real options in a tollroad infrastructure project," *The Quarterly Review of Economics and Finance*, vol. 38, pp. 711-723, 1998.
- [129] C.-L. Tseng and G. Barz, "Short-term generation asset valuation: a real options approach," *Operations Research*, vol. 50, pp. 297-310, 2002.
- [130] R. Rodriguez, "Real option valuation of free destination in long-term liquefied natural gas supplies," *Energy Economics*, vol. 30, pp. 1909-1932, 2008.
- [131] P. Glasserman, Monte Carlo methods in financial engineering, Rozovskii and Yor, Eds., Springer-Verlag, 2004.
- [132] M. Yang and W. Blyth, "Climate Policy Uncertainty and Investment Risk," IEA, 2007.
- [133] H. Bessembinder, J. Coughenour, P. Seguin and M. Smoller, "Mean reversion in equilibrium asset prices: evidence from the future term structure," *The Journal of Finance*, vol. 50, pp. 361-175, 1995.
- [134] R. C. Merton, "Option pricing when underlying stock returns are discontinuous," *Journal of Financial Economics*, vol. 3, pp. 125-144, 1976.
- [135] Y. Ait-Sahalia, "Disentangling diffusion from jumps," *Journal of Financial Economics*, vol. 74, pp. 487-528, 2004.
- [136] F. Longstaff and E. Schwartz, "Valuing American options by simulation: a simple

- least-square appr," *The Review of Financial Studies*, vol. 14, pp. 113-147, 2001.
- [137] S. Shreve, *Stochastic calculus for finance I The binomial asset pricing model*, Springer Finance, 2004.
- [138] S. Shreve, *Stochastic calculus for finance II Continuous-time models*, Springer Finance, 2004, pp. 34-35.
- [139] S. Neftci, *An introduction to the mathematics of financial derivatives*, Second Edition ed., Academic Press, 2000.
- [140] H. Gerber and E. Shiu, "Option pricing by Esscher transforms," *Transactions of Society of Actuaries*, vol. 46, pp. 99-141, 1994.
- [141] F. Esscher, "On the probability function in the collective theory of risk," *Scandinavian Actuarial Journal*, vol. 3, pp. 175-195, 1932.
- [142] P. M. Kahn, "An introduction to collective risk theory and its application to stop-loss reinsurance," *Transactions of Society of Actuaries*, vol. 14, pp. 400-449, 1962.
- [143] Y. Miyahara, "A note on Esscher transformed martingale measures for geometric Levy processes," 2004.
- [144] H. Follmer and M. Scheizer, *Encyclopedia of quantitative finance*, Wiley, 2010, pp. 1200-1204.
- [145] M. Frittelli, "Minimal entropy criterion for pricing in one period incomplete markets," 1995.
- [146] D. Cass and J. Stiglitz, "The structure of investor preferences and asset returns, and separability in portfolio allocation: a contribution to the pure theory of mutual funds," *Journal of Economic Theory*, vol. 2, pp. 122-160, 1970.

- [147] J. Stiglitz, "Theory of valuation: frontiers of modern financial theory," S. Bhattacharya and G. Constantinides, Eds., Rowman & Littlefield Publishers, 1988, pp. 342-356.
- [148] H. Buhlmann, F. Delbaen, P. Embrechts and A. Shiryaev, "No-arbitrage, change of measure and conditional Esscher transforms," *CWI Quarterly*, vol. 9, pp. 291-317, 1996.
- [149] J.-C. Duan, "The GARCH option pricing model," *Mathematical Finance*, vol. 5, pp. 13-32, January 1995.
- [150] T. K. Siu, H. Tong and H. Yang, "On pricing derivatives under GARCH models: a dynamic Gerber-Shiu's approach," *North American Actuarial Journal*, vol. 8, pp. 17-31, 2004.
- [151] H. Gerber and E. Shiu, "Martingale approach to pricing perpetual american options," *Astin Bulletin*, vol. 24, pp. 195-220, November 1994.
- [152] A. Klenke, "Probability Theory," Springer, 2008, pp. 205-215.
- [153] M. Goovaerts and R. Laeven, "Actuarial risk measures for financial derivative pricing," *Insurance: Mathematics and Economics*, vol. 42, pp. 540-547, 2007.
- [154] M. Pereira, C. Epprecht and A. Veiga, "Option pricing via nonparametric Esscher transform," Marseille, France, 2011.
- [155] J. Hull, *Options, futures and other derivatives securities*, 2nd in 1993, 6th ed., Prentice Hall, 2006.
- [156] J. A. Tilley, "Valuing American options in a path simulation model," *Transactions of the Society of Actuaries*, vol. 45, pp. 83-104, 1993.

- [157] J. Carriere, "Valuation of the early-exercise price for options using simulations and non-parametric regression," *Insurance: Mathematics and Economics*, vol. 19, pp. 19-30, 1996.
- [158] L. Stentoft, *Value function approximation or stopping time approximation: a comparison of two recent numerical methods for American option pricing using simulation and regression*, 2008.
- [159] B. Efron, "Bootstrap methods: another look at the jackknife," *The Annals of Statistics*, vol. 7, pp. 1-26, 1979.
- [160] J. Ostrower and P. Vieira, *Jet maker Bombardier finds bigger proves far from better*, 2015.
- [161] K. Forsberg and H. Mooz, "The relationship of system engineering to the project cycle," Chattanooga, TN, 1991.
- [162] D. O. Defense, *Systems engineering fundamentals*, Defense Acquisition University Press, 2001.
- [163] H. L. Harter, "Another look at plotting positions," *Communications in Statistics - Theory and Methods*, vol. 13, no. 13, pp. 1613-1633, 1984.
- [164] L. Makkonen, "Bringing closure to the plotting position controversy," *Communications in Statistics - Theory and Methods*, vol. 37, no. 3, pp. 460-467, 2008.
- [165] N. J. Cook, "Rebuttal of "Problems in the extreme value analysis", " *Structural Safety*, vol. 34, no. 1, pp. 418-426, 1 2012.
- [166] P. Wessa, *Free Statistics Software*, 2015.

- [167] R. Walpole, R. Myers, S. Myers and K. Ye, Probability and Statistics for Engineers and Scientists, 8 ed., Upper Saddle River, NJ: Pearson Prentice Hall, 2007.
- [168] Student, "The probable error of a mean," *Biometrika*, vol. 6, no. 1, pp. 1-25, 3 1908.
- [169] R. C. Veltkamp and M. Hagedoorn, State of the art in shape matching, Springer, 2001.
- [170] M. Hagedoorn, "Pattern matching using similarity measures," 2000.
- [171] H. Alt, C. Knauer and C. Wenk, "Comparison of distance measures for planar curves," *Algorithmica*, vol. 38, pp. 45-58, 2004.
- [172] G. Rote, "Curve with increasing chords," *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 115, no. 1, pp. 1-12, 1 1994.
- [173] R. Geske and H. Johnson, "The American put option valued analytically," *The Journal of Finance*, vol. 39, no. 5, pp. 1511-1524, 12 1984.
- [174] I. J. Kim, "The analytic valuation of American options," *Review of Financial Studies*, vol. 3, no. 4, pp. 547-572, 1990.
- [175] D. J. Higham, An introduction to financial option valuation: mathematics, stochastics and computation, Cambridge University Press, 2004.
- [176] E. S. Schwartz, "The valuation of warrants: implementing a new approach," *Journal of Financial Economics*, vol. 4, pp. 79-93, 1977.
- [177] M. J. Brennan and E. S. Schwartz, "The valuation of American put options," *The Journal of Finance*, vol. 32, no. 2, pp. 449-462, May 1977.
- [178] M. J. Brennan and E. S. Schwartz, "Finite difference methods and jump processes

arising in the pricing of contingent claims: a synthesis," *The Journal of Financial and Quantitative Analysis*, vol. 13, no. 3, pp. 461-474, Sep 1978.

- [179] R. Kangro and R. Nicolaidis, "Far field boundary conditions for Black-Scholes equations," *SIAM Journal on Numerical Analysis*, vol. 38, pp. 1357-1368, 1 2000.
- [180] H. Windcliff, P. A. Forsyth and K. R. Vetzal, "Analysis of the stability of the linear boundary condition For the Black-Scholes equation," *Journal of Computational Finance*, vol. 8, pp. 65-92, 2003.
- [181] G. Barone-Adesi and R. E. Whaley, "Efficient analytic approximation of american option values," *The Journal of Finance*, vol. 42, no. 2, pp. 301-320, Jun 1987.
- [182] N. S. Rasmussen, *Efficient control variates for Monte-Carlo valuation of American options*, 2002.
- [183] N. S. Rasmussen, "Control variates for Monte Carlo valuation of American options," *Journal of Computational Finance*, vol. 9, no. 1, 09 2005.
- [184] N. S. Rasmussen, *Improving the least-squares Monte-Carlo approach*, 2002.
- [185] J. M. Harrison, *Brownian motion and stochastic flow systems*, Krieger Publishing Company, 1985.
- [186] P. Jackel, *Monte Carlo methods in finance*, Wiley Finance, 2002.
- [187] J. Halton, "Algorithm 247: Radical-inverse quasi-random point sequence," *Communications of the ACM*, vol. 7, pp. 701-702, Dec 1964.
- [188] H. Niederreiter, "Low-discrepancy and low-dispersion sequences," *Journal of Number Theory*, vol. 30, pp. 51-70, 1988.
- [189] I. Sobol, "On the distribution of points in a cube and the approximate evaluation of

integrals," *USSR Computational Mathematics and Mathematical Physics*, vol. 7, no. 4, pp. 86-112, May 1966.

- [190] I. M. Sobol, D. Asotsky, A. Kreinin and S. Kucherenko, "Construction and comparison of high-dimensional Sobol' generators," *Wilmott magazine*, pp. 64-79, 2011.
- [191] Microsoft, *INFO: How Visual Basic generates pseudo-random numbers for the RND function*, 2004.
- [192] M. Matsumoto and T. Mishimura, "Mersenne Twister: A 623-dimensionally equidistributed uniform pseudo-random number generator," *ACM Transactions on Modeling and Computer Simulation*, vol. 8, no. 1, pp. 3-30, 1 1998.
- [193] G. Marsaglia and W. W. Tsang, "Some difficult-to-pass tests of randomness," *Journal of Statistical Software*, vol. 7, no. 3, 2002.
- [194] L. Stentoft, "Assessing the least squares Monte-Carlo approach to American option valuation," *Review of Derivatives Research*, vol. 7, pp. 129-168, 2004.
- [195] M. Moreno and J. F. Navas, "On the robustness of least-squares Monte Carlo for pricing American derivatives," *Review of Derivatives Research*, vol. 6, no. 2, 2003.
- [196] S. D. Jacka, "Optimal stopping and the American put," *Mathematical Finance*, vol. 1, no. 2, pp. 1-14, 1991.
- [197] H. Pham, "Optimal stopping, free boundary, and American option in a jump-diffusion model," *Applied Mathematics and Optimization*, vol. 35, pp. 145-164, 1997.
- [198] Y.-K. Kwok, *Mathematical models of financial derivatives*, 2 ed., Springer, 2008.

- [199] P. Carr, R. Jarrow and R. Myneni, "Alternative characterizations of American put options," *Mathematical Finance*, vol. 2, no. 2, pp. 87-106, 1992.
- [200] A. J. Pope, "The statistics of residuals and the detection of outliers," Rockville, Md, 1976.
- [201] D. Bolle and S. Fahrenkrog, A guide to the project management body of knowledge, 3 ed., Project Management Institute, 2004.
- [202] W. Schoutens, Levy processes in finance: pricing financial derivatives, J. Wiley and Sons, Eds., Chichester, UK: John Wiley and Sons, 2003.
- [203] R. Geske, "The valuation of compound options," *Journal of Financial Economics*, vol. 7, pp. 63-81, 1979.
- [204] G. Norris, *More MD-11 improvements revealed*, 1995.
- [205] CFM, *CFM56-5B PIP for the Airbus A320 family*, Retrieved November 23rd 2013.
- [206] K. Thomson and E. T. Schulze, *Delivering fuel and emissions savings for the 777*, 2009.
- [207] Airbus, *Airbus launches Sharklet retrofit for in-service A320 family aircraft*, 2013.
- [208] Aircraft Commerce, *CFM56-5A/5B maintenance analysis and budget*, vol. 50, C. Williams, Ed., 2007.
- [209] Aircraft Commerce, *CFM56-5A/5B modification programmes*, vol. 50, C. Williams, Ed., 2007.
- [210] Aircraft Commerce, *V2500 family modification and upgrade programmes*, 2008.
- [211] C. Meher-Homjo, M. Chaker and H. Motiwala, "Gas turbine performance deterioration," *Proceedings of the 30th turbomachinery symposium*, pp. 139-175,

09 2001.

- [212] S. Lattime and B. Steinetz, "Turbine engine clearance control systems: current practices and future directions," in *38th Joint Propulsion Conference and Exhibit*, Indianapolis, IN, 2002.
- [213] C. Y. Justin and D. N. Mavris, "Aircraft and engine economic evaluation for fleet renewal decision-making and maintenance contract valuation," *Proceedings of the Institution of Mechanical Engineers. Part G, Journal of Aerospace Engineering*, vol. 229, no. 11, pp. 2051-2065, 09 2015.
- [214] W. Gibson and P. Morrell, "Theory and practice in aircraft financial evaluation," *Journal of Air Transport Management*, vol. 10, no. 6, pp. 427-433, 11 2004.
- [215] P. Thokala, J. Scanlan and A. Chipperfield, "Life cycle cost modelling as an aircraft design support tool," *Proc IMechE Part G: Journal of Aerospace Engineering*, vol. 224, pp. 477-488, 2010.
- [216] J. K. Brueckner, D. N. Lee, P. M. Picard and E. Singer, *Product unbundling in the travel industry: the economics of airline bag fees*, 2013.
- [217] R. Klophaus, "The economic of unbundling air travel services," in *Proceedings of 15th HKSTS International Conference*, Hong Kong, 2010.
- [218] T. Economist, *Aircraft leasing, buy or rent?*, 2012.
- [219] L. Wayne, *The real owner of all those planes*, NY, 2007.
- [220] W. E. Gibson, "Airline fleet planning and aircraft investment valuation," in *Air Transport Research Society*, Berkeley, CA, 2007.
- [221] Aircraft Commerce, *A320 family maintenance analysis and budget*, vol. 44, C.

Williams, Ed., 2006.

- [222] S. Chiesa, S. Quer, S. Corpino and N. Viola, "Heuristic and exact techniques for aircraft maintenance scheduling," *Proc IMechE Part G: Journal of Aerospace Engineering*, vol. 223, pp. 989-999, July 2009.
- [223] Aircraft Commerce, *The juggling act of CFM56-5A/B maintenance*, vol. 15, C. Williams, Ed., 2001.
- [224] Aircraft Commerce, *LLP management for short-haul engines*, vol. 34, C. Williams, Ed., 2004.
- [225] Aircraft Commerce, *Owner's and operator's guide: CFM56-5A/-5B*, vol. 50, C. Williams, Ed., 2007.
- [226] S. Ackert, "Engine maintenance concepts for financiers: Elements of turbofan shop maintenance costs," San Francisco, CA, 2011.
- [227] Aircraft Commerce, *CFM56-5A/-5B series specifications*, vol. 50, C. Williams, Ed., 2007.
- [228] O. Rupp, "Economic aspects of maintaining engine efficiency," in *The Engine Yearbook 2005*, P. Copping, Ed., London, UK, Aviation Industry Press, 2004, pp. 64-67.
- [229] Aircraft Commerce, *CFM56-3 maintenance analysis and budget*, vol. 45, C. Williams, Ed., 2006.
- [230] G. Thomas, "Engine maintenance from an OEM perspective," in *The Engine Yearbook 2011*, D. Horwitz, Ed., London, UK, UBM Aviation, 2010, p. 50.
- [231] H. Hanumanthan, "Severity estimation and shop visit prediction of civil aircraft

- engines," Cranfield University, UK, 2009.
- [232] O. Rupp, "Engine maintenance costs," in *The Engine Yearbook 2005*, P. Copping, Ed., London, UK, Aviation Industry Press, 2004, p. 20.
- [233] A. Domitrovic, E. Bazijanac and I. Cala, "Optimal replacement policy of jet engine modules from the aircarrier's point of view," *PROMET Traffic and Transportation*, vol. 20, pp. 1-9, 2008.
- [234] E. V. Zaretsky, R. C. Hendricks and S. Soditus, "Weibull-based design methodology for rotating aircraft engine structures," ISROMAC-9, Ninth International Symposium on Transport Phenomena and Dynamics of Rotating Machinery, Honolulu, HI, 2002.
- [235] S. Ackert, "Basics of aircraft maintenance reserve development and management," San Francisco, CA, 2012.
- [236] Aircraft Commerce, *CFM56-5B/-7 and V.2500 maintenance costs re-examined*, vol. 28, C. Williams, Ed., 2003.
- [237] IATA, "2009 Commercial air transport MRO logistics survey - Airline participant feedback report," IATA, Montreal, QB, 2010.
- [238] M. Wiseman and T.-H. Guo, "An investigation of life extending control techniques for gas turbine engines," in *Proceedings of the American Control Conference*, Arlington, VA, 2001.
- [239] I. Yilmaz, "Evaluation of the relationship between exhaust gas temperature and operational parameters in CFM56-7B engines," *Proc IMechE Part G: Journal of Aerospace Engineering*, vol. 223, pp. 433-440, April 2009.

- [240] S. Deblock, *EU ETS and aviation*, 2012.
- [241] European Commission, *Reducing the climate change impact of aviation*, 2005.
- [242] B. Pearce and D. Pearce, *Setting environmental taxes for aircraft: a case study of the UK*, 2000.
- [243] Aircraft Commerce, *Complying with the EU's ETS*, vol. 70, C. Williams, Ed., 2010.
- [244] J. D. Scheelhaase and W. G. Grimme, "Emissions trading for international aviation - an estimation of the economic impact on selected European airlines," *Journal of Air Transport Management*, vol. 13, pp. 253-263, 2007.
- [245] EUROCONTROL, "Standard inputs for Eurocontrol cost benefit analyses," Brussel, 2013.
- [246] FAA, *Code of Federal Regulations - FAR - Part 121 - Sec 467 - Title 14*, 2010.
- [247] FAA, *Code of Federal Regulations - FAR - Part 121 - Sec 471 - Title 14*, 2010.
- [248] Aircraft Commerce, *Happy families: the A320 vs 737/757*, vol. 15, C. Williams, Ed., 2001.
- [249] EUROCONTROL, "Adjusted unit rates applicable to January 2014 flights," Brussel, 2014.
- [250] IATA, "Airport and en-route aviation charges manual," Montreal, QB, 2005.
- [251] D. B. Kelly, "Forecasting aircraft values: an appraiser's perspective," in *Airfinance Annual*, Chantilly, VA, Avitas, Inc, 2008, pp. 24-30.
- [252] J. G. Wensveen, *Air transportation - A management perspective*, 6th ed., Aldershot, England: Ashgate Publishing, 2007, p. 305.

- [253] Aircraft Commerce, *CFM56-3 and -5A/B maintenance costs*, vol. 8, C. Williams, Ed., 1999.
- [254] F. S. Nowlan and H. F. Heap, *Reliability-centered maintenance*, Dolby Access Press, 1978.
- [255] F. Nowlan, "An evaluation of the relationship between reliability, overhaul period, and economics in the case of aircraft engines," 1960.
- [256] G. Yu, C. Xin-Feng and L. Nan, "Failure analysis on HPT blade of CFM56-7B".
- [257] D. Pascovici, K. G. Kyprianidis, F. Colmenares, S. Ogaji and P. Pilidis, "Weibull distributions applied to cost and risk analysis for aero engines," 2008.
- [258] Aircraft Commerce, *A320 family fuel burn performance*, vol. 44, C. Williams, Ed., 2006.
- [259] Aircraft Commerce, *A320 family values and aftermarket activity*, vol. 44, C. Williams, Ed., 2006.
- [260] Aircraft Commerce, *The different elements of pilot employment costs*, vol. 62, C. Williams, Ed., 2009.
- [261] Aircraft Commerce, *Airbus and Boeing narrowbodies: maintenance cost analysis*, vol. 25, C. Williams, Ed., 2002.
- [262] F. D. Harris, *An economic model of U.S. airline operating expenses*, 2005.
- [263] J. Flottau, *Ryanair Introducing emissions trading surcharge*, 2012.
- [264] Aircraft Commerce, *CFM56-5A/-5B fuel burn performance*, vol. 50, C. Williams, Ed., 2007.
- [265] Aircraft Commerce, *Analysing the options for 757 replacement*, vol. 42, C.

Williams, Ed., 2005.

[266] P. Lironi, *The CFM56-5C*, 2005.

[267] N. O'Keeffe, *GE voices concerns on CFM56 spare engine availability*, 2008.

[268] W. L. F. Corporation, *Form 10-K*, 2009, p. 3.

[269] B. o. T. Statistics, *Form 41*, 2015.

[270] M. Mirza, *Economic impact of airplane turn-times*, Boeing, Ed., 2008.

[271] R. D. Luce, "The choice axiom after twenty years," *Journal of Mathematical Psychology*, vol. 15, pp. 215-233, 1977.

[272] J. Lesourne, *A theory of the individual for economic Analysis*, vol. 1, North Holland Publishing Company, 1977.

[273] D. McFadden, "Conditional logit analysis of qualitative choice behavior," *Frontiers in Econometrics*, pp. 105-142, 1974.

[274] D. McFadden, "Econometric models for probabilistic choice among products," *The Journal of Business*, vol. 53, no. 3, p. 13, 07 1980.

[275] D. Raghavarao, J. B. Wiley and P. Chitturi, *Choice-based conjoint andesign: Models and designs*, Chapman and Hall, 2011.

[276] S. Hamilton, *A320 NEO to have \$7-8 million price premium*, 2010.

[277] W. S. Cleveland, "LOWESS: A program for smoothing scatterplots by robust locally weighted regression," *The American Statistician*, vol. 35, no. 1, p. 54, 2 1981.

[278] J. Renshaw, *Delta's Philadelphia refinery running at 110 percent*, 2015.

- [279] D. K. Yousef, *Emirates will buy 100 more A380s if Airbus upgrades model*, 2015.
- [280] J. Y. Campbell, A. W. Lo and A. C. MacKinlay, *The econometrics of financial markets*, Princeton University Press, 1997.
- [281] C. Group, *U.S. Gulf Coast kerosene-type jet fuel future quotes*, 2015.
- [282] M. Szabo, *Airlines alliances take cartel approach to carbon trading*, 2012.
- [283] BlueNext, *BNS EUA 08-12*, 2012.
- [284] EUREX, *EUA Futures*, 2015.
- [285] A. Einstein, "Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen," *Annalen der Physik*, vol. 322, pp. 549-560, 1905.
- [286] A. Dixit and R. Pindyck, *Investment under uncertainty*, Princeton University Press, 1994.
- [287] R. C. Merton, "On estimating the expected return on the market: an exploratory investigation," *Journal of Financial Economics*, vol. 8, pp. 323-361, 1980.
- [288] A. Dixit, *The art of smooth pasting*, Princeton University Press, 1993.
- [289] R. Cont and P. Tankov, *Financial modelling with jump processes*, Chapman and Hall, 2003.
- [290] P. Tankov and E. Voltchkova, "Jump-diffusion models: a practitioner's guide," *Banque et Marches*, vol. 99, pp. 1-24, 2009.
- [291] S. G. Kou, "A jump-diffusion model for option pricing," *Management Science*, vol. 48, no. 8, pp. 1086-1101, 08 2002.

- [292] S. S. Lee and P. A. Mykland, "Jumps in financial markets: a new nonparametric test and jump dynamics," *Review of Financial Studies*, vol. 21, no. 6, pp. 2535-2563, 9 2008.
- [293] O. E. Barndorff-Nielsen and N. Shephard, "Econometrics of testing for jumps in financial economics using bipower variation," *Journal of Financial Econometrics*, vol. 4, pp. 1-30, 2006.
- [294] T. G. Andersen, T. Bollerslev and F. X. Diebold, "Roughing it up: Including jump components in the measurement, modeling, and forecasting of return volatility," *Review of Economics and Statistics*, vol. 89, no. 4, pp. 701-720, 10 2007.
- [295] S. S. Lee and J. Hannig, "Detecting jumps from Levy jump diffusion processes," *Journal of Financial Economics*, vol. 96, no. 2, pp. 439-479, 2010.
- [296] R. Cont and P. Tankov, "Calibration of jump-diffusion option-pricing models: a robust non-parametric approach," *Journal of Computational Finance*, vol. 7, no. 3, pp. 1-49, 09 2004.
- [297] G. J. Jiang and R. C. Oomen, *A new test for jumps in asset prices*, 2005.
- [298] R. Cont and P. Tankov, "Calibration of jump-diffusion option-pricing models: a robust non-parametric approach," *Journal of Computational Finance*, vol. 7, no. 3, pp. 1-49, 09 2004.
- [299] P. Bjerksund and S. Ekern, "Real options in capital investment: models, strategies, and applications," 3rd ed., Praeger, 1995, pp. 207-220.
- [300] L. Breiman, Probability, R. O'Malley, Ed., Society for Industrial and Applied Mathematics, 1992, pp. 303-305.

- [301] D. Duffy, Introduction to C++ for financial engineers: an object-oriented approach, J. Wiley, Ed., Wiley Finance, 2006.
- [302] P. Barker, Java methods for financial engineers, Springer-Verlag, 2007.
- [303] G. E. P. Box and M. E. Muller, "A note on the generation of random normal deviates," *The Annals of Mathematical Statistics*, vol. 29, pp. 610-611, 1956.
- [304] A. M. Turing, "Rounding-off errors in matrix processes," *Quarterly Journal of Mechanics and Applied Mathematics*, pp. 287-308, 1948.
- [305] P. Wilmott, S. Howison and J. Dewynne, The mathematics of financial derivatives - A student introduction, Press Syndicate of the University of Cambridge, 1995.
- [306] D. E. Knuth, The art of computer programming, 2 ed., vol. 2, M. Harrison, Ed., Addison-Wesley Professional, 1981.

VITA

CEDRIC Y. JUSTIN

Cedric Justin was born in Rouen, France. He attended public schools in Bonsecours, Franqueville-St-Pierre and Rouen before receiving a Master of Science in Engineering degree from the Ecole Nationale Supérieure des Mines de Nancy, in Lorraine, France. In 2005 he came to the Georgia Institute of Technology to pursue his long-time aspiration of becoming an aerospace engineer, receiving a Master of Sciences in Aerospace Engineering from the Guggenheim School of Aerospace Engineering in 2006. While working on his PhD dissertation, he graduated in 2012 with a Master of Sciences in Quantitative and Computational Finance offered by the Milton Steward School of Industrial Engineering at Georgia Tech. Cedric Justin is an instrument-rated commercial pilot and enjoys buzzing around the skies in a seaplane or sailing when not working on his research.