Order Number 9123983

A mean-variance serial replacement decision model

Brown, Matthew J., Ph.D.

The University of Michigan, 1991

U·M·I
300 N. Zeeb Rd.
Ann Arbor, MI 48106

A MEAN-VARIANCE SERIAL REPLACEMENT DECISION MODEL

by
Matthew J. Brown

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Industrial and Operations Engineering)
in The University of Michigan
1991

Doctoral Committee:

Associate Professor Jack R. Lohmann, Chair
Associate Professor Anastassios N. Perakis
Assistant Professor Medini R. Singh
Associate Professor Candace A. Yano

I learned this, at least, by my experiment; that if one advances confidently in the direction of his dreams, and endeavors to live the life which he has imagined, he will meet with a success unexpected in common hours. He will put some things behind, will pass an invisible boundary; new, universal, and more liberal laws will begin to establish themselves around and within him; or the old laws be expanded, and interpreted in his favor in a more liberal sense, and he will live with the license of a higher order of beings. In proportion as he simplifies his life, the laws of the universe will appear less complex, and solitude will not be solitude, nor poverty poverty, nor weakness weakness. If you have built castles in the air, your work need not be lost; that is where they should be. Now put the foundations under them.

Henry David Thoreau
Walden
1846

RULES REGARDING THE USE OF MICROFILMED DISSERTATIONS

To my parents.

ACKNOWLEDGEMENTS

Many people helped me to achieve more academically than I ever dreamed was possible. The completion of this dissertation represents the culmination of ten collegiate years of full-time study. Completion would not have been possible without the help I received from my doctoral committee. I wish to thank Professor Jack R. Lohmann, not only for serving as the committee chair, but also for not abandoning me and making all the trips back and forth between Washington, D.C. and Ann Arbor. Additionally, I wish to thank Professors Anatassios N. Perakis, Medini R. Singh, and Candace A. Yano for their suggestions, insights, and comments. They helped keep me motivated, on-track, and honest.

The completion of this dissertation also would not have been possible without financial support in the form of fellowships, teaching and research assistantships, fringe benefits, and tuition waivers. Financial support was provided by the IOE Department, the Institute of Industrial Engineers, Professor Jack R. Lohmann, and Professor James B. Martin.

In considering a career in academia, I spent many fascinating hours studying the lifestyle of a professor. In addition to the members of my doctoral committee, my interest in an academic career was influenced by John R. Birge, Vivian L. Carpenter, Paul W. Durance, Walton M. Hancock, Gary D. Herrin, W. David Kelton, W. Monroe Keyserling, Jeffery K. Liker, James B. Martin, Chan S. Park, Stephen M. Pollock, and Robert L. Smith.

Last but not least are two special acknowledgements. First, I want to thank my parents for all they have done for me and for all the calls on Sunday night just to say hello. Second, a heartfelt thanks to Professor Candace A. Yano for all her support, encouragement, guidance, and intellectual participation over the last eight years. She was the first to encourage me to consider a Ph.D. program, a part of almost every major decision I made along the way, the recipient of many of my long e-mails, the person who was always there to provide emotional support and motivation when I needed it, and my role model. Since I could not have finished without her, I hope that I will be able to do something for her in the future beyond simply being "amusing."

TABLE OF CONTENTS

DEDICATION ii

ACKNOWLEDGEMENTS iii

LIST OF FIGURES vii

LIST OF TABLES viii

LIST OF APPENDICES xi

CHAPTER

I. INTRODUCTION 1

1.1 Overview 4

II. LITERATURE REVIEW 6

2.1 Background and Definitions 7
2.2 Project Selection Methods Under Uncertainty 10

III. MEAN-VARIANCE DECISION MODEL 14

3.1 Decision Environment and Major Assumptions 15
3.2 Five Alternative Decision Procedures 19
3.2.1 Random 19
3.2.2 Traditional 20
3.2.3 Expected Value 21
3.2.4 Certain Monetary Equivalent 21
3.2.5 Expected Utility 22
3.3 Performance Measures 28

IV. COMPUTATIONAL EXPERIMENTS 32

4.1 The Test Problem Generator 33
4.1.1 Test Problem Generator Parameters 34

4.1.2 Setting the Level of Risk Aversion 37
4.1.3 Setting the Correlation Coefficients 39
4.1.4 Measure of Uncertainty 41
4.1.5 Generating the NPVs 43
4.2 The Experiments 46
4.2.1 Experiment One 47
4.2.2 The Cluster Heuristic and Upper Bounding Procedures 53
4.2.3 Experiment Two 57
4.2.4 Experiment Three 62
4.2.5 Experiment Four 71
4.2.6 Experiment Five 79
4.3 Summary of Parameter Effects 83

V. SUMMARY AND EXTENSIONS 88

5.1 Summary 91
5.2 Extensions 94

APPENDICES **96**

BIBLIOGRAPHY **114**

LIST OF FIGURES

Figure

3.1 Branch and Prune Procedure 26

4.1 MV-Efficient Nodes for a Sample EX1 Problem, m = 15% 52

4.2 MV-Efficient Nodes for a Sample EX1 Problem, m = 30% 52

A.1 Decision Tree 98

LIST OF TABLES

Table

3.1	Random Sequence Performance Measure Variance	29
3.2	Sample Utility Performance Values	31
4.1	Correlation Coefficient Matrix for Three Asset Types	40
4.2	Correlation Coefficient Sign for Negative Case	41
4.3	EX1 Parameter Ranges	48
4.4	Maximum Number of MV-Efficient Nodes Required	49
4.5	EX1 Analysis of Variance for m=15%	50
4.6	EX2 Parameter Distributions	58
4.7	EX2 Problem Set 1 - Uncorrelated Assets	59
4.8	EX2 Problem Set 2 - Uncorrelated Assets, Higher Uncertainty	59
4.9	Analysis of Variance for EX3 TRAD	64
4.10	Analysis of Variance for EX3 EV	65
4.11	Analysis of Variance for EX3 CME	66
4.12	Average Upper Bound Utility Performance	67
4.13	Analysis of Variance for Positive Correlation	68
4.14	Analysis of Variance for Negative Correlation	69
4.15	EX4 Overall Results	76
4.16	Summary of EX4 Significant Factors	77

4.17 EX5 Overall Results 81

4.18 Summary of EX5 Significant Factors 82

B.1 Analysis of Variance for EX4 TRAD, Exponential Utility 101

B.2 Analysis of Variance for EX4 TRAD, Logarithmic Utility 101

B.3 Analysis of Variance for EX4 TRAD, Power Utility 102

B.4 Analysis of Variance for EX4 EV, Exponential Utility 102

B.5 Analysis of Variance for EX4 EV, Logarithmic Utility 103

B.6 Analysis of Variance for EX4 EV, Power Utility 103

B.7 Analysis of Variance for EX4 CME, Exponential Utility 104

B.8 Analysis of Variance for EX4 CME, Logarithmic Utility 104

B.9 Analysis of Variance for EX4 CME, Power Utility 105

B.10 Analysis of Variance for EX5 Mean Only, (0.00,0.25) 105

B.11 Analysis of Variance for EX5 Mean Only, (0.25,0.50) 106

B.12 Analysis of Variance for EX5 Mean Only, (0.50,0.75) 106

B.13 Analysis of Variance for EX5 Mean Only, (0.75,1.00) 107

B.14 Analysis of Variance for EX5 Variance Only, (0.00,0.25) 107

B.15 Analysis of Variance for EX5 Variance Only, (0.25,0.50) 108

B.16 Analysis of Variance for EX5 Variance Only, (0.50,0.75) 108

B.17 Analysis of Variance for EX5 Variance Only, (0.75,1.00) 109

B.18 Analysis of Variance for EX5 Mean and Variance, (0.00,0.25) 109

B.19 Analysis of Variance for EX5 Mean and Variance, (0.25,0.50) 110

B.20 Analysis of Variance for EX5 Mean and Variance, (0.50,0.75) 110

B.21 Analysis of Variance for EX5 Mean and Variance, (0.75,1.00) 111

B.22 Analysis of Variance for EX5 Mean,Same Direction, (-0.50,-0.25) . . 111

B.23 Analysis of Variance for EX5 Mean,Same Direction, (-0.25,0.00) . . 112

B.24 Analysis of Variance for EX5 Mean,Same Direction, (0.00,0.25) . . . 112

B.25 Analysis of Variance for EX5 Mean,Same Direction, (0.25,0.50) . . . 113

LIST OF APPENDICES

Appendix

A. EXAMPLE OF HOW SELECTING THE ASSET WITH THE HIGHEST PROBABILITY OF BEING THE OPTIMAL FIRST ASSET CAN BE SUBOPTIMAL 97

B. ANALYSIS OF VARIANCE TABLES FOR EX4 AND EX5 100

CHAPTER I

INTRODUCTION

Questions about what equipment to replace, what new technology to adopt, and when to do it arise in a variety of situations. Studies have indicated that replacement projects consume a significant portion of a firm's capital budget (one study estimates as much as 50 percent [15]). As Park and Sharp [46] note, "how the capital is to be allocated is one of the most important and most difficult decisions that any firm must make." Thus, one key to the success of a firm is to make intelligent replacement decisions.

A major subset of replacement problems, and the focus of this dissertation, involve replacement of a single asset with another single asset. These problems are commonly called serial replacement problems. Typically, for such problems, the decision to be made is whether to either keep an asset that is currently providing a service or replace it with one of several alternative new assets. The service is usually expected to be required much longer than any asset currently available can remain in service. Thus, similar decisions will have to be made in the future and incorporated into the current analysis because they impact the current decision. The objective is to determine the "most economical" sequence of assets to provide the required service to the horizon. A sequence of assets is completely described by specifying for each asset in the sequence: the asset type, its time of installation, and its length of service.

Since assets available in the future can differ technologically and monetarily from assets currently available, a mechanism is needed to characterize these future assets. It is often assumed that assets evolve over time and that future versions are related to the assets currently available by a function of the time they are installed. This requires specification of only the monetary consequences that would result from installing, operating, and replacing each of the assets available currently and the functions to describe the change in the monetary consequences for future assets. Uncertainty about the monetary consequences can be modeled with stochastic cash flows.

Much of prior replacement research has assumed typically that either the cash flows that describe the monetary consequences resulting from an asset providing service are deterministic or that the decision maker's (DM's) objective is to maximize expected value (EV). This dissertation develops a model that incorporates a DM's risk aversion directly into the replacement analysis, in the form of a utility function. It also studies the model under a wide range of scenarios and provides some insight regarding the conditions under which use of an EV model may be undesirable.

Utility theory, as Currim and Sarin [10] note "is the dominant paradigm for decision making under risk in a variety of disciplines, such as economics, finance, and marketing." Recently, Oakford, Bhimjee, and Salazar [41], Thompson and Thuesen [55], and Ramis, Thuesen, and Barr [50] have used utility within a capital budgeting framework. However, utility has yet to be considered in replacement models. A model incorporating utility is desirable because:

- DMs often view the impact of monetary gains and losses differently [54]. As losses become greater, the impact of each additional dollar lost becomes greater.

Similarly as gains increase, the impact of each additional dollar gained becomes less. Thus, the use of an EV criterion may be inconsistent with the DM's views.

- Risk is considered directly, thus adding a dimension that traditional sensitivity analysis cannot address and EV only partly addresses.

- A utility-based model can solve a wider range of problems since EV is a special case of expected utility (EU).

Replacement problems with an objective of maximizing EU can be difficult to solve for at least three reasons. First, the monetary consequences which result from installing, operating, and replacing assets in the future must be forecasted. With rapid changes in technology and the need to adopt products and services to a quickly changing market, forecasting the required data can be difficult. Second, the number of different asset sequences that could provide the required service is typically so large that total enumeration is not a viable solution procedure. Dynamic programming provides a computationally efficient approach to find an optimal solution for both deterministic models and stochastic models that use EV as the monetary objective. However, for stochastic models with an EU objective, pruning partial sequences is more difficult since utility is not additive. Finally, the method used to evaluate the attractiveness of sequences of assets must account for the temporal cash flows. Discounted cash flow analysis has been the most popular method used to determine an equivalent measure of worth for both deterministic models and stochastic models with an EV monetary objective. However, for stochastic models with an EU objective, evaluating a given sequence of assets is not as easy. The difficulty arises from the presence of multiple random variables in each time period. These random variables can be correlated across time periods and assessing a utility function for multiple periods is typically more difficult than for a single period.

The research objectives of this dissertation were threefold. The first objective was to develop a procedure to find the sequence of highest EU. This was accomplished by expanding the longest path deterministic dynamic program used to find the EV sequence, into a branch and prune algorithm using mean-variance (MV) dominance. The second objective was to gain insight about replacement decisions under uncertainty by studying a wide range of problems and examining the impact of a number of different model parameters. The third objective was to evaluate and compare the performance of several alternative decision procedures by performing a series of computational experiments.

1.1 Overview

Chapter II reviews the replacement literature. Due to the broad range of replacement problems, it is helpful to describe the models in terms of a classification scheme. Chapter II classifies replacement models based on five characteristics: the number of assets providing service, the horizon, how the relevant cash flows are modeled, asset reliability assumptions, and the DM's monetary objective. Methods that have been used in a variety of decision environments to select projects under uncertainty are then discussed.

The MV model, which incorporates utility directly, is discussed in Chapter III. The major assumptions made by the model are first described. Five alternative decision procedures that vary in level of sophistication are then described. The decision procedures are random, traditional, EV, certain monetary equivalent, and EU. The chapter concludes by developing the performance measures used for the computational experiments to evaluate and compare the different decision procedures.

Chapter IV describes the computational experiments. A total of five experiments were performed. The experiments were designed to examine the ability of the EU decision procedure to solve a wide range of problems, evaluate and compare the performance of the alternative decision procedures under a variety of conditions, study the impact of both constant and decreasing risk aversion, and examine the sensitivity of the EU sequence to changes in the forecasts.

Finally, Chapter V provides a research summary and some possible extensions to the MV model.

CHAPTER II

LITERATURE REVIEW

Many types of assets that provide a service or produce a product are replaced over time. Replacement occurs when an asset fails and cannot be repaired, when the cost of keeping an asset operational is prohibitive, when changes in technology make an asset inferior or obsolete, or simply when a change is desired. A replacement decision typically involves a choice between either keeping a set of one or more assets currently in service or replacing them with a new set of assets; or if there is no set of assets currently in service, a choice between installing one of several alternative sets of new assets. In this dissertation, the term "service" may refer to both a service, such as transportation provided by a vehicle, or the production of products. The DM's objective typically is to determine the sequence of assets that will provide the most economical, either lowest cost or most profitable, service for the period of time under consideration. The replacement decision domain is not limited to studies of machines found in typical manufacturing operations, but includes virtually any asset that provides service and is replaced.

Replacement research has been conducted actively since the late 1940s. Thus, it is not surprising that a wide range of models have been developed. The models differ in terms of assumptions, realism, focus, scope, modeling flexibility, etc. The next section describes some major characteristics that can be used to classify replacement models. Different methods to select projects under uncertainty that are used in a variety of decision environments are then discussed. These methods could be incorporated into replacement models to select among alternative sequences of assets under uncertainty.

2.1 Background and Definitions

While detailed replacement classification schemes exist [36], five characteristics appear sufficient to classify most models. They are: the number of assets replaced at a given time, the horizon, how the relevant cash flows are modeled, the assumptions made about the reliability of the assets, and the DM's monetary objective. Each characteristic is described below including a brief summary of the relevant replacement model literature. A more detailed summary of the literature can be found in Bean, Lohmann, and Smith [5] or Ganesh [15].

The **number of assets replaced at a given time** determine whether a replacement model is a serial or a parallel model. "Serial" replacement models assume a single asset replaces another single asset, while "parallel" or "multiple" replacement models assume a set of assets replace another set of assets. A set of assets can be modeled serially if the entire set is always replaced at the same time. The majority of the replacement models address serial rather than parallel problems. Two notable exceptions are the work of Leung and Tanchoco [30], and VanderVeen [56]. Leung and Tanchoco looked at single period model to study replacement of a subset of machines within an integrated system, such as a flexible manufacturing system. VanderVeen formulated the general problem as a mixed integer program, developed a set of dominance criteria and a decision variable encoding scheme to reduce the number of variables and constraints, and evaluated the performance of five heuristics.

The length of time service is required is referred to as the **horizon** and it can be either "finite" or "infinite." The most well known serial model assumes an infinite horizon and identically repeating future assets [16]. These two assumptions, while making the problem easy to solve, frequently do not model reality well. A dynamic programming approach which relaxed the repeatability assumption but required a finite horizon was developed by Wagner [58] in 1975. A more general dynamic programming approach was introduced in 1984 by Oakford, Lohmann, and Salazar [42]. While still dealing only with finite horizons, the model allowed more than one asset type to be available at each point in time. This model was extended to the infinite horizon case in 1985 by Bean, Lohmann, and Smith [5]. In 1986 Lohmann [35] extended the Bean, Lohmann, and Smith model to allow for stochastic cash flows by combining simulation and dynamic programming.

The horizon is important because the first asset type in the optimal sequence may depend on the length of the horizon. While no service will be required forever, infinite horizons are often used to try to avoid end-of-horizon effects. For a DM unsure about the appropriate horizon, the equivalent finite horizon may be able to be determined [5]. This value provides the length of the horizon beyond which the optimal first asset would not change. Sometimes DMs will avoid using a replacement model because they feel they cannot forecast the relevant cash flows to the horizon. Instead, the problem is inappropriately modeled as a one time decision. Methods have been developed to provide a bound on the resulting error. For example, Bean, Lohmann, and Smith [6] have developed an error bounding procedure for a deterministic serial replacement model.

"Deterministic" models assume the **relevant cash flows** are known with certainty [42]. "Stochastic" models treat some or all the cash flows as uncertain, requiring a DM to provide probability distributions for the uncertain cash flows [35]. Deterministic models are by far the most common, primarily because they tend to be easier to solve than stochastic models.

Replacement models fall into two categories based on the **assumptions about an asset's reliability** during its service life. Some models assume an asset to have a maximum life over which it can provide service. During this life the asset is assumed to never fail to the point it has to be replaced before the maximum life (e.g., Bean, Lohmann, and Smith [5]). Alternatively, other models have been developed which recognize that assets sometimes must be replaced due to failure before reaching some maximum service life. These models combine simple reliability ideas and replacement and are sometimes called "equipment inspection and replacement models." Often, to model the reliability of an asset, a set of states is defined to characterize the possible conditions of that asset. For example, the states could represent a machine's condition as new, excellent, good, fair, poor, and failed. The state of the machine in the next period given its current state is defined in terms of a transition probability matrix. Each possible decision (i.e., set of keep or replace decisions for each state) has a separate matrix associated with it. Since the state of the system will not be known until the DM reaches that point in time, the optimal sequence cannot be identified a priori. Instead, a DM needs to know what the best course of action is given the current state. Thus, the objective is to determine the optimal "policy," where a policy specifies the best decision for each possible state. Such problems can be formulated as a Markov Decision Problem (MDP) and are discussed by Denardo [12], Dreyfus and Law [13], and Ross [51].

The equipment inspection and replacement models generally do not offer the modeling flexibility of those which assume assets do not unexpectedly fail before their maximum service life. Flexibility must be compromised due to the structure required to make the solution procedure tractable. For example, by simply allowing the reward received from an asset's service to depend on the period the asset is installed, so that rewards are not stage invariant, nonstationary policies and added solution complexity may result. Further complexity is added when the DM desires to use utility as the basis for making decisions, as Jaquette [25] describes for a general MDP with an objective of maximizing expected utility for a DM with constant risk aversion. For these reasons, the MV model described in the next chapter does not consider the case of an asset that must be replaced unexpectedly before it reaches its maximum service life. However, the impact of failure without replacement can be considered by including this cost in the cash flows.

The **DM's monetary objective** typically has been to maximize either present value or future value. No utility-based replacement models were found in the literature. However, methods that consider uncertainty have been used in many other decision environments, and some of the related literature is discussed in the next section.

2.2 Project Selection Methods Under Uncertainty

Decision theory provides both descriptive and prescriptive methods for dealing with uncertainty. To study how DMs define and react to uncertainty in practice, a number of surveys have been taken. March and Shapira [38] provide an overview of many of these surveys and point out that while the terms "risk" and "uncertainty" are often used interchangeably in the literature, DMs tend to view risk in terms of the magnitude of possible loss. Just because a project has a wide range of possible outcomes (i.e., high uncertainty), does not mean it would be considered risky as might be the case if all the outcomes were positive. Similarly, if a project had a narrow range of possible outcomes, but all resulted in large losses, it may be viewed as quite risky. In this dissertation, uncertainty is used to refer to the variance of monetary outcomes.

A number of surveys have verified that DMs are risk averse, at least under certain conditions. For example, Swalm [53] presents plots of actual utility functions which exhibit the concave shape expected for risk averse DMs. MacCrimmon and Wehrung [37] found that DMs who were the most risk averse tend to be the most mature. Laughhunn, Payne, and Crum [29] found that when faced with only non-ruinous losses, most DMs were risk seeking in trying to overcome below-target returns. However, when ruinous losses were involved, most DMs were risk averse.

Different methods have been developed for project selection under uncertainty. These methods can be roughly grouped into three categories based on the level of sophistication. The simplest methods are based on the mean. One approach is to specify the expected value for each cash flow and then analyze the problem as if the cash flows were deterministic [22]. This approach is often combined with sensitivity analysis and is one of the most commonly used approaches [16]. A closely related approach uses the cash flow distributions. These distributions are combined across all periods to obtain the discounted convolution [3, 22, 23, 24, 44, 45, 57]. Once this net present value (NPV) distribution is obtained, projects can be compared based on the mean NPV. Another approach is to adjust the discount rate to reflect the level of risk associated with the proposal [1, 7]. The greater the risk, the higher the discount rate used. Once an appropriate discount rate is determined, the analysis can proceed using either of the two approaches just discussed.

The second category of methods are based on the mean and variance of the cash flow distributions. One approach is combine the mean and variance into a single measure by subtracting some fraction of the variance from the mean [46, 55]. The basic idea is that the more risk averse the DM, the larger the fraction of the variance subtracted from the mean. These single measures can then be treated as deterministic values and the problem solved. In a similar fashion, Park and Wu Yeh [47] use the mean, variance, and skewness to develop a criterion for a capital budgeting problem. Another approach based on the mean and variance is to use the NPV distribution to select the project with the greatest mean subject to some risk hurdle [1] (e.g., probability of losing more than $100,000 must be less than ten percent). Finally, a MV-dominance approach can be used [34, 39, 59]. A project is MV-efficient if no other project has a greater mean and equal variance, or a greater than or equal mean and a smaller variance. If more than one project is MV-efficient, some other method must be used to make a final decision.

The third category of methods are based on the cash flow distributions and typically require specification of the DM's utility function. Stochastic dominance can be used to determine if the multi-period distribution of one project dominates the multi-period distribution of another [33], or sometimes if one NPV distribution dominates another. Bawa [4] provides a comprehensive bibliography of the stochastic dominance literature. Rarely will stochastic dominance identify the single best project, due to the strong nature of the dominance conditions that must be satisfied. However, stochastic dominance can frequently be used to screen out the most unattractive opportunities, leaving a set of efficient proposals for the DM to consider further. It also does not require specification of a utility function. The conditions under which stochastic dominance and MV-dominance will yield the same set of efficient projects has been examined by Porter and Gaumnitz [48] and Levy [32]. A different approach, the discounted certainty equivalent approach, requires a specific utility function. The certain monetary equivalent for each cash flow is computed and the analysis is performed as if these values were deterministic [54, 55]. Finally, a time-weighted or horizon EU approach can be used. For the time-weighted approach, the EU is computed for each period and then multiplied by a weighting factor [27, 54, 55]. Periods farther out in time are given less weight. The values for all the periods are summed and this sum is then used as the single measure of desirability. The horizon approach is based on the NPV distribution, requiring only a single period utility function. Brockett and Golden [9] provide a list of the utility functions commonly used for modeling.

Surveys have also been taken to determine which methods are used in practice. Gurnani [17], in reviewing a number of different surveys, found the percentage of firms that explicitly account for risk varied from 15 percent to 78 percent depending on the survey. The most popular methods were to use risk-adjusted discount rates or to shorten the payback period cutoff. Farragher [14], in studying practices of non-industrial firms, found the activity that was rated as most difficult was assessing risk, followed by adjusting procedures for risk. Farragher also found the most popular methods were to adjust discount rates or shorten the payback period.

CHAPTER III

MEAN-VARIANCE DECISION MODEL

The MV decision model envisions a DM who seeks to identify the best sequence of assets to provide a desired service over time. The DM is assumed to: 1) have alternative asset types from which to choose, and 2) be uncertain about the cash flow outcomes that would result from installing and operating each asset types available currently. The MV model is similar to a stochastic replacement model developed by Lohmann [35]. However, there are two primary differences. First, the monetary objective of Lohmann's model is to maximize expected NPV, while the MV model uses the more general objective of maximizing EU. Second, the output the DM receives from each model differs. Lohmann's approach provides the DM with point estimates of the probability that each of the assets currently available is the optimal first asset, along with an estimate of the distribution for the corresponding service lives and NPVs. This results in no clear decision rule for the best course of action to take currently. Thus, for example, if the asset type with the highest probability of being optimal were selected, it could, in fact, be suboptimal even in terms of expected NPV. (An example is provided in Appendix A.) The MV model provides as output the optimal sequence of assets which clearly defines the optimal current decision.

Many different decision procedures have been suggested to solve a wide variety of replacement problems. In the next section, the decision environment and major assumptions that characterize the replacement problems that can be considered by the MV model are described. Next, the five alternative decision procedures considered in this research are discussed. The chapter concludes by describing the performance measures used in the computational experiments to evaluate and compare the five alternative decision procedures.

3.1 Decision Environment and Major Assumptions

A1: The DM is assumed to not be subject to capital rationing and the economic value of the service provided by an asset can be characterized by its NPV distribution. For simplicity, all cash flows were assumed to be expressed after taxes. Without capital rationing constraints, an individual replacement project can be considered independently of all other projects that would otherwise be competing for the available capital. This does not imply that the DM is unconcerned with capital budgeting, only that the current concern is selecting the best alternative for a given project. Once the best alternative is selected, then the DM can consider the capital rationing component and select the best set of replacement projects, and other productive investments, available at the current time. If a utility-based capital budgeting approach, such as the method of Thompson and Thuesen [54], is used, then choosing the best alternative for individual replacement projects based only on EV could result in suboptimal capital allocation.

Summarizing the value of the service provided by an asset by its NPV eliminates the need to develop a decision procedure which explicitly considers individual period cash flows. Not only may it be difficult to forecast the overall individual period cash flows (the convolution of all the different cash flow components for each period), but these will often be correlated across time, only adding to the complexity. Using the NPV would be reasonable for projects in which individual period cash flows were unlikely to cause the DM financial ruin. Since the DM typically will be uncertain about the cash flows to describe each currently available alternative, the overall NPVs are modeled as random variables.

A2: In order to be able to determine the NPV distribution of a sequence of assets, the NPV distribution of each asset in the sequence is assumed to be normally distributed. By making this assumption, the NPV of the entire sequence will also be normally distributed. Few distributions other than the normal are preserved under addition. The ability to easily determine the NPV distribution of a sequence is critical in keeping the procedure used to find the sequence of highest EU tractable. The ease of having to specify only the mean and variance is another advantage of using the normal distribution.

Even if the individual period component cash flow distributions are skewed or bimodal, the resulting NPV may still be approximately normal. Hillier [24] noted that the NPV distribution is the discounted aggregation of many cash flow streams, and while these most likely will not be independent and identically distributed, there are versions of the Central Limit Theorem which exist that require less stringent conditions to hold. This implies:

- The present value of each individual cash flow component (e.g., maintenance expense) may be closer to being normal than the distributions for the cash flows in each period would suggest, because the present value of a component is the discounted sum of these random variables.

- The NPV distribution, being the sum of the present values for all the individual cash flow components, may be closer to normal than is implied by the distribution of the present value of the individual cash flow components.

Hillier concluded that it is often reasonable to assume the NPV distribution to be normal, even when the periodic distributions for each individual cash flow component are far from normal.

The results of Hillier suggest that forecasting the NPVs to be normally distributed is reasonable for a problem which has a number of uncertain cash flow components, as the MV model envisions. For example, sales revenues, labor costs, setup costs, scrap costs, rework costs, warranty costs, material costs, inventory costs, maintenance expenses, utility expenses, and salvage values could all be uncertain. Since many of the uncertain cash flow components have cash flows in each period, random variables will be summed both across time periods and cash flow components. The resulting NPV distribution should be closer to normal than both the individual period cash flows or the cash flow components. Forecasting the NPVs as normally distributed would probably not be valid when the service life of an asset is short (e.g., one period) and there are only a few uncertain cash flow components.

A3: The NPV of an asset being added to a sequence is assumed to be either independent of the NPVs of all other assets in the sequence, or correlated only with the NPV of the immediately preceding asset in the sequence . The latter case will be referred to as "Markov-correlation." Thus, correlation is across time periods, not across assets available in the same period such as stock portfolio problems are typically modeled. There is probably little pragmatic need to develop a more elaborate method to model correlation because Markov-correlation captures the major effects, and the DM must forecast all the higher orders of correlation, which may quite difficult.

Correlation can be helpful in modeling the impact of the external environment on the NPVs. For example, consider an automobile replacement problem. The DM may forecast the NPVs assuming the automobile would be operated mainly on paved surfaces. However, if the automobile is operated frequently on gravel surfaces, then the maintenance costs may increase. If future automobiles were also expected to be operated in an environment similar to the current automobile, then positive correlation could be used to capture the increased maintenance costs for future automobiles.

A4: A single asset is assumed to be in service at all times and there is a finite number of independent assets from which to choose at each decision point. Each asset is assumed to have a known, deterministic maximum service life beyond which it must be replaced. Any monetary consequences of an asset failing before the maximum service life are assumed to be accounted for in the forecasted cash flows.

A5: When considering two investment opportunities with the same expected return, the DM is assumed to prefer the investment with the lower variance. To reflect this risk aversion, the DM's objective is to maximize EU. The DM's utility function is assumed to be a continuous, concave, monotonically increasing function of money. If the DM is only interested in the MV-efficient set of sequences, then the exact form and parameters for the utility function are not required. Expected utility is calculated using the NPV distribution for a sequence of assets, requiring the DM to assess a single period utility function. Such a function will typically be easier to assess than a temporal function. Methods to assess utility functions are discussed by Berger [8] and Keeney and Raiffa [27]. For the computational experiments, the DM's utility function is assumed to be of a specific functional form, so that the alternative decision procedures can be compared.

3.2 Five Alternative Decision Procedures

Many decision procedures have been suggested for solving replacement problems. In the computational experiments, five alternative decision procedures were used, including three of the most common from the literature. The alternative decision procedures include the commonly suggested traditional (TRAD), expected value (EV), and certain monetary equivalent (CME) approaches, in addition to the random and expected utility (EU) approaches. Each decision procedure is described in the next five subsections.

3.2.1 Random

As described further in the next section, the performance measures developed for the computational experiments use the random sequence as the benchmark of performance. Random decision procedures have been similarly used as a benchmark by Baksh [1], Bean, Lohmann, and Smith [5], and VanderVeen [56]. The random sequence is generated using a three step process. First, the asset type is selected using a discrete uniform distribution with endpoints of one and the number of asset types. Second, the service life is selected using a discrete uniform distribution with endpoints of one and the maximum service life for the asset type just selected. The process is repeated until the sequence's service life is greater than or equal to the horizon. Then, if the service life of the last asset selected results in a sequence of assets providing service beyond the horizon, the service life is truncated so the total service provided by the sequence equals the horizon.

3.2.2 Traditional

The most commonly suggested replacement procedure assumes an infinite horizon and identically repeating future assets [16]. Using data for the assets available currently, the DM selects the asset type and service life pair that has the highest annual equivalent value, and then this pair is repeated forever. For simplicity, this approach will be referred to as traditional. The assumptions of an infinite horizon and identically repeating future assets, while making the problem easy to solve, frequently do not model reality well. Bean, Lohmann, and Smith [5] studied a slight modification of traditional replacement. Called sequential traditional replacement, the highest annual equivalent value asset type and service life pair is selected at each replacement epoch, using updated cost forecasts. Bean, Lohmann, and Smith found that when technology and inflation change, the performance of the sequential approach is better than the classical traditional approach, although not as good as more sophisticated decision procedures.

The TRAD sequence is found using the sequential traditional approach. The procedure begins by determining at time $T = 0$ the asset type and service life pair with the highest expected annual equivalent value. Let $\mu(J,T,N) = E[NPV(J,T,N)] = E[NPV(J,0,N)] \times F(J,T) \,/\, (1+m)^T$ be the NPV EV for asset type J, installed at time T, providing N periods of service, and $(A/P,m,N)$ be the capital recovery factor. Then the expected annual equivalent value at time T is calculated as $\mu(J,T,N) \times (A/P,m,N)$. Given the highest value asset type and service life pair, the process is repeated at time $T + N_{\text{highest value pair}}$. Any service life for which $(T + N)$ exceeds the horizon is not considered. The procedure terminates when T equals the horizon.

3.2.3 Expected Value

The EV sequence is that sequence of assets with the highest mean NPV. Deterministic dynamic programming is used to find the EV sequence [42]. The dynamic programming algorithm proceeds in a forward direction, with the stages being the time periods, and the states being the (J, T, N) triplets. The optimal value function, $ENPV^*(T)$, is the maximum possible NPV EV for any sequence providing service for the first T periods. The functional equation is given by

$$ENPV^*(T) = Max\{ENPV^*(T') + \mu(J, T', T - T') :$$
$$T - T' = 1, 2, \ldots, \bar{N}_J,\ T' \geq 0,\ J = 1, 2, \ldots, \bar{J}\}$$

where $\bar{N}_J$ is the maximum service life for asset J and $\bar{J}$ is the total number of asset types. The boundary condition is $ENPV^*(T = 0) = 0$.

3.2.4 Certain Monetary Equivalent

Deterministic dynamic programming is also used to find the CME sequence. The only difference between the EV and CME approaches lies in how the optimal value function is calculated. Let $\sigma^2(J, T, N) = Var[NPV(J, T, N)] = Var[NPV(J, 0, N)] \times F(J,T)^2 \ / \ (1 + m)^{2T}$ be the NPV variance for asset type J, installed at time T, providing N periods of service, $f_{npv}(w|\mu, \sigma^2)$ be the NPV density function, $U(w)$ be the DM's utility function, and $EU(J, T, N) = \int U(w)\ f_{npv}(w|\mu(J, T, N), \sigma^2(J, T, N))\ dw$ be the EU for asset type J, installed at time T, providing N periods of service. For the computational experiments, the value of $EU(J, T, N)$ was calculated using numerical integration with the limits of integration set at plus and minus ten standard deviations from the mean. The CME for asset type J, installed at time T, providing N periods of service is then $CME(J, T, N) = U^{-1}(EU(J, T, N))$. The optimal value function, $CME^*(T)$, is the maximum possible CME for any sequence providing service for the first T periods. The functional equation is given by

$$CME^*(T) = Max\{CME^*(T') + CME(J, T', T - T') :$$
$$T - T' = 1, 2, \ldots, \bar{N}_J,\ T' \geq 0,\ J = 1, 2, \ldots, \bar{J}\}$$

The boundary condition is $CME^*(T = 0) = 0$.

In the case of assets being correlated, the method to calculate the variance is modified in an attempt to capture the covariance. Suppose asset type J_1, installed at time T_1, for N_1 periods of service is followed by asset type J_2, installed at time T_2, for N_2 periods, and that J_2 is correlated with J_1. The variance is then calculated as $\sigma^2(J_2, T_2, N_2) + 2\rho_{J_1J_2}\sigma(J_1, T_1, N_1)\sigma(J_2, T_2, N_2)$, with $\rho_{J_1J_2}$ being the correlation coefficient for J_2 following J_1. If the result is negative, the variance is set equal to zero and $CME(J_2, T_2, N_2) = \mu(J_2, T_2, N_2)$.

3.2.5 Expected Utility

The EU sequence is generated using a branch and prune procedure, or if necessary, the cluster heuristic which is described in the next chapter. If the branch and prune procedure is unable to find a sequence before exceeding the allocated computer memory, the cluster heuristic is invoked to find a "good" utility sequence. The branch and prune procedure is similar to the forward dynamic programming algorithms used to find the EV and CME sequences, except that instead of a single best partial sequence always being found at each stage, a MV-efficient set of partial sequences is found. A sequence is MV-efficient if no other sequence has a greater mean and equal variance, or a greater than or equal mean and a smaller variance.

The branch and prune procedure is a multiple objective dynamic program. The objectives are to maximize the mean and minimize the variance of the NPV. To do so, the branch and prune finds all the MV-efficient sequences and then uses a utility function to select the sequence of highest EU. The set of MV-efficient sequences can equivalently be called the Pareto-optimal, non-dominated, or admissible set of paths. Because of the computationally intensive nature of finding all the Pareto-optimal paths, no study found in the literature identified all the Pareto-optimal paths, except for small problems (e.g., three stages with no more than three states at each stage) used for illustrative purposes. Instead a variety of other approaches have been used to solve multiple objective problems. One approach is to find the path with the highest probability of being shorter (longer) than all other paths [52], or the path with the highest probability of achieving a given target [21]. A second approach is to find the MINSUM or MINMAX path [11, 19]. The MINSUM path minimizes the weighted sum of the deviations from a target value for each objective, while the MINMAX path minimizes the weighted maximum deviation from any target value. Finally, a heuristic can be used to find a "good" subset of Pareto-optimal paths [2, 20]. Thus, this research is the first study to attempt to find all the Pareto-optimal paths for problems considered to be of realistic size.

Stochastic dominance rules from portfolio selection theory can be used to prove that a sequence is MV-efficient if no other sequence has a greater mean and equal variance, or a greater than or equal mean and a smaller variance. While the proofs below assume the random variables are normally distributed, the results can also be proved for random variables from a restricted class of two-parameter distributions [18, 28]. Following the approach of Park and Sharp [46], given two random variables X_1 and X_2 with cumulative distribution functions $F_1(x)$ and $F_2(x)$, $F_1(x)$ dominates $F_2(x)$ (or equivalently, X_1 dominates X_2) if $E[U(X_1)] \geq E[U(X_2)]$ for every utility function in the class of utility functions and strict inequality holds for a least one function in the class. For a DM with a nondecreasing, finite valued utility function, $F_1(x)$ dominates $F_2(x)$ by first-order stochastic dominance (FSD) if and only if $F_1(x) \leq F_2(x)$ for all x and strict inequality holds for some x. Adding the restriction that the DM is risk averse, $F_1(x)$ dominates $F_2(x)$ by second-order stochastic dominance (SSD) if and only if $\int_{-\infty}^{x} F_1(t)dt \leq \int_{-\infty}^{x} F_2(t)dt$ for all x and strict inequality holds for some x.

Result 1 (Equal Variance): For a DM with a nondecreasing, finite valued utility function comparing $X_1 \sim Normal(\mu_1, \sigma^2)$ and $X_2 \sim Normal(\mu_2, \sigma^2)$, X_1 will be preferred if $\mu_1 > \mu_2$.

Proof: If $\mu_1 > \mu_2$, then $Z_1 = (x - \mu_1)/\sigma < Z_2 = (x - \mu_2)/\sigma$ for all x. This implies $F_1(x) < F_2(x)$ for all x, and thus by FSD, X_1 dominates X_2.

Result 2 (Smaller Variance): For a risk averse DM with a nondecreasing, finite valued utility function comparing $X_1 \sim Normal(\mu_1, \sigma_1^2)$ and $X_2 \sim Normal(\mu_2, \sigma_2^2)$, X_1 will be preferred if $\mu_1 \geq \mu_2$ and $\sigma_1 < \sigma_2$.

Proof: Since $\mu_1 \geq \mu_2$ and $\sigma_1 < \sigma_2$, the cumulative distribution functions $F_1(x)$ and $F_2(x)$ will intersect. Let the point of intersection be denoted as x_0. For $x < x_0$, $Z_1 = (x - \mu_1)/\sigma_1 < Z_2 = (x - \mu_2)/\sigma_2$ which implies $F_1(x) < F_2(x)$, and thus $\int_{-\infty}^{x_0^-}(F_2(t) - F_1(t))dt > 0$. For $x > x_0$, $Z_1 > Z_2$ which implies $F_1(x) > F_2(x)$, and thus $\int_{x_0^+}^{x}(F_2(t) - F_1(t))dt < 0$. To determine the sign of $\int_{-\infty}^{x}(F_2(t) - F_1(t))dt$ for $x > x_0$, note that since μ_1 - $\mu_2 \geq 0$, μ_1 - $\mu_2 = \int x(f_1(x) - f_2(x))dx \geq 0$. Using integration by parts with $u = x$ and $dv = (f_1(x) - f_2(x))dx$, $du = dx$, $v = \int f_1(x)dx$ - $\int f_2(x)dx = 1$ - $1 = 0$, so uv - $\int v du = 0$ - $\int[\int(f_1(x) - f_2(x))dx]du =$ -$\int(F_1(x) - F_2(x))dx = \int(F_2(x) - F_1(x))dx$. Thus, $\int(F_2(x) - F_1(x))dx \geq 0$. If $\int(F_2(x) - F_1(x))dx \geq 0$, $\int_{-\infty}^{x_0^-}(F_2(t) - F_1(t))dt > 0$, and $\int_{x_0^+}^{x}(F_2(t) - F_1(t))dt < 0$, then it must be the case that $\int_{-\infty}^{x}(F_2(t) - F_1(t))dt \geq 0$ for all x and obviously strict inequality holds for all $x <$

x_0. Thus by SSD, X_1 dominates X_2.

The branch and prune procedure proceeds in a forward direction, with the stages being the time periods. The states are the (J, T, N) triplets. At any given stage there can be more than one node (partial sequence) that is MV-efficient. In order to be able to backtrack to find the EU sequence, each MV-efficient node is labeled. At each time T, a MV-efficient set of sequences is built using a two step process. In the first step, the service life N is fixed and all the asset types available at $T' = T - N$ are compared. The purpose of this step is to eliminate assets which are MV-dominated by others for the given N. Since assets can evolve over time, this process must be repeated for each T. The result is a MV-efficient set of assets installed at time T' providing N periods of service. One MV-efficient set is formed for each possible N, with the total number of sets equal to the smaller of T and the maximum $\bar{N}_J$ for any J. In the case when assets are correlated, the MV-efficient set is formed using the variance of the asset to be added plus the covariance component. For example, if J_2 is the asset to be added and it follows J_1, then the variance term is given by $\sigma^2(J_2, T_2, N_2)$ + $2\rho_{J_1J_2}\sigma(J_1, T_1, N_1)\sigma(J_2, T_2, N_2)$, where $\rho_{J_1J_2}$ is the correlation coefficient for J_2 following J_1. Negative values, unlike the case for the CME procedure, are maintained since the variance of an entire sequence should be nonnegative for logical DMs. If sequences with negative variances are possible based on the DM's forecasts, then the DM should adjust the correlation coefficients or variances specified.

The second step is to add the MV-efficient set of assets for each possible N to the MV-efficient set of sequences at time T', to form an MV-efficient set of sequences at time T. This process is illustrated in Figure 3.1. For a given N, the normally distributed NPV of each MV-efficient asset is added to the normally distributed NPV of each MV-efficient sequence at time T', and the set of MV-efficient sequences

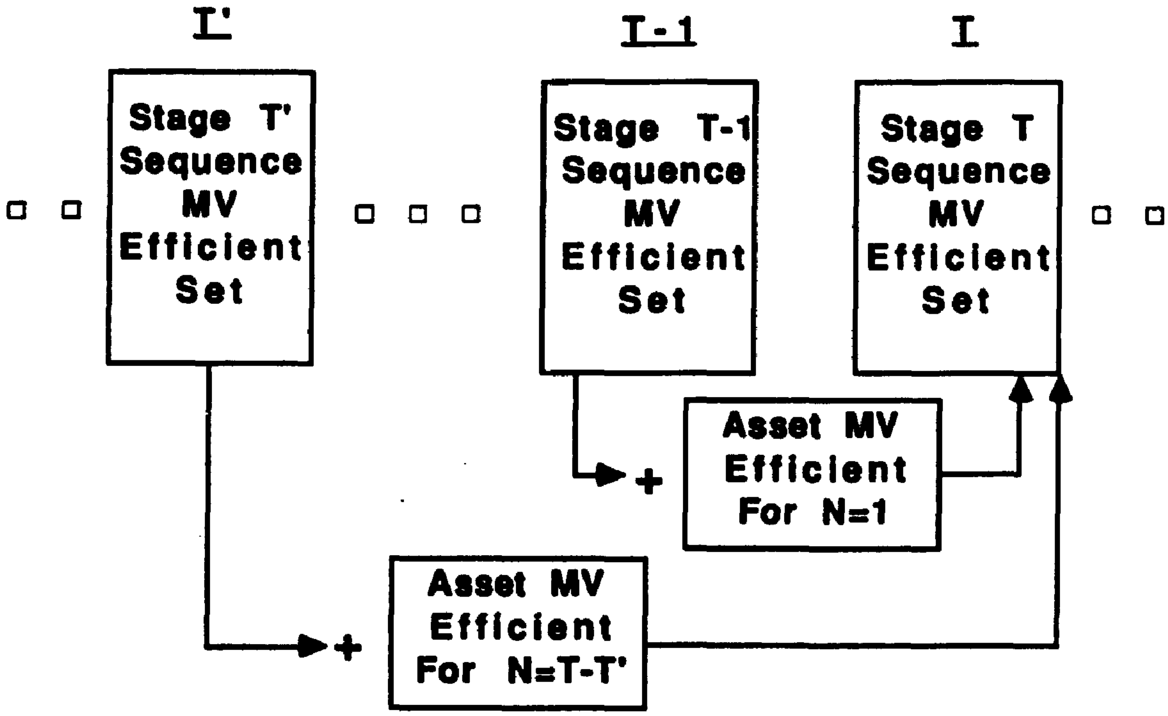

Figure 3.1: Branch and Prune Procedure

at time T is updated. The updating involves adding the resulting sequence to the MV-efficient set, or adding the sequence and then pruning other previously MV-efficient sequences now MV-dominated, or simply pruning the sequence. For the computational tests, the second step started with the maximum possible N and then moved towards $N = 1$. The rationale was that sequences in which the last asset had a short service life would tend to be dominated by sequences in which the last asset had a longer service life, due to the first cost being spread over a longer period of time. If that were the case, then many branches could be pruned without having to be added to the MV-efficient set. Once the MV-efficient set for the final stage has been formed, the EU for each efficient sequence is found using numerical integration. The EU sequence is then found by backtracking from the node at the final stage with the highest EU.

For independent assets, the branch and prune is an optimal algorithm. However, for the case of Markov-correlation, the branch and prune is only a heuristic procedure. For this case, the branch and prune's performance was evaluated against an upper bound as part of the computational experiments. The upper bounding procedures are described in the next chapter. The branch and prune is not an optimal algorithm because the optimal path can be pruned if it is not MV-efficient at each replacement epoch. Consider the following example with a horizon of two periods. At time zero, two asset types, 1 and 1$'$, are available, both with identical one period service lives. At time one, asset type 2 for one period of service only is available. Let $\mu_1 = \mu_{1'}$, $\sigma_1^2 = 64$, $\sigma_{1'}^2 = 81$, $\sigma_2^2 = 100$, $\rho_{12} = 0.5$, and $\rho_{1'2} = 0.1$. Comparing the sequence of asset type 1 followed by asset type 2, [(1,0,1),(2,1,1)], with [(1',0,1),(2,1,1)] the means are identical and thus can be ignored. For [(1,0,1),(2,1,1)], the sequence variance is 64 + 100 + 2(0.5)(8)(10) = 244. For [(1',0,1),(2,1,1)], the sequence variance is 81 + 100 + 2(0.1)(9)(10) = 199. Thus, the optimal sequence is [(1',0,1),(2,1,1)]. However, the branch and prune procedure would select [(1,0,1),(2,1,1)] because at time 1, (μ_1, σ_1^2) MV-dominates $(\mu_{1'}, \sigma_{1'}^2)$.

For the computational tests, only problems with finite horizons were considered. For the infinite horizon case, a stopping rule to find an equivalent finite horizon is not clear for the EU maximizer. If replacement were to occur in a given period, and if assets were assumed to be independent, then one of the MV-efficient partial sequences would match the optimal sequence through the given period. However, the matching partial sequence could not be identified a priori. If assets were assumed to be correlated, then there is no guarantee that any of the MV-efficient partial sequences would match the optimal sequence through the given period. Thus, in general it is impossible to identify an equivalent finite horizon. It may be possible to identify an equivalent finite horizon in several special cases when assets are assumed to be independent, such as if there is only a single MV-efficient node at each stage under consideration, or if all the MV-efficient nodes at each stage are part of sequences that begin with the same asset type.

3.3 Performance Measures

It is difficult to develop a good performance measure to compare the sequences identified using the alternative decision procedures. Because the EU of a sequence can be positive, zero, or negative, a simple fraction such as EU(sequence being compared) / EU(EU sequence) can be difficult to interpret and may be undefined. The difficulty of interpretation results because such a fraction does not have the conventional range of zero to one, but rather negative infinity to positive infinity. In related research, to avoid a similar difficulty, Bean, Lohmann, and Smith [5] used a random sequence as a benchmark level of performance by which other decision procedures were compared. For the MV model, using a random sequence as a benchmark results in a performance measure of the form

$$\frac{EU(\text{sequence being compared}) - EU(\text{random})}{EU(\text{EU sequence}) - EU(\text{random})} \tag{3.1}$$

By setting (3.1) to zero should the random sequence outperform the sequence being compared, (3.1) will range from zero to one.

However, (3.1) may be unstable if based on a single random sequence. If the EU for random sequences were to vary considerably, interpretation of (3.1) may be difficult (i.e., the performance measure would tend to be quite noisy). Three alternative modifications to (3.1) were studied empirically. The first alternative was to use the sequence with highest EU out of a set of random sequences, the second was to use the average EU of the set, and the last was to use the first random sequence

Table 3.1: Random Sequence Performance Measure Variance

Problem Set	Number of Observations	Best of 100 Variance	Avg. of 100 Variance	First of 100 Variance	Best of 100 Avg. Ratio
1+	11	0.1235	31.3520	393.2683	0.7920
1-	19	0.0909	989.9619	86.9267	1.2691
2+	4	0.0519	0.1069	0.2609	0.7972
2-	26	0.0350	1.2158	0.4409	1.1200
3+	9	0.6773	0.9063	0.5870	0.6888
3-	21	0.0142	43.9928	14.9985	1.1543
4+	15	0.0026	2.5805	2.5805	0.9501
4-	15	0.0284	1.5660	3.3647	1.1424

generated. The ratio EU(random) / EU(EU sequence) was computed for each of four sets of 30 problems. The variance of this ratio was then used as a measure of stability. The lower was the variance, the higher was the stability. The first problem set had independent assets with "low" uncertainty, the second problem set had independent assets with "high" uncertainty, and the last two problem sets had low uncertainty with negative and positive correlation, respectively. Random sequence sets of size 10, 100, and 1000 were considered. While the results for a size of 100 were clearly better than for a size of 10, sets of size 100 and 1000 did not produce statistically different results. The results for a size of 100 are shown in Table 3.1. The four problem sets were divided into two categories based on the sign of the EU sequence's mean NPV. For the positive EU case, the measure was always less than or equal to one, and for the negative EU case it was always greater than or equal to one. This was done to avoid some of the problems of interpretation discussed above. The positive and negative categories are identified in the "Problem Set" column in Table 3.1 by the "+" and "-" symbols. Note from the "Number of Observations" column that more of the problems generated were of the negative EU case than the positive EU case. A comparison of the variance for the three alternatives shows the best of 100 random sequences to have the lowest variance, except for the case, "3+."

Because it appeared to be the most stable, the best random sequence out of 100 was the alternative used for the computational experiments. The resulting "utility performance" measure was the following.

$$\frac{EU(\text{sequence being compared}) - EU(\text{best of 100 random})}{EU(\text{EU sequence}) - EU(\text{best of 100 random})} \tag{3.2}$$

The value of (3.2) was set to zero if the random sequence had an equal or higher EU value than the sequence being compared. In order to evaluate the quality of the EU decision procedure for the cases when it was not guaranteed to be optimal (i.e., when assets were correlated or the cluster heuristic had to be used), the "EU to upper bound utility performance" measure was developed:

$$\frac{EU(\text{EU sequence}) - EU(\text{best of 100 random})}{EU(\text{upper bound}) - EU(\text{best of 100 random})} \tag{3.3}$$

The value of (3.3) was set to zero if the random sequence had an equal or higher EU value than the EU sequence.

Table 3.2 shows the values of (3.2) for three problems taken from the computational experiments. For each problem, the sequence type, mean, variance, EU, and utility performance are displayed. For the first problem, despite the EV and EU sequences being very close in mean, variance, and EU, the performance for the EV sequence was 0.0. This value resulted because the best of 100 random had an EU value greater than 97.2172. For the second problem, both the EV and CME sequences had positive performance values. A comparison of the EV sequences for the first two problems shows that the performance value was lower for the first problem (0.0 vs. 0.1801) despite a smaller absolute EU difference (0.0078 vs. 0.5661) as compared to the matching EU sequence. A comparison of the EV sequence for the first problem and the CME sequence of the second problem shows that a similar absolute EU difference (0.0078 vs. 0.0216) resulted in a much lower performance

Table 3.2: Sample Utility Performance Values

Problem	Sequence	Mean	Variance	EU	Utility Performance
1	EV	213.3180	1564.4730	97.2172	0.0000
	EU	213.3053	1550.9050	97.2250	
2	EV	73.4461	810.9745	39.4817	0.1801
	CME	73.1684	577.2316	40.0262	0.9686
	EU	73.1353	565.2512	40.0478	
3	EV	-65.8346	73.3013	-227.7407	0.9505
	EU	-65.8932	63.6304	-226.9904	

value (0.0 vs. 0.9686) for the EV sequence. For the third problem, the EV sequence's performance was about the same as the CME sequence's performance in the second problem (0.9505 vs. 0.9686), but the absolute EU difference was about 35 times larger. As these three examples illustrate, the performance measure was rather insensitive to the absolute EU difference.

The best random sequence out of 100 might be interpreted as representing an intelligent DM who uses experience and judgement in making replacement decisions, but no sophisticated techniques. As the last column of Table 3.1 shows, the average ratio for the best sequence out of 100 ranged from 1.1200 to 1.2691 when the EU for the EU sequence was positive, and from 0.6888 to 0.9501 when negative. Thus, had the DM selected the best of 100 random sequences, the resulting EU would have been, on average, 69% to 95% of the maximum EU possible.

CHAPTER IV

COMPUTATIONAL EXPERIMENTS

A series of computational experiments were conducted to gain insight about replacement decision making under uncertainty and to compare the alternative decision procedures. A test problem generator with a set of parameters was used to randomly generate problems so that a wide range of potential problems could be considered. In total, five experiments were performed. The first experiment indicated that the branch and prune procedure could solve optimally many problems when assets were assumed to be independent. However, for some problems the allocated computer memory was exceeded. In the second experiment, the cluster heuristic developed to overcome the memory problem and the heuristic's associated upper bounding procedure were evaluated. The results showed both to perform well. Some key model parameters were then studied in the third experiment in an attempt to determine some of the conditions that influence the performance of the alternative decision procedures. The fourth experiment considered three different types of utility functions and found that the corresponding sequences of highest EU were similar when the risk aversion levels were similar. Finally, the fifth experiment examined the sensitivity of the EU sequence to changes in the NPV forecasts. The results indicated that changes in the mean have a much greater effect than similar changes in the variance.

The remainder of the chapter is divided into three sections. First, the test problem generator used for the experiments is described. Next, the design and numerical results for each of the five experiments are discussed. The chapter concludes with a summary of the parameter effects.

4.1 The Test Problem Generator

Different approaches were considered in attempting to meet the research objectives of gaining insight into replacement decision making under uncertainty and comparing the performance of the alternative decision procedures. First considered was use the of problems from the literature. However, no applicable problem sets could be found. Most replacement research has considered only deterministic cash flows and the work done with stochastic cash flows assumed the DM desired to maximize EV. A second approach considered was to collect real data. The difficulty of this approach lies in the amount of time required to collect data for enough problems to be able to draw some general conclusions. A third approach, and the one used here, was to randomly generate a wide range of problems that are believed to be representative of typical replacement problems.

In developing the test problem generator, the key question was how to assign values to the model parameters. Some parameters can be set by drawing a value directly from a distribution, while others require more extensive calculations. The first subsection below lists the all the parameters required as input by the sequence generator, and specifies the ranges for those parameters that can be set simply by drawing a value. Then, the four subsequent subsections discuss how the remaining parameters were set. These parameters include the level of risk aversion, correlation coefficients, level of uncertainty, and NPVs.

4.1.1 Test Problem Generator Parameters

The MV model requires as input the discount rate, level of risk aversion, horizon, number of asset types, maximum service life for each asset type, NPV means and variances for all asset type and service life combinations, rate of technological change for all asset types, and correlation coefficients. The parameters which can be set by drawing a value directly from a distribution include the discount rate, horizon, number of asset types, maximum asset service life, rate of technological improvement, and asset monetary difference.

For the computational experiments, the after tax **discount rate** ranged from 15% to 30%. Real firms may use a discount rate that exceeds 30%. However, only problems with lower discount rates were considered since the higher rates reduce significantly the impact of future decisions on the current decision. With very high rates, it becomes almost irrelevant what assets are installed beyond the first asset, thus reducing the need for a sophisticated procedure.

The **horizon** ranged from 10 to 50 periods. While real problems could have horizons less than ten periods, such problems will typically be easy to solve and provide less insight.

The **number of asset types** ranged from 2 to 7. Two is the minimum (unless the DM's objective is only the optimal length of service), while seven was considered a reasonable maximum. For real decisions, some alternatives may prove to be technologically infeasible, others often can be eliminated without detailed analysis. Additionally, the time and effort required to collect the necessary information for a large number may be overwhelming.

The **maximum asset service life** ranged from 2 to 14 periods. Perhaps more important than the range is the relationship between the asset service life and the horizon. With short horizons the service life could actually exceed the horizon, while for longer horizons a number of replacements could be necessary.

For the computational experiments, assets available in the future are assumed to be related monetarily to assets currently available in a deterministic manner. The NPV for future versions of asset type J, installed at time T, providing N periods of service is given by

$$NPV(J,T,N) = \frac{NPV(J,0,N)F(J,T)}{(1+m)^T}$$

where $F(J,T)$ is the deterministic **rate of technological improvement** and m is the discount rate. By relating future assets to assets currently available in a deterministic manner, the set of assets available at each time period can be generated easily. However, the decision procedures are set up to accept forecasts for any finite number of assets available at each point in time if the DM desires to provide such detail. For the computational experiments, the DM was assumed to be able to provide a deterministic discount rate that represents the minimum rate of return expected for future but as yet unknown productive opportunities. The discount rate is also commonly called a growth rate, the cost of capital, or the minimum attractive rate of return. If all future productive opportunities were known completely, there would be no need for a discount rate [41]. However, in the more realistic situation of incomplete information, the DM will know typically only the proposals available currently, but not precisely what future opportunities will become available. For purposes of this research, the discount rate represents the minimum rate of return expected from these future but as yet unknown productive opportunities. A common approach to determine its value is to examine historical records and determine the lowest average rate of return generated by productive investments for similar replacements. Oakford, Salazar, and Diguilio [43] provide a more detailed description of factors to consider in determining a reasonable value for the discount rate.

The rate of technological improvement ranged from 0% to 30% per period. Thus, the $F(J,T)$'s are assumed to be geometric functions of the form $(1+r)^T$, where r is the rate of technological improvement. Ganesh [15], after examining a number of previous studies and data from industry, concluded that the geometric is "a reasonable representation of productivity improvement." With a rate of up to 30%, technology may be increasing either faster or slower than the discount rate, allowing a wide range of problems to be considered.

The **asset monetary difference** ranged from 0% to 30%. This parameter is used to limit how monetarily different the asset types are. If problems are randomly generated with large monetary differences among the alternative asset types, then these problems will typically be easy to solve, may not require a sophisticated decision procedure, and will provide less insight. The asset monetary difference is used when computing the NPVs and the rate of technological improvement. For each problem a base set of NPVs is computed. The base set is then perturbed, using the asset monetary difference value, to generate the NPVs for the individual asset types. Additionally, a base rate of technological improvement is perturbed to generate the rate of technological improvement for the individual asset types. The methods used to generate the NPVs are explained in more detail below. An upper limit of 30% was selected because problems with differences any greater will typically be easy to solve.

For most of the experiments, the ranges for the parameters were split into sub-ranges, and the combined sub-ranges might fail to span the entire range listed above. All such deviations are noted. It turns out that the ranges selected for some of the parameters is not of great importance, as the results and observations from the five experiments will illustrate.

4.1.2 Setting the Level of Risk Aversion

Before the DM's level of risk aversion can be set appropriately, an estimate of the monetary range that the utility function must span is needed. The NPV distribution for the EV sequence can be used to provide this estimate. The EV sequence always corresponds to the MV-efficient sequence with the highest mean and variance found by the EU decision procedure. Since the NPV for each asset in the EV sequence is assumed to be normally distributed, the EV sequence's NPV distribution is also normal. An estimate for the relevant range can be made by truncating the EV sequence's NPV distribution 3.5 standard deviations above and below the mean. The upper endpoint will, with probability 0.9999, never be exceeded since it is based on the maximum mean and maximum variance. However, the lower estimate may be too high. A more accurate and elaborate method to estimate the lower endpoint would be to find the minimum NPV lower endpoint of all the MV-efficient sequences that the EU decision procedure found.

Given the relevant range the utility function must span, in practice the DM could assess subjectively an appropriate utility function. For example, the lower endpoint could be assigned a utility of zero, the upper endpoint a utility of one, and then the points in between assessed via lotteries. While such an approach could be used in practice, the test problems were designed in a manner considered to be representative of a wide range of DMs. Except for the fourth experiment, it was assumed that the DM's utility function was an exponential of the form $U(w) = (1 - e^{-cw})/c$, where c denotes the level of risk aversion. By definition, $r(w) = -U''(w)/U'(w)$ is the risk aversion function [49]. The exponential form was selected for two reasons. First, it has constant risk aversion and thus the initial wealth of the DM does not need to be known. Second, the more general exponential function of the form $U(w) = a - be^{-cw}$ reduces to the above form if the approach of Thompson and Thuesen [55] is followed by setting $U(0) = 0$ and $U'(0^+) = U'(0^-) = 1$. This eliminates the need to set more than one parameter. Kallberg and Ziemba [26] provide results for portfolio selection problems which indicate that utility functions of differing functional forms and parameter values yield similar optimal portfolios if they have similar absolute risk aversion (at least for short horizons).

The approach used to set c for the experiments was based on the EV sequence's NPV. The endpoints for the range of monetary consequences were set by truncating this NPV distribution 3.5 standard deviations above and below the mean. The endpoint of largest absolute value ($MAXAB$) was used to set $c = ln(z)/MAXAB$ where z ranged from 1.5 to 20. The more risk averse the DM, the greater the value for z. Unless the EV sequence's mean NPV was zero, the utility function was assessed over a range wider than the relevant range. The range for z was based on data collected from students assessing their personal utility. Using lotteries, fifteen students assessed their utility function over a range from \$0 to \$100. Each student's utility function was then fit to an exponential utility function of the form $U(w) = 1 - e^{-cw}$. Regression was used to determine the value of c giving the best fit. The resulting c values were were consistent with both examples in Thompson and Thuesen [55] and results in Kallberg and Ziemba [26]. The range for z ($= e^{cMAXAB}$) was determined by setting $MAXAB = 100$ and using the minimum and maximum values found for c.

4.1.3 Setting the Correlation Coefficients

The MV model assumes assets in a sequence to be either independent or Markov-correlated. Thus, the test problem generator generates correlation coefficients for assets that are adjacent in a sequence. Two different schemes were considered to set the correlation coefficients. Both assumed that assets of the same type had the highest correlation coefficient and that the correlation coefficient was identical for all the possible associated service lives. The first assumption was made since it seemed unlikely that two different asset types would be more positively correlated than two assets of the same type. The second assumption was made to reduce the amount of input data required and because there was little to gain from the additional complexity. However, the MV model is set up to accept a specific value for each asset type and service life pair if the DM desires to provide such detail.

The first scheme was quite simple. For the positive correlation case, the coefficient for sequential assets of the same type was set equal to $0.1\bar{J}$, where $\bar{J}$ is the number of asset types. For sequential assets of different types, the coefficient was set equal to $0.1\bar{J} - (0.1 \times du(1, \bar{J} - 1))$, where $du(1, \bar{J} - 1)$ is a realization from a discrete uniform distribution with endpoints of 1 and $\bar{J} - 1$. For the negative correlation case, the coefficient for assets of the same type was set to 0.0. For sequential assets of different types, the coefficient was set equal to $-(0.1 \times du(1, \bar{J} - 1))$. A sample set of coefficients for three asset types is shown in Table 4.1 where asset type J_2 follows type J_1.

There was, however, a problem with this scheme. It was possible to have sequences with covariance terms that sum to a value more negative than the sum of all the positive variance terms. To avoid such illogical negative variances, a more sophisticated scheme was used to set the correlation coefficients for the computa-

Table 4.1: Correlation Coefficient Matrix for Three Asset Types

Positive			
	J_2		
J_1	1	2	3
1	0.3	0.2	0.1
2	0.1	0.3	0.2
3	0.2	0.1	0.3

Negative			
	J_2		
J_1	1	2	3
1	0.0	-0.1	-0.2
2	-0.2	0.0	-0.1
3	-0.1	-0.2	0.0

tional experiments. Let $u(a,b)$ represent a random draw from a uniform distribution with endpoints a and b. Then for the positive correlation case, the coefficient for sequential assets of the same type was set equal to $u(0,1)$. The remaining values for the coefficients for the current asset type followed by an asset of a different type were set by multiplying $u(0,1)$ times the coefficient just set for sequential assets of the same type. For the negative correlation case, the coefficient for sequential assets of the same type was set equal to zero. The remaining values for the coefficients for the current asset type followed by an asset of a different type were then set equal to $u(0,0.5)$, with the sign based on the following observation for any partial sequence of three assets. If the first two assets are positively correlated, and the second and third are also positively correlated, then to be consistent the coefficient for a sequence with the first asset followed by the third asset should also be positive. Table 4.2 shows all the possible patterns for any partial sequence of three assets. The first column in Table 4.2 is the sign of the correlation coefficient for the first and second asset in the sequence, the second column the sign for the second and third, and the last column for the first and third. Note that for more than two asset types, all the coefficients

Table 4.2: Correlation Coefficient Sign for Negative Case

1st and 2nd	2nd and 3rd	1st and 3rd
+	+	+
+	-	-
-	+	-
-	-	+

cannot be negative, meaning the negative case will typically be a mixture of positive and negative values. Even with all combinations of three assets being consistent in sign, to assure no sequences would have negative variances, the maximum absolute coefficient value was limited to 0.5.

4.1.4 Measure of Uncertainty

Previous capital budgeting work that has examined risk has found that until the level of uncertainty is "high" enough, there is little or no difference in the selections made by risk and non-risk adjusted procedures [1]. Thus, an important characterization of any problem is the level of uncertainty. Since the MV model requires as input the NPV distribution for each asset type and service life pair, the level of uncertainty should be characterized based on these inputs. In their capital budgeting work, Thompson and Thuesen [55] characterized the amount of uncertainty in terms of the variance of the initial capital investment. Since replacement problems typically have a first cost, uncertainty could be based on the first cost variance. However, such an approach can result in widely varying values for the NPV coefficient of variation. For example, suppose the first cost is 1.0, the variance is set to three times the first cost, the revenues for the first two periods equal 0.6, the horizon is two periods, the asset has no salvage value, and the discount rate is zero. The resulting NPV mean equals 0.2 and the variance equals 3.0. The coefficient of variation is 8.66. Now for the same example, assuming the horizon is four periods and the revenues are 0.6 in each of these periods, the resulting NPV mean equals 1.4 and the variance again equals 3.0. The resulting coefficient of variation is 1.24.

For the experiments, a slightly different approach was used. Uncertainty was characterized by the median NPV coefficient of variation across all asset types and service lives. This value was estimated for the test problem generator using a sample set of problems. Both the first cost and the annual costs were allowed to be uncertain, since a DM will typically also have uncertainty about the costs in future periods. A model that limits uncertainty to only the first cost is unnecessarily restrictive. For both cost parameters, the standard deviations were assumed to be uniformly distributed between $uncert_a$ and $uncert_b$ times the realization for the cost parameter mean. For the first two experiments, uncertainty was studied at only a low level and thus a single ($uncert_a$, $uncert_b$) pair of values was used. For the last three, uncertainty was studied at both a low and a high level, requiring two ($uncert_a$, $uncert_b$) pairs of values. Low uncertainty was meant to represent the point at which uncertainty was just large enough to start causing a few EV sequences to differ from the associated EU sequences, while high uncertainty was meant to represent a point at which the EV and EU sequences frequently differed. The values used for $uncert_a$ and $uncert_b$ are listed for each experiment. Since the NPV coefficient of variation has a lower bound of zero, the distribution will tend to be positively skewed. Therefore, the median rather than the average ratio was used.

The NPV coefficient of variation represents the overall, or absolute level, of uncertainty. Another important concept is that of relative uncertainty. This refers to the amount of difference in the variance across asset types and service lives. Even when the overall level of uncertainty is high, if there is very little relative uncertainty, then the EV and EU sequences will tend to match. There has to be sufficient reduction in the variance before a sequence with a lower mean will be selected, even for highly risk averse DMs.

4.1.5 Generating the NPVs

Two methods, the "regular" method and the "annual equivalent" (AE) method, were used to generate the NPVs. The regular method was used for all but the fourth experiment, which, due to special requirements, necessitated using the AE approach. For both methods, once the number of asset types, maximum service life for each asset type, rate of technological improvement for the first asset type, and horizon were randomly drawn, a base set of NPVs was generated. The base set was then perturbed to generate the NPVs for all the other asset types. Since the MV model assumes the NPVs are normally distributed, only mean and variance terms had to be generated.

For the regular method, the NPVs were calculated based on three cash flow components: the first cost, the annual operating cost (or revenue if this cost was negative), and the salvage value. The mean first cost was used as the basis for generating all the other cash flows. The asset's mean first cost was assumed to be uniformly distributed between 1 and 100, which allowed a range of problems to be generated. Once a value for the mean first cost was drawn, the standard deviation of the first cost was assumed to be uniformly distributed between $uncert_a$ and $uncert_b$ times this realization. The mean annual cost for the first period was assumed to be uniformly distributed between -0.75 and 0.75 times the mean first cost realization. This allowed problems to be generated in which annual costs exceed revenues (or revenues are not considered) as well as the reverse case. The standard deviation of the first period's annual cost was assumed to be uniformly distributed between $uncert_a$ and $uncert_b$ times the first period's mean annual cost realization. Annual costs for future periods were assumed to change geometrically based on the first period (i.e., $AC(T) = (1+g)^{T-1} \times AC(T=1)$ where AC is the annual cost, T the time period, and g the rate of geometric change). To allow for both increasing and decreasing cost patterns, g was assumed to be uniformly distributed between -0.5 and 0.5. The salvage value was assumed to decline geometrically with g uniformly distributed between -1.0 and 0.0 (i.e., $SV(T) = (1+g)^T \times FC$ where SV is the salvage value and FC the first cost). Given the first cost, annual costs, and salvage value, the NPV means and variances for the base set were then computed for each possible service life up to the maximum of any asset type.

For the AE method, the base asset NPVs were computed using a randomly selected value for the mean AE. The mean AE was assumed to be uniformly distributed between 100 and 10,000. Once a value, ae, for the mean AE was drawn, the base mean NPV for each possible asset type and service life combination was set equal to $ae \times u(0.75, 1.25) \times (P/A, m, N)$, where $(P/A, m, N)$ is the present value factor for N periods of service and discount rate m, calculated as shown in (4.1). The value $u(0.75, 1.25)$ was used to generate some economic difference among the different service lives for the same asset type. The standard deviation of the AE was assumed to be uniformly distributed between $uncert_a$ and $uncert_b$ times the mean AE realization. Given a realization, ae_{std}, for the standard deviation of the AE, the variance of the base NPV for each possible asset type and service life combination was set equal to $(ae_{std})^2 \times u(0.75, 1.25) \times (P_v/A, m, N)$, where $(P_v/A, m, N)$ is the present value variance factor for N periods of service and discount rate m, calculated as shown in (4.2). Note that $(P_v/A, m, N)$ is not just $(P/A, m, N)^2$. Ignoring the $u(0.75, 1.25)$

factor, use of $(P/A, m, N)^2$ for the variance will cause a bias towards shorter service lives. The bias occurs because such a formula results in a variance higher than it should for longer service lives (e.g., for the same asset type, replacing every year will have a lower variance than replacing every three years). A bias will also occur if the technological improvement factor, $F(J,T)$, for each asset type J is not part of the formulas. The correct expressions are:

$$(P/A, m, N) = \sum_{T=1}^{N} \frac{F(J, T-1)}{(1+m)^T} \tag{4.1}$$

$$(P_v/A, m, N) = \sum_{T=1}^{N} \frac{F(J, T-1)^2}{(1+m)^{2T}} \tag{4.2}$$

The NPVs and rate of technological improvement for the first asset type were set equal to the base set. The NPVs for the other asset types were then computed by randomly perturbing the base NPVs. The monetary difference between the assets was assumed to be distributed between 0.00 and 0.05 (low case) or 0.10 and 0.30 (high case), with the case (low/high) randomly selected except for the two experiments. For each asset type and service life pair, two realizations of the monetary difference (one for for the mean and one for the variance) were drawn. The sign for each (positive or negative) was also randomly selected. The NPVs were then computed by multiplying the base NPVs by one plus these realizations. Given the NPVs, the final step was to set the rate of technological improvement for the other asset types. For each asset type, a realization of the monetary difference was drawn. The rate of technological improvement was then set by multiplying the rate for the base asset by one plus this realization.

Several extensions could be made in generating the NPVs. First, the regular method could be extended to allow for more than three cash flow components and the rate of technological improvement could be specified by cash flow component.

For example, first costs could rise by six percent per period. Annual operating costs could be broken into machine maintenance costs and raw material costs, with machine maintenance costs rising nine percent per period and raw material costs rising by two percent per period. A second extension would be to include changeover effects. Changeover effects would arise if the cost of installing and operating an asset was dependent on the immediately preceding asset type in the sequence. To model changeover effects, the DM could specify the cash flows dependent on the immediately preceding asset type.

4.2 The Experiments

Five experiments were performed. The numbering of the experiments corresponds to the order in which they were performed. The objectives of the experiments were:

- To gain insight about replacement decision making under uncertainty.
- To compare the performance of the alternative decision procedures.

The first two experiments evaluated the ability of the branch and prune procedure and cluster heuristic to solve a wide range of problems when assets were assumed to be independent. The results indicated that these two procedures could find either optimal or nearly optimal sequences. Thus, the last three experiments focused on identifying some of the factors that influence the performance of the alternative decision procedures, including the impact of correlated assets. The subsections below describe the purpose, conclusions, design, and results including a number of observations for each of the five experiments.

4.2.1 Experiment One

The purpose of the first experiment (EX1) was to determine if the branch and prune procedure could solve optimally a wide range of problems when assets were assumed to be independent. The main conclusions of EX1 were:

- The branch and prune procedure solved optimally many, but not all, of the independent asset problems. Problems with long horizons, a large number of asset types, and long service lives occasionally caused the procedure to exceed the available computer memory.
- The MV-efficient nodes were found to cluster.

EX1 Experimental Design and Results

EX1 was a 2 x 2 x 2 x 2 factorial design with five replications. The performance measure used was the maximum number of MV-efficient nodes that occurred in any stage. The number of nodes was selected as the performance measure because EX1 focused on finding optimal sequences and not on comparing the alternative decision procedures. The four factors were the number of asset types, maximum asset service life, asset monetary difference, and horizon. The distributions used for the low and high cases for each factor are shown in Table 4.3, where DU is used to denote a discrete uniform distribution and U a continuous uniform distribution. Uncertainty was limited to one level, with $uncert_a = 0.1$ and $uncert_b = 0.5$. Based on a sample of ten problems (244 mean and variance pairs), the median NPV coefficient of variation was 0.4076, with an average of 0.5906, minimum of 0.0950, maximum of 10.0002, and variance of 0.9140.

EX1 was run using discount rates of 15% and 30%. The data collected from the

Table 4.3: EX1 Parameter Ranges

Parameter	Low	High
Assets	DU(2,3)	DU(4,5)
Service life	DU(2,5)	DU(6,9)
Difference	U(0.0,0.05)	U(0.1,0.3)
Horizon	DU(10,20)	DU(20,30)

five replications are shown in Table 4.4. The header for each of the sixteen cells indicates the level of the four factors in the order listed above, with L being the low case and H the high case. The m=15% results are shown first, with the m=30% results following. A review of Table 4.4 shows a wide range in the required maximum number of MV-efficient nodes, even within the same cell.

An analysis of variance for the m=15% case is shown in Table 4.5. For the "Source of Variation" column, the factors are coded as follows: A is the number of assets, S the service life, D the asset difference, and H the horizon. To be significant, F_0 must be at least $F(\alpha = 0.10, \nu_1 = 1, \nu_2 = 64) \approx 2.79$. The results show the service life and horizon factors are significant.

EX1 Observations

EX1 yielded several insights. First, despite the fact that the asset monetary difference (factor D in Table 4.5) was not significant, the hardest problems to solve optimally were those for which several assets were almost identical. This result is neither surprising nor of great concern because if two assets are very similar economically then it matters less which is selected. While the NPVs can differ slightly, the key was that future versions of the assets change almost identically (e.g., a technological rate of increase of 0.0475 for asset type one and 0.0480 for asset type two).

Second, the analysis of variance shown in Table 4.5 indicated that both the hori-

Table 4.4: Maximum Number of MV-Efficient Nodes Required

$m = 15\%$

LLLL	LHLL	HLLL	HHLL	LLLH	LHLH	HLLH	HHLH
68	23	16	3	33	51	3	13
4	6	9	12	2	90	110	68
3	14	6	6	1976	3	1277	151
51	113	11	5	4	39	744	9
22	3	9	9	2	130	15	82
LLHL	**LHHL**	**HLHL**	**HHHL**	**LLHH**	**LHHH**	**HLHH**	**HHHH**
5	7	7	84	964	9	81	9
7	34	11	17	57	110	99	20
6	49	21	10	19	65	144	140
9	13	224	9	2	11	13	21
3	311	1202	47	34	24	867	6

$m = 30\%$

LLLL	LHLL	HLLL	HHLL	LLLH	LHLH	HLLH	HHLH
55	17	20	3	33	26	2	11
2	8	16	10	208	24	3	7
2	9	5	5	++	3	++	287
50	81	11	7	5	27	++	9
27	4	10	3	2	179	15	118
LLHL	**LHHL**	**HLHL**	**HHHL**	**LLHH**	**LHHH**	**HLHH**	**HHHH**
9	7	8	76	1208	10	521	11
7	44	11	22	26	173	247	50
4	95	36	18	19	323	884	89
6	12	676	12	2	14	15	49
3	346	++	23	136	27	++	14

Note: ++ indicates the m=30% case required more nodes than the corresponding m=15% case.

Table 4.5: EX1 Analysis of Variance for m=15%

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
A	18422.5	1	18422.5	0.17
S	498332.5	1	498332.5	4.72*
D	2247.2	1	2247.2	0.02
H	316009.8	1	316009.8	2.99*
AS	49104.1	1	49104.1	0.46
AD	23529.8	1	23529.8	0.22
AH	6480.0	1	6480.0	0.06
SD	7144.2	1	7144.2	0.07
SH	250432.2	1	250432.2	2.37
DH	179551.3	1	179551.3	1.70
ASD	35448.2	1	35448.2	0.34
ASH	25776.2	1	25776.2	0.24
ADH	26136.5	1	26136.5	0.25
SDH	82818.5	1	82818.5	0.78
ASDH	33048.5	1	33048.5	0.31
Error	6758670	64	105604.2	
TOTAL	8313151	79		

zon and service life were significant. Longer horizons made optimal sequences more difficult to obtain. Results for horizons longer then 30 periods showed that the number of MV-efficient nodes required could exceed the allocated computer memory. Shorter service lives also made optimal sequences more difficult to obtain. One explanation is that longer service lives will in general have greater annual equivalent value and, thus, tend to dominate shorter service lives. For example, buying a car and then selling it after just one or two years will typically result in higher annual costs than keeping it for six or seven years. Without any longer, dominating service lives, the branch and prune procedure cannot eliminate all the shorter combinations. While shorter service lives created more nodes on the MV-efficient list, it was the longer service lives that greatly increased the required CPU times. The correlation between the maximum MV list size and CPU times, while positive, is not one.

Third, higher discount rates resulted in problems that were on average more difficult to solve. Comparing the m=15% and m=30% cases in Table 4.4 shows that 42 problems were more difficult with the higher discount rate, 13 were the same, and 25 were easier. Table 4.4 also shows that the most difficult problems to solve were even more difficult with the higher discount rate (e.g., cells HLLH or HLHH). A higher discount rate tends to significantly diminish any differences among the assets. With no assets dominating, the nodes expand more rapidly. Exceeding the allocated computer memory was only a problem for EX1 beyond about 25 periods and at that point the additional value of a given asset was small. From a practical point of view, while it probably makes little difference which assets are selected at such later time periods, it does make finding the EU sequence more difficult.

Finally, the MV-efficient nodes were found to cluster, with the breaks in the variance. Plots of the MV-efficient list for two problems from EX1 are shown in Figure 4.1 and Figure 4.2. The clustering pattern was noted for both problems with medium sized (e.g., 250) MV-efficient lists as well as for large sized (e.g., 2000) lists. Additionally, clustering was more pronounced for higher discount rates. This can be seen by comparing Figure 4.1 and Figure 4.2, in which the first problem used m = 15%, while the second m = 30%. The fact that the MV-efficient nodes were found to cluster was used to develop the cluster heuristic to find "good" EU sequences. The cluster heuristic was invoked if the branch and prune's MV-efficient list exceeded an upper limit. Thus, the EU decision procedure consisted of a combination of the branch and prune procedure and the cluster heuristic. Clustering was also used to develop the cluster heuristic's associated upper bounding procedure, by creating a set of nodes for each stage using the maximum mean and minimum variance within each cluster.

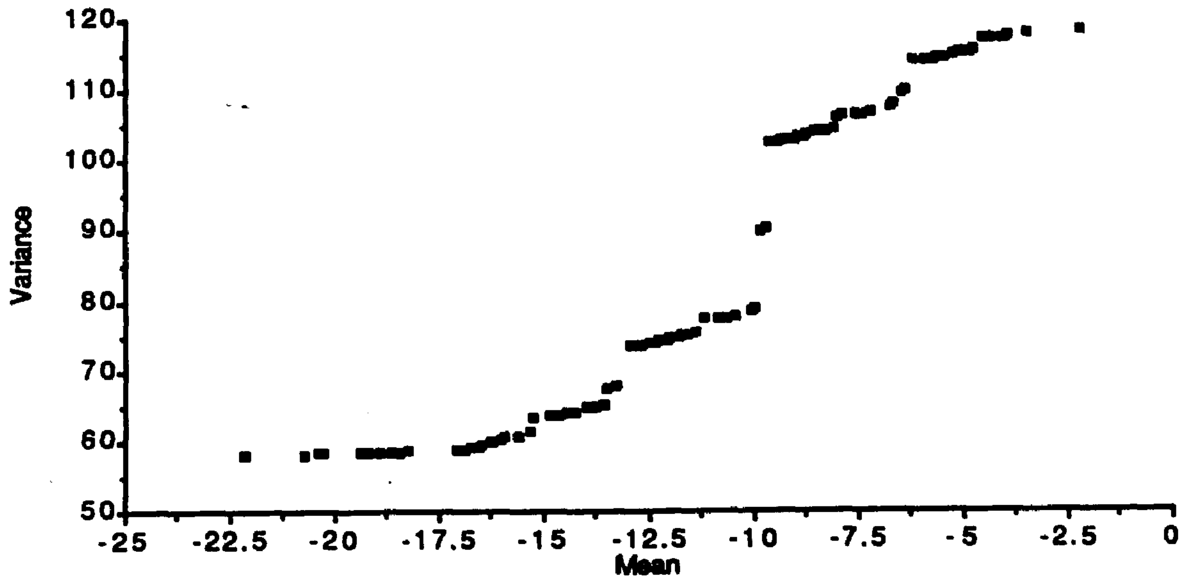

Figure 4.1: MV-Efficient Nodes for a Sample EX1 Problem, m = 15%

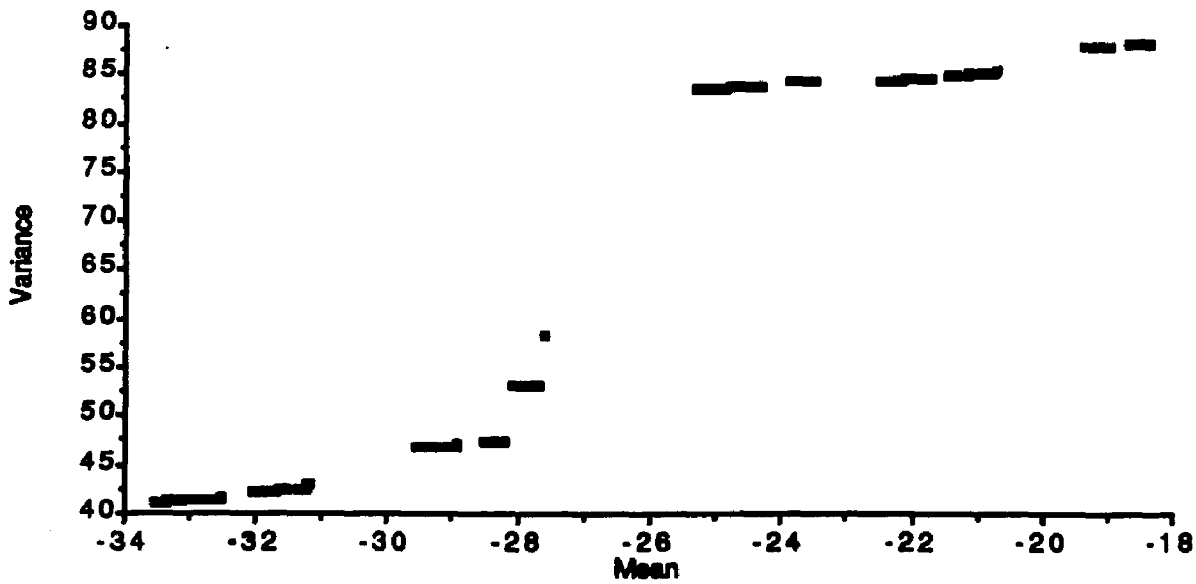

Figure 4.2: MV-Efficient Nodes for a Sample EX1 Problem, m = 30%

4.2.2 The Cluster Heuristic and Upper Bounding Procedures

The cluster heuristic's objective was to partition nodes into clusters and identify one representative node from each cluster. Since the nodes within a cluster have fairly similar means and variances, the hope was that such an approach could find "good" sequences in a much more computationally efficient manner than the branch and prune. However, before such a heuristic could be implemented, a method was needed to identify the clusters. Two methods were considered. The first method was to calculate the EU for each node and then to determine when a "large" change in the EU occurred for two adjacent nodes. A large change would signal movement to a different cluster. This method was not used because all the EU values would have had to be calculated using numerical integration–a rather CPU intensive operation. The second method, which was the one used, exploited the fact that breaks in the clusters occur in the variance. Intuitively, for two adjacent nodes within the same cluster with with means μ_1 and μ_2 and variances σ_1^2 and σ_2^2, $\gamma = (\mu_1 - \mu_2) \ / \ (\sigma_1 - \sigma_2)$ will tend to be "large" because the denominator will be close to zero. However, γ for two adjacent nodes in different clusters will be "relatively smaller," since the denominator will be much larger while the numerator stays about the same. By noting when γ becomes relatively smaller, the first node in the next cluster can be identified. The breakpoint value for γ was set experimentally as described below.

The cluster heuristic and branch and prune procedure are identical with one exception. Once the MV-efficient list expands beyond an upper limit, nodes are eliminated. Since the list is built for a given stage before any nodes are eliminated, the maximum size can be considerably larger than the upper limit. The rationale for waiting until the list is built is to eliminate the worst nodes overall, not just the worst nodes thus far. Nodes are selected for elimination using a measure based on first order stochastic dominance for truncated normal distributions [31]. The heuristic may eliminate all but one node, or only eliminate enough nodes to fall under the upper limit. For two truncated normal distributions with cumulative distribution functions $F_1(x)$ and $F_2(x)$, if $\mu_1 > \mu_2$ and $\sigma_1 > \sigma_2$ then $F_1(x)$ dominates $F_2(x)$ by first order stochastic dominance if and only if $\gamma = (\mu_1 - \mu_2) \ / \ (\sigma_1 - \sigma_2) > \delta$, where δ denotes the number of standard deviations from the mean that the distributions are truncated (assumed to be symmetric, although rules also exist for nonsymmetric truncation). Thus, the cluster heuristic keeps only those nodes that could not be eliminated if the normal distributions were truncated δ standard deviations from the mean.

Once invoked, the cluster heuristic starts with the MV-efficient node of highest mean and variance and ends with the MV-efficient node of lowest mean and variance. The heuristic compares sequentially the NPVs for each node on the MV-efficient list (which is kept in descending order of mean and variance) with the next node on the list. The comparison involves computing a value for γ. If γ is greater than the cutoff value, δ, the second node is eliminated and the first and third nodes are then compared. If the second node is not eliminated, the second and third nodes are then compared. Thus, when $\gamma < \delta$, a new cluster has been found. Note that the single representative node selected by the heuristic will always be the node of highest mean and variance within each cluster. Once the list is processed, if the number of nodes is still greater than the upper limit, δ is cut in half and the process repeated until the resulting list is smaller than the limit. Different δ values were tried using several MV-efficient lists from EX1. Most lists were reduced on a single pass to a size smaller than the upper limit when a value of 10 was used, and hence the initial value of δ was 10 standard deviations.

The EU decision procedure is not guaranteed to find an optimal sequence when either the cluster heuristic is used or assets are assumed to be correlated (unless all the correlation coefficients are equal to the same positive value). To evaluate the performance of the EU decision procedure in these cases, a bounding approach was taken. Finding optimal sequences using total enumeration was considered computationally unreasonable based on the number of numerical integrations that would be required for a typical problem. Two different upper bounding procedures were developed, one for independent assets and the other for correlated assets. While the bounding procedure for the correlated case is also applicable to independent case, independence allows a tighter bound to be developed.

The independent asset upper bounding procedure is identical to the cluster heuristic with two exceptions. First, when nodes are eliminated, the NPV variance of the remaining node is set equal to the variance of the node eliminated (which cannot be greater) to create a pseudo node. This pseudo node is better than either of the two nodes just considered because it MV-dominates both. At the last stage, the node of highest EU is found and that EU value is the upper bound. Second, the cutoff value δ used is larger. The rationale for a higher cutoff value is to get a tighter bound. At the horizon, if the sequence of highest EU is not bounded at any stage, then the resulting sequence is the optimal sequence. The more nodes eliminated, the less likely this is to occur.

The cluster heuristic and independent asset upper bounding procedure are not run separately from the branch and prune procedure. Rather, the branch and prune procedure evolves into the cluster heuristic and upper bounding procedures when the number of nodes on the MV efficient list exceeds the upper limit. From that point on in time, two MV-efficient lists are kept, one for each procedure. While lists of up to about 2000 nodes were found to be possible before exceeding the available computer memory, an upper limit of 300 nodes was used for the computational experiments. This limit was selected because:

- Results of EX1 indicated MV-efficient lists of size 300 often clustered.
- CPU times increased significantly for larger cutoff values.
- Larger limits resulted in lists for the next stage that exceeded the available computer memory, since nodes were not eliminated until the complete list for a given stage was built (for an upper limit of 300, the largest list found was about 1900 nodes).

The initial cutoff value for the independent asset upper bounding procedure was set at 50 standard deviations.

The correlated asset upper bounding procedure uses the maximum mean and a lower bound on the minimum variance to calculate a bound. If all the correlation coefficients are positive, an upper bound can also be found by solving the problem assuming the assets are independent, and then that bound with the smaller EU used. The maximum mean is given by the EV sequence mean. A lower bound on the minimum variance is found by solving a shortest path problem using deterministic dynamic programming. The dynamic program proceeds in a forward direction, with the stages being the time periods, and the states being the last asset type in the sequence and the last asset's minimum and maximum standard deviation for any service life. The optimal value function, $SeqVar^*(T)$, is a lower bound on the minimum variance for any sequence providing service for the first T periods. The key is to find a tight lower bound on the minimum covariance. To do this, at each stage for each asset type J, a lower bound on the minimum variance is found for any sequence in which the last asset is of type J. Let $SeqVar^*(J,T)$ be this lower bound, $StdMin(J,T)$ be the minimum possible and $StdMax(J,T)$ be the maximum possible standard deviation for asset type J in any sequence providing T periods of service for which J is the last asset. Then:

$$StdMin(J,T) = Min\{\sigma(J,T',T-T') : \ T-T' = 1,2,\ldots,\bar{N}_J, \ T' \geq 0\}$$

$$StdMax(J,T) = Max\{\sigma(J,T',T-T') : \ T-T' = 1,2,\ldots,\bar{N}_J, \ T' \geq 0\}$$

and the functional equation is

$$\begin{aligned} SeqVar^*(J,T) = Min\{ & SeqVar^*(J',T') + \sigma^2(J,T',T-T') + \\ & 2\rho_{JJ'}StdMin(J',T')\sigma(J,T',T-T'), \\ & SeqVar^*(J',T') + \sigma^2(J,T',T-T') + \\ & 2\rho_{JJ'}StdMax(J',T')\sigma(J,T',T-T') : \\ & T-T' = 1,2,\ldots,\bar{N}_J, \ T' \geq 0, \ J' = 1,2,\ldots,\bar{J}\} \end{aligned}$$

At each stage, $SeqVar^*(T) = Min\{SeqVar^*(J,T) : \ J = 1,2,\ldots,\bar{J}\}$. The boundary conditions are $SeqVar^*(J, T=0) = 0$ for all J.

4.2.3 Experiment Two

The purpose of the second experiment (EX2) was to evaluate the performance of the cluster heuristic and the independent asset upper bounding procedure. The main conclusions of EX2 were:

- The cluster heuristic and independent asset upper bounding procedure worked well over a wide range of problems in which assets were assumed to be independent. The average EU to upper bound utility performance was at least 0.9492.

Table 4.6: EX2 Parameter Distributions

Parameter	Distribution
Discount Rate	U(0.15,0.30)
Horizon	DU(20,50)
Asset Types	DU(2,7)
Asset Service Life	DU(2,14)
Asset Monetary Difference	U(-0.1,0.1)

- The utility performance measure slightly overestimated the EV performance, since 9 out of 60 times the best of 100 random sequences outperformed the EV sequence, but the bias was not considered serious.

EX2 Experimental Design and Results

The test problem generator used for EX2 was the same used for EX1, including the uncertainty endpoints, except the low/high input parameters were eliminated. Instead, these parameters were randomly drawn from the distributions shown in Table 4.6. The low/high parameters were eliminated because the purpose of EX2 was to study the performance across a wide range of problems and not to study individual parameter effects. To increase the likelihood that the the cluster heuristic would be invoked, problems with short horizons and large asset monetary differences were not considered.

Two sets of 30 problems in which the assets were independent were run for EX2. The first set had low uncertainty while the second set had high uncertainty (the problems were generated by multiplying the regular NPV variances for the asset types by 100). For each set, a number of different statistics were compiled. The results are displayed in Tables 4.7 and 4.8. Each table is broken into three sections. The first section displays how many problems required the cluster heuristic, how

Table 4.7: EX2 Problem Set 1 - Uncorrelated Assets

For the EU sequences:
10/30 required the cluster heuristic and upper bounding procedure.
9/10 cluster heuristic sequences were optimal.
0.9999 was the average EU to upper bound utility performance for the 10 cases.
For the EV sequences:
26/30 EV sequences matched the EU sequence.
1/30 random sequences outperformed the EV sequence.
0.9369 was the average utility performance for the 30 EV sequences.
For the CME sequences:
30/30 CME sequences matched the EU sequence.
0/30 random sequences outperformed the CME sequence.
1.0000 was the average utility performance for the 30 CME sequences.

Table 4.8: EX2 Problem Set 2 - Uncorrelated Assets, Higher Uncertainty

For the EU sequences:
7/30 required the cluster heuristic and upper bounding procedure.
3/7 cluster heuristic sequences were optimal.
0.9492 was the average EU to upper bound utility performance for the 7 cases.
For the EV sequences:
16/30 EV sequences matched the EU sequence.
8/30 random sequences outperformed the EV sequence.
0.6475 was the average utility performance for the 30 EV sequences.
For the CME sequences:
30/30 CME sequences matched the EU sequence.
0/30 random sequences outperformed the CME sequence.
1.0000 was the average utility performance for the 30 CME sequences.

many of the cluster heuristic sequences were optimal, and what the average EU to upper bound utility performance (as given by (3.3)) was for the cluster heuristic. For EX2, the average EU to upper bound utility performance was computed using only those problems in which the heuristic was required, as opposed to basing it on the average of all 30 problems in the set (using a value of 1.0 for all problems that did not require the heuristic). The second section shows how many EV sequences matched the EU sequence, how many times the best of 100 random sequences outperformed the EV sequence, and the average utility performance (as given by (3.2)) of the EV sequences. The number of times the EV sequence was outperformed by the best of 100 random sequences was monitored to check if the utility performance measure would be biased. Since the utility performance measure value is set to zero if the best of 100 random sequences has equal or higher EU, if this were to occur frequently, then the performance of the EV decision procedure would tend to overestimated. Finally, the third section displays the same three statistics for the CME sequences. The results indicate:

- The cluster heuristic's average EU to upper bound utility performance was at least 0.9492.

- For the 60 problems, 17 required the cluster heuristic. Of those 17, 12 heuristic sequences were known to be optimal. For the other 5 sequences, the average EU to upper bound utility performance was 0.9287, indicating the upper bound is generally quite tight.

- The EV decision procedure was outperformed by the best of 100 random 9 times out of 120, while the CME decision procedure was never outperformed. When the EV decision procedure was outperformed, the utility performance was set to 0.0–the same value as if the random and EV sequences identified had equal EU. Thus, on average, the utility performance was slightly higher for the EV decision procedure than if negative performance values had been allowed.

- The EV decision procedure's utility performance ranged from 0.6475 to 0.9369, with the worst performance for the high uncertainty problem set.

- The CME decision procedure's utility performance was 1.0, suggesting the CME to be an excellent decision procedure.

The two sets of problems required about 12 real-time hours and 7 CPU hours to run on an Apollo DN3000 workstation.

EX2 Observations

Three insights were gained from EX2. First, a particular asset type and service life pair frequently repeated in the EU sequence. In extreme cases, the same asset type with a one period service life repeated out to the horizon. Typically however, the first few asset type and service life pairs would vary, followed by a repeating pair, and ending with a different repeating pair. The repeating pairs have been observed in other replacement research done with similar models [15, 35].

Second, the EV decision procedure's utility performance was better for the low uncertainty case as compared to the high uncertainty case. This was expected for a risk averse DM and illustrates the idea that the amount of uncertainty much reach a certain level before the EV and EU sequences fail to match.

Third, the cluster heuristic ran more frequently than the independent asset upper bounding procedure. For example, while the bounding procedure only needed to run for every third stage, the heuristic had to be run for every stage. By being able to create a node having the maximum mean and minimum variance within each cluster, the bounding procedure was able to MV-dominate many nodes the heuristic was forced to continue to consider.

4.2.4 Experiment Three

The results of EX1 and EX2 indicated that the combination of the branch and prune procedure and cluster heuristic could find either optimal or nearly optimal sequences for a wide range of problems when assets were assumed to be independent. The third experiment (EX3) narrowed the focus and examined the impact of several model parameters on the performance of the alternative decision procedures. Specifically, the purpose of EX3 was to study the effects of risk aversion, correlation, uncertainty, and the horizon on the utility performance of the TRAD, EV, and CME decision procedures, and to study the EU decision procedure's EU to upper bound utility performance when assets were assumed to be correlated. The main conclusions of EX3 were:

- Compared to the EU decision procedure, the CME decision procedure was, on the average, the best since it performed at better than 99%, while the EV decision procedure followed at at 86%, and the TRAD decision procedure, at 64%, was the worst.

- The TRAD decision procedure performed poorly regardless of the level of the four factors.

- The EV decision procedure performed better for negative correlation and for short horizons combined with low risk aversion.

- The CME decision procedure was a good heuristic for constant risk aversion because the individual asset CMEs sum to nearly the overall sequence CME.

- The EU decision procedure's EU to upper bound utility performance was, on average, 76% for positive correlation and 60% for negative correlation, with uncertainty having a significant effect on the performance.

EX3 Experimental Design and Results

Due to the insights gained in the first two experiments, only four parameters were selected for further study in EX3. The parameters selected were those considered to have a significant effect on the performance of the decision procedures. They included the level of risk aversion, type of correlation, amount of uncertainty, and length of the horizon. Other parameters that could have been selected, such as the discount rate, number of asset types, maximum asset service life, or rate of technological improvement, were considered to influence more the difficulty (either in terms of the MV-efficient list size or CPU processing times) of finding the sequence of highest EU, than the utility performance of the decision procedures.

EX3 was a 2 x 3 x 2 x 2 factorial design with five replications, for a total of 120 problems. Risk aversion, uncertainty, and the horizon were studied at two levels (low/high). Correlation was studied at three levels (negative, positive, and zero or independent). For the low level of risk aversion, z was uniformly distributed between 1.5 and 5.0, while for the high level, between 16.5 and 20.0. Uncertainty for the low level used $uncert_a = 0.1$ and $uncert_b = 0.4$. Based on a sample of ten problems (244 mean and variance pairs), the median NPV coefficient of variation was 0.3369, with an average of 0.4884, minimum of 0.0820, maximum of 8.0297, and variance of 0.5876. Uncertainty for the high level used $uncert_a = 0.6$ and $uncert_b = 1.2$. Based on a

Table 4.9: Analysis of Variance for EX3 TRAD

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	0.004420	1	0.0044	0.021
C	0.155744	2	0.0779	0.368
U	0.076403	1	0.0764	0.362
H	0.367339	1	0.3673	1.738
RC	0.208925	2	0.1045	0.494
RU	0.156136	1	0.1561	0.739
RH	0.072749	1	0.0727	0.344
CU	0.065151	2	0.0326	0.154
CH	0.354630	2	0.1773	0.839
UH	0.260782	1	0.2608	1.234
RCU	0.566443	2	0.2832	1.340
RCH	0.754542	2	0.3773	1.785
RUH	0.026470	1	0.0265	0.125
CUH	0.071150	2	0.0356	0.168
RCUH	0.849309	2	0.4247	2.009
Error	20.28776	96	0.2113	
TOTAL	24.2780	119		

sample of ten problems (330 mean and variance pairs), the median NPV coefficient of variation ratio was 0.7521, with an average of 2.4423, minimum of 0.2674, maximum of 29.5582, and variance of 21.3342. Finally, the horizon was uniformly distributed between 10 and 30 periods for the low case, and between 30 and 50 periods for the high case.

The sequence generator for EX1 and EX2 used the random, EV, CME, and EU decision procedures. The TRAD decision procedure was added in EX3. The rationale for adding the TRAD decision procedure was twofold. First, the TRAD decision procedure is probably the most popular method suggested for solving replacement problems. Second, in EX2 it was observed that certain asset type and service life pairs often repeat, suggesting that the TRAD decision procedure may perform well.

The overall average utility performance for the 120 problems was 0.6396 for the

Table 4.10: Analysis of Variance for EX3 EV

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	0.159741	1	0.1597	1.572
C	0.161145	2	0.0806	0.793
U	0.392006	1	0.3920	3.859*
H	0.008514	1	0.0085	0.084
RC	0.077346	2	0.0387	0.381
RU	0.020780	1	0.0208	0.205
RH	0.314450	1	0.3145	3.095*
CU	0.024397	2	0.0122	0.120
CH	0.138219	2	0.0691	0.680
UH	0.007466	1	0.0075	0.074
RCU	0.012331	2	0.0062	0.061
RCH	0.086060	2	0.0430	0.424
RUH	0.000008	1	.0000	.000
CUH	0.454670	2	0.2273	2.238
RCUH	0.374584	2	0.1873	1.844
Error	9.752125	96	0.1016	
TOTAL	11.9839	119		

TRAD decision procedure, 0.8581 for the EV decision procedure, and 0.9937 for the CME decision procedure. The EU decision procedure's average EU to upper bound utility performance was equal to 0.9855 for the independent case, 0.7560 for the positive correlation case, and 0.6026 for the negative correlation case. The problems required about 36 real-time hours and 24 CPU hours to run on an Apollo DN3000 workstation. An analysis of variance was performed for the TRAD, EV, and CME decision procedures. The results are shown in Tables 4.9 through 4.11. For the column "Source of Variation," the factors are coded as follows: R is the level of risk aversion, C the correlation, U the uncertainty, and H the horizon. F_0 must be at least $F(\alpha = 0.10, \nu_1 = 1, \nu_2 = 96) = 2.77$ or $F(0.10, 2, 96) = 2.37$ to be significant. Significant factors are followed by an asterisk in the F_0 column of the tables. The results indicate:

Table 4.11: Analysis of Variance for EX3 CME

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	0.000376	1	0.0004	0.353
C	0.005581	2	0.0028	2.611*
U	0.003008	1	0.0030	2.814*
H	0.003048	1	0.0030	2.852*
RC	0.001632	2	0.0008	0.764
RU	0.001028	1	0.0010	0.962
RH	0.000039	1	.0000	0.037
CU	0.007006	2	0.0035	3.277*
CH	0.003328	2	0.0017	1.557
UH	0.001805	1	0.0018	1.689
RCU	0.001021	2	0.0005	0.478
RCH	0.000542	2	0.0003	0.254
RUH	0.000431	1	0.0004	0.403
CUH	0.004275	2	0.0021	2.000
RCUH	0.000243	2	0.0001	0.114
Error	0.102617	96	0.0011	
TOTAL	0.1360	119		

- The TRAD decision procedure performed poorly under all conditions since there are no significant factors in Table 4.9

- U and RH are significant in Table 4.10 for the EV decision procedure. The EV decision procedure performed better, on average, for low U (0.9152) than for high U (0.8009). It also performed better, on average, for for low R and low H (0.9373) than for high R and low H (0.7620). Results were about the same, on average, for low R, high H and high R, high H (0.8518 and 0.8812). These last results suggest the the EV decision procedure is less sensitive to the level of risk aversion as the horizon increases.

- C, U, H, and CU are significant in Table 4.11 for the CME decision procedure, although because the CME performed so well there is very little total error.

Table 4.12: Average Upper Bound Utility Performance

R	U	H	Positive Corr. Avg. Upper Bound Utility Performance	Negative Corr. Avg. Upper Bound Utility Performance
l	l	l	0.9406	0.7533
l	h	l	0.6164	0.3121
h	l	l	0.8485	0.7601
h	h	l	0.4872	0.3174
l	l	h	0.7956	0.8733
l	h	h	0.8785	0.6121
h	l	h	0.8955	0.7162
h	h	h	0.5859	0.4761
Overall			0.7560	0.6026

The EU decision procedure's EU to upper bound utility performance was, on average, 0.7560 for positive correlation and 0.6026 for negative correlation. Table 4.12 shows the average values for each cell, with "l" being the low case and "h" the high case. The performance for some cells is very good, but for others quite poor. An analysis of variance for positive correlation is shown in Table 4.13 and indicates uncertainty to be significant. The performance was better for low uncertainty. Table 4.14 shows the analysis of variance for negative correlation. Again uncertainty is significant, as well as the horizon. Thus, these results suggest that as uncertainty increases the EU decision procedure's EU to upper bound utility performance decreases. However, these results may mask the effect of uncertainty on the tightness of the bound compared to the optimal solution. As discussed below, higher uncertainty may result in the variance used for the bound being considerably smaller than the variance of the optimal solution.

Table 4.13: Analysis of Variance for Positive Correlation

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	0.10709	1	0.10709	1.73
U	0.52019	1	0.52019	8.42*
H	0.04315	1	0.04315	0.70
RU	0.11538	1	0.11538	1.87
RH	0.00051	1	0.00051	0.01
UH	0.13164	1	0.13164	2.13
RUH	0.07895	1	0.07895	1.28
Error	1.97708	32	0.06178	
TOTAL	2.97397	39		

EX3 Observations

Three insights were gained from EX3. First, despite the frequency of sequences with assets that repeat, the TRAD decision procedure performed poorly. This may have been due to the fact that assets generally repeated later in the sequence when they had less impact.

Second, as in EX2, the CME decision procedure performed amazingly well. The most likely reason is that the CMEs are "almost" additive for the exponential utility function with constant risk aversion. With constant risk aversion, the uncertainty penalty is independent of the mean. For example, using c = 0.008483 for a mean of 100 and variance of 100,000, the CME = -324.193. For the same c and variance, but with a mean of -100, the CME = -524.193. For both cases, the CME is just the mean minus 424.193. With the individual asset NPVs in a sequence assumed to be normally distributed, the variances (plus any covariance terms) sum to the variance of the sequence. As long as the individual asset variances are of the same order of magnitude as the total variance for the sequence, the sum of the penalties for individual assets will be close to the total penalty for the sequence. For EX3,

Table 4.14: Analysis of Variance for Negative Correlation

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	0.04932	1	0.04932	0.89
U	1.19904	1	1.19904	21.73*
H	0.17881	1	0.17881	3.24*
RU	0.00024	1	0.00024	.00
RH	0.05817	1	0.05817	1.05
UH	0.09144	1	0.09144	1.66
RUH	0.00032	1	0.00032	0.01
Error	1.76542	32	0.05517	
TOTAL	3.34276	39		

since the number of assets in a sequence could never be more than 50 (the maximum horizon) and was typically much less, the variance for each of the individual assets was similar in order of magnitude to the sequence variance, especially when there was no correlation. Continuing with $c = 0.008483$, suppose a sequence was made up of 10 assets, had a mean equal to $MEAN$, and variance equal to 100.0. If the assets in the sequence were independent, and each had a mean of ($0.1 \times MEAN$) and variance of 10.0, then the CME for each asset is ($0.1 \times MEAN$) - 0.04241939. For the entire sequence, the CME = $MEAN$ - 0.42419379. Adding the CMEs for the 10 individual assets results in a value of $MEAN$ - 0.4241939, or a difference of just 0.00000011.

To verify the CMEs are almost additive for the exponential utility function with constant risk aversion, an experiment was run using 100 randomly generated problems. The first 50 were generated using the regular NPV method, while the last 50 were generated using the AE NPV method. All the assets were assumed to be independent. For each problem, 100 random sequences were generated, and for each sequence the following two performance measures were computed, with n being the number of assets in the sequence:

$$\text{CME difference} = \frac{|\sum_{j=1}^{n} \text{CME}(asset_j) - \text{CME(sequence)}|}{|\text{CME(sequence)}|}$$

$$\text{EU difference} = \frac{|\text{U}(\sum_{j=1}^{n} \text{CME}(asset_j)) - \text{EU(sequence)}|}{|\text{EU(sequence)}|}$$

The average CME difference for the 50 regular NPV problems (based on a total of 5000 sequences) was 0.000141%, while for the 50 AE NPV problems the value was 0.000362%. The average EU difference values were 0.000005% for the 50 regular NPV problems and 0.000007% for the 50 AE NPV problems.

For each problem, the best random sequence out of 100 based on the EU for the sequence, was compared to the best random sequence out of 100 based on the sum of the individual asset CMEs. For the 50 regular NPV problems, the sequences matched all 50 times. For the 50 AE NPV problems, 49 out of 50 times the sequences matched. To verify that it was not trivial to select the best out of 100 sequences, the average difference between the best sequence and the next best (that did not match) sequence was also computed both in terms of CME and EU as follows:

$$\text{CME difference} = \frac{|\text{CME(best sequence)} - \text{CME(next best sequence)}|}{|\text{CME(best sequence)}|}$$

$$\text{EU difference} = \frac{|\text{EU(best sequence)} - \text{EU(next best sequence)}|}{|\text{EU(best sequence)}|}$$

The average CME difference was 4.38% for the 50 regular NPV problems and 0.82% for the 50 AE NPV problems. For the EU difference the values were 5.74% and 0.72%. Both indicate that selecting the best sequence was not trivial, especially for the AE NPV problems.

Overall, the results strongly support the theory that the CME decision procedure is an excellent heuristic for a DM with constant risk aversion because the CMEs for the individual assets sum to almost the sequence CME. While the CMEs are almost additive, the EUs are not. A small experiment run with 10 problems (1000 sequences total) found the average absolute EU difference between the sum of the individual assets and the entire sequence to be 38.9%.

The third observation was the EU decision procedure's EU to upper bound utility performance was probably better than the results of EX3 indicated. Generating tight bounds was difficult, especially in the negative case. For the 40 problems with positive correlation, the best upper bound for 36 was found by solving the problem assuming the assets to be independent. However, in the negative case, the bound had to be developed using the maximum mean and a lower bound on the minimum variance. For negative correlation coefficients, the lower bound for the minimum variance was driven down towards zero, underestimating the true variance by a considerable margin for a typical problem. As the results of EX5 will show, the EU sequence is fairly stable for small changes in the variance forecasts. Thus, although it is possible for the EU decision procedure to prune the optimal sequence, even if this occurs, the sequence found should still typically be "good."

4.2.5 Experiment Four

As discovered in EX2 and EX3, the CME decision procedure is an excellent heuristic for a DM with constant risk aversion. To study the impact of decreasing risk aversion, experiment four (EX4) was performed. For EX4, in addition to the exponential utility function with constant risk aversion used for EX2 and EX3, the logarithmic and power utility functions which have decreasing risk aversion were also used. The purpose of EX4 was to determine if the EU sequences were similar for the three different types of utility functions, and to study the performance of the CME decision procedure for DMs with decreasing risk aversion. The main conclusions of

EX4 were:

- The EU sequences for the exponential utility function were very similar to the EU sequences for both the logarithmic and power utility functions when the risk aversion levels were similar.

- The CME decision procedure performed poorly for utility functions with decreasing risk aversion, and thus, in general, is not a good heuristic.

EX4 Experimental Design and Results

The sequence generator for EX4 used three different types of utility functions:

- The exponential, $U(w) = (1 - e^{-cw})/c$, with constant risk aversion, $r(w) = c$.

- The logarithmic, $U(w) = \ln(w + b)$, $(w + b) > 0$, with decreasing risk aversion, $r(w) = 1/(w + b)$.

- The power, $U(w) = (w - w_0)^{\beta}$, $w_0 < w$ for all w, $0 < \beta < 1$, with with decreasing risk aversion, $r(w) = (1 - \beta)/(w - w_0)$.

To compare intelligently the impact of the different utility functions, the level of risk aversion needs to be similar for all three. Since $r(w)$ for the logarithmic and power utility functions depends on the value of w, the $r(w)$'s for the three utility functions can only match at one point. Given the method used for EX2 and EX3 to estimate the relevant range for the utility function, this single point logically should fall within that range. Three different approaches were considered for EX4. The first approach was to match the risk aversion levels at the lower endpoint. However, then the DM would be less risk averse over the entire range of outcomes and comparing such utility functions would provide less insight. The second approach, matching at the upper endpoint, had the reverse problem.

The third approach was to set the $r(w)$'s equal at the mean NPV for the EV sequence, using c as the anchor. Then for approximately half of the relevant range the level of risk aversion for the logarithmic and power utility functions would be lower than the exponential, and higher for the other half. However, with the normally distributed NPVs having infinite tails, the values for b, β, and w_0 that make the $r(w)$'s equal could fail to satisfy the restrictions on the logarithmic and power utility functions. The $r(w)$'s are equal for the exponential and logarithmic utility functions when $b = 1/c - w$, for given values of c and w. Even if the lower tail were truncated (say at the mean minus 3.5 standard deviations), there could still be cases in which the value of b would not be large enough to satisfy $(w + b) > 0$. For example, for an EV sequence with a mean NPV of 0 and variance of 816.32, truncating at 3.5 standard deviations yields a range of -100 to 100, with $c = \ln(1.5)/100 = 0.00405$ to $c = \ln(20)/100 = 0.02996$. For $c = 0.02996$, $b = 33.38$ and if $w = -100$ then $(w + b) < 0$.

The $r(w)$'s are equal for the exponential and power utility functions when $\beta = 1 - c(w - w_0)$, for given values of c, w, and w_0. The problem in this case is that β may be negative. Continuing with $c = 0.02996$ and assuming $w_0 = -101$, $\beta = 1 - 0.02996(0 + 101) = -2.02596$. In general for truncated distributions, very low risk aversion levels (e.g., $c = 0.00405$) will typically satisfy $(w + b) > 0$ and $0 < \beta < 1$, but higher levels of risk aversion can create problems.

To avoid as many of these problems as possible, two changes were made for EX4. First, the test problem generator was changed to generate only positive mean NPVs, since negative mean NPVs are more likely to create problems for the logarithmic and power utility functions. Positive mean NPVs were generated by using the AE method, and by drawing the mean AE from a strictly positive uniform distribution (with endpoints of 100 and 10,000). The regular NPV method could have been modified to only produce positive NPV means by, for example, making the first cost a large positive value. However, such modifications change the basic structure and may affect the relative rankings of the asset type and service life pairs. Second, the numerical integration procedure used for the sequence generator was changed to truncate the normally distributed NPVs at ± 3.5 standard deviations from the mean.

However, these two changes were still not sufficient to ensure the logarithmic and power utility function restrictions would always be satisfied. Thus, the numerical integration routine was modified to detect if the restrictions were violated and to correct any violations found. To do so, the values for b and β were set as follows:

- $b = \max\{1/c - \mu_{ev}, -(w_0 - 1)\}$, where $w_0 = \mu_{ev} - 3.5\sigma_{ev}$ is the lower NPV endpoint, μ_{ev} is the mean of the NPV for the EV sequence, and σ^2_{ev} is the variance of the NPV for the EV sequence. The numerical integration routine set $(w + b) = 1^{-20}$ for any $(w + b) < 1^{-20}$.

- $\beta = \max\{1 - c(\mu_{ev} - w_0), 0.001\}$, where $w_0 = \mu_{ev} - 3.5\sigma_{ev}$ is the lower NPV endpoint, μ_{ev} is the mean of the NPV for the EV sequence, and σ^2_{ev} is the variance of the NPV for the EV sequence. The numerical integration routine set $(w - w_0) = 0$ for any $(w - w_0) < 0$.

When b or β could not be set so that the level of risk aversion for the logarithmic and power utility functions matched the exponential at μ_{ev}, the resulting level of risk aversion was always lower at μ_{ev}.

EX4 used the same 2 x 3 x 2 x 2 factorial design with five replications that was used for EX3. The four factors (level of risk aversion, uncertainty, correlation, and horizon) were also the same, with the ranges the same except for uncertainty. Uncertainty had to be changed because the NPVs were generated using the AE method. The AE uncertainty endpoints were set such that the resulting median NPV coefficient of variation was approximately the same as the value used in EX3. Uncertainty for the low level used $uncert_a = 0.5$ and $uncert_b = 1.0$. Based on a sample of ten problems (323 mean and variance pairs), the median NPV coefficient of variation was 0.3266 (vs. 0.3369 for EX3), with an average of 0.3800, minimum of 0.1356, maximum of 1.1543, and variance of 0.0322. Uncertainty for the high level used $uncert_a = 1.0$ and $uncert_b = 2.8$. Based on a sample of ten problems (431 mean and variance pairs), the median NPV coefficient of variation was 0.7630 (vs. 0.7521 for EX3), with an average of 0.9361, minimum of 0.2604, maximum of 3.5340, and variance of 0.2824.

For each problem, the TRAD, EV, CME, and EU sequences were found for each of the three different utility functions. Thus, twelve potentially different sequences were found in total. Using these twelve sequences, a total of thirteen performance statistics were calculated. The first nine statistics were the utility performance values (as given by (3.2)) for the TRAD, EV, and CME sequences by utility function. Then, four special performance statistics were calculated. These special statistics were designed to compare the CME and EU sequences across utility functions, to determine how different the sequences selected were. To do so, exponential utility was used as the base and all the sequences were compared using their associated exponential utility. The first two statistics compared the exponential CME sequence to the logarithmic and power CME sequences, while the last two compared the

Table 4.15: EX4 Overall Results

Regular Utility Performance			
Utility Function	**TRAD**	**EV**	**CME**
Exponential	0.4141	0.8965	0.9859
Logarithmic	0.4298	0.9258	0.5406
Power	0.4515	0.9305	0.3001
Special Performance			
Comparison		**CME**	**EU**
Exponential-Logarithmic		0.5320	0.9802
Exponential-Power		0.2710	0.9801

exponential EU sequence to the logarithmic and power EU sequences. Specifically, letting *exp* denote the exponential utility function, the first two are of the form:

$$\frac{EU_{exp}(\text{log or power CME sequence}) - EU_{exp}(\text{best of 100 } exp \text{ random})}{EU_{exp}(exp \text{ CME sequence}) - EU_{exp}(\text{best of 100 } exp \text{ random})}$$

while the last two are of the form:

$$\frac{EU_{exp}(\text{log or power EU sequence}) - EU_{exp}(\text{best of 100 } exp \text{ random})}{EU_{exp}(exp \text{ EU sequence}) - EU_{exp}(\text{best of 100 } exp \text{ random})}$$

For both, if the best of 100 random outperformed the sequence being compared, the value was set to zero.

The sequence generator required about 21 real-time hours, and 13 Apollo DN3000 CPU hours to solve the 120 problems. The average value for each performance measure is shown in Table 4.15. Comparing the overall results for EX3 with the overall exponential results for EX4 shows:

- The TRAD decision procedure's utility performance drops from 0.6396 for EX3 to 0.4141 for EX4.

- The EV decision procedure's utility performance does slightly better for EX4 (0.8965 vs. 0.8518).

Table 4.16: Summary of EX4 Significant Factors

Sequence	Exponential Significant Factor(s)	Logarithmic Significant Factor(s)	Power Significant Factor(s)
TRAD	R,C,H,RCUH	C,H,RCUH	RCUH
EV	R,C,U,H,RU,CU,CH,UH,CUH	R,C,U,H,CU,CH,UH,CUH	C,U,H,CU,CH,UH,CUH
CME	-	R,U,RU	R,U,RU

- The CME decision procedure's utility performance drops slightly from 0.9937 for EX3 to 0.9859 for EX4.

Comparing the overall results of Table 4.15 across the three utility functions indicates the average performance was about the same except for the CME decision procedure. The CME decision procedure results diminish for the logarithmic and power utility functions. Looking at the special performance statistics in Table 4.15, for the CME decision procedure the average values are close to the corresponding average regular values indicating different sequences were selected for the three utility functions, especially for the exponential and power functions. The high EU values provide strong evidence that very similar EU sequences were being selected for the three utility functions.

An analysis of variance was performed for each of the nine regular utility performance measure statistics. The results are displayed in Tables B.1 through B.9. A summary of the significant factors is shown in Table 4.16. Comparing the significant factors for each type of sequence across the three utility functions shows:

- The TRAD decision procedure's significant factors differ only slightly. The significant factors for the logarithmic utility function are a subset of the exponential factors. The significant factors for the power utility function are a subset of the logarithmic factors.

- The EV decision procedure's significant factors also differ only slightly. The significant factors for the logarithmic utility function are a subset of the exponential factors, and the significant factors for the power utility function are a subset of the logarithmic factors.

- The CME decision procedure's significant factors differ, but the significant factors for the logarithmic utility function match those of the the power utility function.

EX4 Observations

EX4 yielded three insights. First, the EV decision procedure's utility performance was better for the logarithmic and power utility functions than for the exponential as the level of risk aversion increased. For low risk aversion, the EV decision procedure's utility performance was 0.9594 when the exponential utility function was used, but decreased to 0.9554 when the logarithmic utility function was used, and decreased to 0.9519 when the power utility function was used. However, for high risk aversion, the EV decision procedure's utility performance was 0.8335 when the exponential utility function was used, but increased to 0.8962 when the logarithmic utility function was used, and increased to 0.9091 when the power utility function was used. One possible explanation is that in order to satisfy the parameter restrictions, the logarithmic and power utility functions were forced to use a level of risk aversion lower than that used by the exponential utility function. As found in EX3, the utility performance of the EV decision procedure tends to decrease as the level of risk aversion increases.

Second, the CME decision procedure was not a good heuristic for decreasing risk aversion. The special CME performance statistics show that the selected sequences differed for the three utility functions. However, since the EU sequences for the three were very similar and the CME decision procedure performed well when the exponential utility function was used, the utility performance of the CME decision procedure for the logarithmic and power functions would probably have increased if the CME decision procedure had assumed constant risk aversion for these cases.

Third, the regular and AE methods to generate NPVs, produced problems with different characteristics even when the uncertainty levels were matched. In general, the AE method generated problems with lower relative uncertainty. When relative uncertainty was very low, the EV and EU sequences matched no matter how high the absolute uncertainty for the problems considered in this research. CPU times for the problems generated using the AE method, were only about half those for the regular method. This would indicate that the dominating asset type and service life pairs were easier to identify when the AE method was used.

4.2.6 Experiment Five

As the NPV forecasts are being developed by a DM, inconsistencies in the forecasts may arise. For example, demand forecasts made by the marketing department may differ from those made by the production department. The sequence of highest EU could depend on which forecast is used. To gain some insight into the impact of changes in the forecasts, the fifth experiment (EX5) was performed. The purpose of EX5 was to study the sensitivity of the sequence of highest EU to changes in the mean and/or variance of the NPVs. A base set of NPVs was used to find the sequence of highest EU, and this sequence was used as the basis for comparison. The mean and/or variance of each of the base NPVs was then perturbed, and the perturbed values were used to find a new sequence of highest EU. This new sequence was then compared to the "base" sequence (the sequence of highest EU found for the base set of NPVs). The main conclusions of EX5 were:

- Perturbing the mean caused the selected sequence of highest EU to differ quite dramatically from the base sequence, while perturbing the variance by the same amount had much less impact.

- Perturbing both the mean and variance provided results similar to perturbing only the mean.

- If the mean was always perturbed in the same direction, the magnitude of the difference between the selected sequence of highest EU and the base sequence remained relatively constant for all levels of perturbation.

EX5 Experimental Design and Results

EX5 used the same 2 x 3 x 2 x 2 factorial design with five replications that was used for both EX3 and EX4. The four factors (level of risk aversion, uncertainty, correlation, and horizon) were the same, and the ranges used were the same as for EX3. For EX5, only the exponential utility function was used. The same set of problems was solved four different times. In the first case, only the means were perturbed. For the second case, only the variances were perturbed. In the third case, the means and the variances were both perturbed. Finally, for the fourth case, only the means were perturbed and only in the same direction. Four different levels of perturbation were used. For the first three cases, the base values were perturbed uniformly by $\pm(0.0, 0.25)$, $\pm(0.25, 0.50)$, $\pm(0.50, 0.75)$, and $\pm(0.75, 1.0)$. Each mean and/or variance was perturbed independently. For the fourth case, the base means were perturbed uniformly by $(-0.5, -0.25)$, $(-0.25, 0)$, $(0, 0.25)$, and $(0.25, 0.5)$. Each mean was perturbed independently. Thus, for each problem a total

Table 4.17: EX5 Overall Results

Case	Perturbation Level			
	(0.00,0.25)	**(0.25,0.50)**	**(0.50,0.75)**	**(0.75,1.00)**
1. Mean	0.4023	0.2530	0.2701	0.2030
2. Variance	0.8465	0.7670	0.7301	0.7048
3. Mean and Variance	0.4140	0.2555	0.2407	0.1579
	(-0.50,-0.25)	**(-0.25,0.00)**	**(0.00,0.25)**	**(0.25,0.50)**
4. Mean, Same Direction	0.5347	0.5381	0.5293	0.5071

of 17 potentially different potentially different sequences of highest EU had to be found (the base sequence plus one sequence for for each of the four perturbed sets of NPVs for each of the four cases). To avoid using the computationally intensive EU decision procedure, the CME decision was used in EX5 to find the sequence of highest EU. The rationale was that in EX3, the CME decision procedure was found to perform well when the exponential utility function was used, and finding 17 EU sequences for each of the 120 problems would have required over 100 real-time hours based on the EX3 times.

The performance measure used for EX5 compared the CME sequence found using the perturbed forecasts with the CME sequence found using the original forecasts (the base sequence). The best of 100 random sequences found using the original forecasts was used as the benchmark. To make the comparison, the EU using the base NPVs was calculated for CME sequence found using the perturbed forecasts. The resulting performance measure was then:

$$\frac{EU(\text{CME sequence using perturbed forecasts}) - EU(\text{best of 100 random})}{EU(\text{CME base sequence}) - EU(\text{best of 100 random})}$$

If the best of 100 outperformed the CME sequence found using the perturbed forecasts, the value was set to zero.

Each of the four runs required about 11 real-time hours and 10 CPU hours on

Table 4.18: Summary of EX5 Significant Factors

Case	(0.00,0.25) Significant Factor(s)	(0.25,0.50) Significant Factor(s)	(0.50,0.75) Significant Factor(s)	(0.75,1.00) Significant Factor(s)
1. Mean	-	RCU	-	CU
2. Variance	H,RC,CU,UH,RCU,CUH	U,H	R,U,RU	R,U,H,RU
3. Mean and Variance	R,C,RH,CH,RCH	R,U,H,RCU	R,RCUH	-
	(-0.50,-0.25)	(-0.25,0.00)	(0.00,0.25)	(0.25,0.50)
4. Mean, Same Direction	RH,CH	C,RC,RH,RCH	C,RCH	RH,UH,RCH

an Apollo DN3000 workstation. The average results for the 120 problems are shown in Table 4.17. For the first case, perturbing the means by up to 25% resulted in a performance value of 0.4023, indicating the perturbed and base sequences were on average quite different. Beyond 25%, the performance value falls to less than 0.30. Comparing the first two cases across levels, shows that perturbing only the means has a greater impact than perturbing only the variances. For the third case, the performance values are similar to the first case except for the 75% to 100% level. At this level, it appears the variance had enough of an impact to cause the additional difference. Finally, for the fourth case, the performance values were about 0.53 for all four levels.

An analysis of variance was performed for each of the four cases at each of the four levels of forecast change. The results are displayed in Tables B.10 - B.25. A summary of the significant factors is shown in Table 4.18. Comparing the significant factors for each type of sequence across the three utility functions shows no strong pattern of significant factors across the different levels for each of the four cases.

EX5 Observations

As the results of EX5 show, perturbing the means by as little as 25% caused the perturbed and base sequences to differ quite dramatically. While if the means were always perturbed in the same direction the results were better, the performance was still only in the 50% range. In practical terms, even if the sequences differ, as long as the first asset types match, the DM has less incentive to try to resolve inconsistencies. However, if the first asset type is sensitive to changes in the forecasts, then inconsistencies present a problem. Since, for a given level, perturbing the variances has a much smaller impact than perturbing the means, increasing the variance (uncertainty) may be the easiest way to resolve inconsistencies.

4.3 Summary of Parameter Effects

The MV model parameters were found to have varying levels of impact. The parameters were the level of risk aversion, type of correlation, level of uncertainty, horizon, asset monetary difference, number of asset types, maximum asset service life, discount rate, and rate of technological improvement. The impact for each is described below. The summary includes the parameter's effect on the performance of the alternative decision procedures, its influence on the difficulty of finding the sequence of highest EU both in terms of the maximum size of the MV-efficient set and the CPU time required, and any unusual or unexpected behavior that was observed.

The **level of risk aversion** had an impact on the performance of the EV decision procedure as well as on the CME procedure when the DM had decreasing risk aversion. However, the amount of impact was influenced by the level of uncertainty. As was found in EX3, the EV decision procedure performed best for low risk aversion and for low uncertainty. EX4 illustrated that the EU sequences obtained using three different utility functions were very similar when the risk aversion levels for all three were similar. However, for the logarithmic and power utility functions, which reflect decreasing risk aversion, the performance of the CME procedure was poor. Since the branch and prune procedure is based on MV-dominance and does not use the level of risk aversion until a choice must be made at the final stage, the difficulty of finding the sequence of highest EU was not affected by the level of risk aversion.

Correlation had an impact on both the performance of the EV and CME decision procedures, and increased the CPU times compared to the independent case. In EX3, the EV decision procedure was found to perform best for negative correlation and worst for positive correlation. Negative correlation reduced the variance (uncertainty) causing the EV and EU sequences to more closely match. For constant risk aversion, the CME decision procedure always found an optimal sequence when assets were assumed to be independent. However, the CME performance decreased slightly when the assets were assumed to be correlated, especially when the correlation was negative. To obtain the EU sequence when assets were assumed to be correlated, all the covariance terms had to be calculated which increased the CPU times.

As **uncertainty** increased, the performance of the EV procedure decreased unless the level of risk aversion was very low. The level of uncertainty was expressed in terms of the median NPV coefficient of variation. As was found in EX2, the EV and EU sequences began to differ when the median NPV coefficient of variation was about 0.3 and the NPVs were generated in the regular manner. For the AE method of generating the NPVs, differences were not observed until ratio was slightly higher due to the lower relative uncertainty. The sequence of highest EU was easier to find when uncertainty was high since larger monetary differences typically existed among the asset type and service life pairs.

The **horizon** influenced the performance of the EV and CME decision procedures as well as the difficulty of obtaining the sequence of highest EU. As the horizon increased, the CME decision procedure sometimes selected suboptimal assets late in the sequence. At these later time periods, the effects of discounting often reduced the mean NPVs to near zero. Since the utility function was assumed to be almost linear near zero, the CME decision procedure was, in effect, selecting assets based on the highest EV. In EX3, the combined effect of the horizon and the level of risk aversion was found to be a significant factor for the EV decision procedure. The EV decision procedure's utility performance was best for low risk aversion and short horizons. As was expected, longer horizons typically resulted in both higher numbers of MV-efficient sequences and higher CPU times.

Both the **asset monetary difference** and **number of asset types** were found to have little impact on the performance of the alternative decision procedures, but to impact the maximum number of MV-efficient nodes and CPU times. If the asset monetary difference was large, then an asset type and service life pair could dominate. Typically, the larger the differences between the asset type and service life pairs, the smaller the CPU times that were required. Since a few asset type and service life pairs often repeated in the sequences, the number of asset types made little difference in the performance of alternative decision procedures. However, as the number of asset types increased, CPU times increased because more computations had to be performed at each stage.

The **maximum asset service life** had little impact on the performance of the alternative decision procedures, again because a few asset type and service life pairs often repeated. However, more so than any other parameter except the horizon, the maximum asset service life significantly influenced the CPU times. The CPU times for a maximum life of 14 periods were typically 20 times those for a maximum life of 3 periods.

The **discount rate** had little influence on the performance of the alternative decision procedures, but as it increased, the sequence of highest EU become harder to obtain. As EX1 showed, higher discount rates diminished the differences among the asset type and service life pairs, especially if the rate of technological improvement was low. With only small differences, the MV-efficient set grew rapidly and clustered. Higher discount rates resulted in stronger clustering.

Since the **rate of technological improvement** was always selected randomly, the amount of influence on the performance of the alternative decision procedures cannot be estimated. For deterministic cash flows, Bean, Lohmann, and Smith [5] found the TRAD decision procedure to perform better when technology improved slowly rather than rapidly. The impact of the rate of technological improvement on the difficultly of obtaining EU sequences was easier to estimate. Problems in which the rate of technology was greater than the discount rate turned out to be easy to solve because the effects of discounting did not diminish the differences in the asset type and service life pairs over time. Slow technological change meant that any differences would tend to be diminished much faster, with the MV-efficient set size and CPU times increasing.

The method used to handle technological improvement was observed to have two weaknesses. To forecast easily NPVs of future assets, the computational experiments assumed the DM inputs the NPVs for just the assets currently available. Future assets were then assumed to evolve deterministically from the base set (i.e., $NPV(J,T,N) = NPV(J,0,N) \times F(J,T) \ / \ (1+m)^T$). The first weakness is that for a given asset type, if a particular combination of mean and variance dominates, then it will dominate at all points in time. For example, consider two different sequences that provide service for the first three periods. In the first sequence, asset type J is replaced each year with a new version of asset type J. In the second sequence, asset type J is keep in service for the full three periods. If the first sequence MV-dominates the second sequence, then it will also MV-dominate the second sequence at any common starting period in the future. If the $F(J,T)$s are similar across different asset types, then dominance will tend to hold at all points in time as well. This pattern of MV-dominance may be one of the reasons why particular asset type and service life pairs often repeated in the sequences for the test problems.

The second weakness of the method used to handle technological improvement arises when computing the mean and variance for assets with negative mean NPVs. Using $NPV(J,T,N) = NPV(J,0,N) \times F(J,T) \ / \ (1+m)^T$, the resulting mean is then $\mathrm{E}[NPV(J,T,N)] = \mathrm{E}[NPV(J,0,N)] \times F(J,T) \ / \ (1+m)^T$, while the variance is given by $\mathrm{Var}[NPV(J,T,N)] = \mathrm{Var}[NPV(J,0,N)] \times F(J,T)^2 \ / \ (1+m)^{2T}$. When the mean NPV is positive and technology improves (i.e., $F(J,T) > 1.0$), both the mean and the variance of the NPV increase. When the mean NPV is negative, and technology improves (i.e., $F(J,T) < 1.0$), the mean NPV again increases (becomes less negative). However, the mean NPV can never become positive. The result is that the absolute increase in the mean NPV is smaller (and decreasing in time) as compared to the positive mean NPV case. Additionally, the variance of the NPV decreases instead of increasing. This weakness could be mitigated by expanding the model to allow technological improvement for multiple cash flow components (e.g., capital costs, operating expenses, revenues) rather than for just the overall NPV [42].

CHAPTER V

SUMMARY AND EXTENSIONS

This dissertation examined replacement decision making under uncertainty. The research objectives were to develop a procedure to find the sequence of highest EU, gain insight about replacement decisions under uncertainty, and compare the performance of several alternative decision prócedures with the EU decision procedure. The MV model is the first utility-based replacement model to be developed. A utility-based model is desirable because DMs often view the impact of monetary gains and losses differently and such a view would typically be inconsistent with an EV criterion that treats equally the impact of both. Considering risk directly adds a dimension that traditional sensitivity analysis cannot address. Moreover, a utility-based model can solve a wider range of problems since EV is a special case of EU.

The MV model allows the number of assets available for selection at each point in time to vary, and offers the the DM considerable flexibility in specifying the NPV distributions. If the DM assumes that assets evolve over time and that future versions are related to the assets currently available by a function of the time they are installed, the NPVs can be made up of any number of cash flow components. Each cash flow component can change over time due to technology and inflation, and the discount rate may vary by period. The MV model is also the first replacement model to consider asset correlation across time. However, with correlation the branch and prune procedure is only a heuristic and finding tight bounds is difficult.

This research makes two additional contributions beyond development of the first utility-based replacement model. First, the average performance of the EU decision procedure was found to be better than that of the TRAD and EV decision procedures, and better than that of the CME decision procedure for a DM with decreasing risk aversion. The TRAD decision procedure corresponds to making the best decision considering only the assets available currently. This approach is taken frequently by DMs who feel it is too difficult to forecast the monetary outcomes for assets available in the future. In EX3, the TRAD decision procedure's average performance was just 0.64. Furthermore, no factors were found to have a significant effect on the performance. Thus, the TRAD decision procedure performs poorly even for DMs with low risk aversion and for problems with short horizons relative to the EU decision procedure. In EX3, the EV decision procedure's average performance was 0.86, with uncertainty and the level of risk aversion combined with the horizon having a significant effect on the performance. When uncertainty, characterized by the median NPV coefficient of variation, reached about 0.3, the EV and EU sequences began to differ. The EV decision procedure performed best (0.94) when both the level of risk aversion and the horizon were low. The CME decision procedure's average performance was 0.99 in EX3 for a DM with constant aversion to risk. Thus, a DM with constant risk aversion would have little incentive to use the more computationally intensive EU decision procedure. However, the CME decision procedure's average performance fell less to than 0.55 in EX4 for a DM with decreasing aversion to risk.

The second contribution was that finding all the MV-efficient (or Pareto-optimal) sequences before applying a decision criterion was found to be a viable solution approach for replacement problems of realistic size. The MV model considered problems with horizons of up to 50 periods, the number of asset types available at each period as high as 7, and asset service lives of up to 14 periods. No previous study had tried to find all the Pareto-optimal solutions except for small problems used for illustrative purposes. Rather, some decision criterion other than EU was selected in advance, or only a subset of the Pareto-optimal solutions was found. The advance selection requires the problem to be resolved each time the decision criterion changes.

Several factors which increased the solution difficulty (either the maximum size of the MV-efficient list or the CPU times) were identified. These factors were the horizon, maximum asset service life, discount rate, and amount of difference between the rate of technological improvement across asset types. The longer the horizon and maximum service life, the more sequences possible. Higher discount rates and smaller differences in the rate of technological improvement caused the monetary differences across asset types and service lives to decrease. The MV-efficient sequences were found to cluster and clustering was used to develop a heuristic. The average performance of the cluster heuristic as compared to an upper bound was 0.99 in EX3 for independent assets. Other heuristics that have been used to find subsets of all the Pareto-optimal solutions do not rely on clustering [2, 20]. Thus, the cluster heuristic represents another alternative. However, it may only be applicable for problems in which discounting is used, given that clustering was found to be more pronounced for higher discount rates.

5.1 Summary

Replacement projects consume a significant portion of many firms' capital budget. As the marketplace becomes more competitive, decisions about what equipment to replace, what new technology to adopt, and when to do it become increasingly important. Forecasting the data required to make intelligent replacement decisions is typically difficult and filled with uncertainty. Utility theory, a popular method in many fields for dealing with decision making under risk, has been incorporated into capital budgeting models. However, utility has yet to be directly considered in replacement models. The MV model was developed to study replacement decision making for a DM who desires to maximize EU. It envisions the DM trying to identify the best sequence of assets to provide a desired service over time. The DM will typically have a number of different asset types from which to choose, and be uncertain about the cash flow streams that characterize each asset type. The MV model makes the following major assumptions:

- The DM is not subject to capital rationing and the economic value of the service provided by an asset can be characterized by its NPV distribution.
- The NPV distribution for the service provided by an asset is normally distributed.
- The NPVs of assets in a given sequence are independent or Markov-correlated across time. If Markov-correlated, the NPV of an asset being considered for addition to a sequence is correlated only with the NPV of the immediately preceding asset in the sequence.

- A single asset is assumed to be in service at all times and there are a finite number of independent assets from which to choose at each decision point.

- The DM's objective is to maximize EU over the specified, deterministic horizon. For a given sequence, EU = $\int U(w) f_{npv}(w|\mu, \sigma^2) dw$, where the utility function, $U(w)$, is a continuous, concave, monotonically increasing function of money, and f_{npv} is the density function for the sequence NPV.

For the computational experiments, three alternative solution procedures were compared to the EU decision procedure. The alternative solution procedures were the three most common in the literature: TRAD, EV, and CME. The TRAD sequence was formed by selecting the asset type and service life pair with the highest expected annual equivalent value at each replacement epoch. The EV sequence had the maximum expected NPV, while the CME sequence had the maximum total CME, found by summing the CME for each asset in the sequence. Both sequences were found by solving a longest path problem using deterministic dynamic programming. The sequence of highest EU was found using a branch and prune procedure based on MV-dominance. The branch and prune procedure is similar to the dynamic programming algorithms used to find the EV and CME sequences, except that instead of identifying a single best partial sequence at each time period, a MV-efficient set of partial sequences is found. A sequence is MV-efficient if no other sequence has a greater mean and equal variance, or a greater than or equal mean and a smaller variance.

To allow a wide range of problems to be considered, a test problem generator with a set of variable parameters was built. The parameters included the horizon, number of alternative asset types, asset annual cost pattern, rate of asset technological improvement, maximum asset service life, level and type of correlation across sequential assets, DM's level of risk aversion, amount of uncertainty in the DM's forecasts, and discount rate. Experiments were then performed to determine if the branch and prune could optimally solve a wide range of problems, under what conditions did the EU, EV, and CME sequences differ, what was the impact of different types of utility functions, and how sensitive was the EU sequence to changes in the forecasts. Three main results were obtained from the first two experiments. First, the MV-efficient nodes for the branch and prune procedure could sometimes exceed the available computer memory. Second, the MV-efficient nodes were found to cluster and clustering was used to develop a heuristic. Third, the EU decision procedure performed well when assets were assumed to be independent. The third experiment studied the impact of risk aversion, correlation, uncertainty, and the horizon on the performance of the alternative decision procedures. Results showed the EU, EV, and CME sequences matched for low uncertainty and low risk aversion. The CME procedure was an excellent decision procedure for constant risk aversion because the CME of each asset sums to almost the CME for the entire sequence. Correlation, higher levels of risk aversion, and higher uncertainty tended to cause the EU, EV, and CME sequences to differ. The fourth experiment compared three different types of utility functions and found the EU sequences tended to match for similar levels of risk aversion. The CME was also found to perform quite poorly for decreasing risk aversion. The final experiment found the EU sequence to be much more sensitive to changes in the means than the variances.

5.2 Extensions

The MV model could be extended in several different ways. First, for a specific problem, the expected value of perfect information could be approximated using Monte Carlo simulation. To do this, the EU sequence would first be found. Then for each simulation replication, a realization for each of the DM's forecasts would be drawn. These realizations would be used to find the optimal sequence as well as the realized value of the EU sequence. The difference between the optimal and EU sequence NPV would represent a realization of the value of perfect information. After n simulation replications, the expected value of perfect information could be estimated. Any change in the DM's forecasts would require the simulation to be repeated. To avoid having to repeat the simulation, an analytical approach, such as used by Moore and Chen [40] for a capital budgeting problem, may be able to be used for very simple problems.

Second, an attempt could be made to develop additional dominance criteria to eliminate some MV-efficient nodes in the branch and prune procedure, such that for the independent asset case the procedure is still optimal. Given that the EU to upper bound utility performance of the cluster heuristic was on average 0.95 in EX2, it is clear that eliminating nodes using stochastic dominance rules for truncated normal distributions is very effective. Since in most problems the effect of discounting reduces the impact of assets later in the sequence, it may be possible at some point to start eliminating certain partial sequences without compromising optimality. Insight may to be able to be gained by studying the cases when suboptimal selections are made. Such additional criteria could improve the efficiency of the branch and prune procedure and reduce or even eliminate the need for the cluster heuristic.

Third, research could be conducted from a utility-based capital budgeting perspective to study the EV and EU decision procedures for selecting individual replacement projects. For example, the capital budgeting approach of Thompson and Thuesen [54] could be used. Computational experiments could then be run to try to determine the conditions under which choosing the best alternative for individual replacement projects based on EV results in suboptimal capital allocation.

Fourth, two additional assumptions could be made so that an additive temporal utility function would be appropriate. For a temporal utility function, the EU of a sequence would be the weighted sum of the EU for the cash flow in each period. This approach could then be compared to determining the EU based on the NPV of the sequence, as was done in this research. The additional assumptions would be that utility independence across time periods is appropriate for the DM, and that the DM can supply the appropriate cash flow forecasts for each time period. By considering a wide range of possible problems, the comparison could provide additional insight into replacement decision making under uncertainty as well as identify some of the conditions under which the two approaches differ. In their utility-based capital budgeting research, Thompson and Thuesen [55] compared a temporal approach based on a project's cash flows in each period with an approach based on the project's NPV.

APPENDICES

APPENDIX A

EXAMPLE OF HOW SELECTING THE ASSET WITH THE HIGHEST PROBABILITY OF BEING THE OPTIMAL FIRST ASSET CAN BE SUBOPTIMAL

Consider a DM faced with a choice of keeping an asset D currently in service or replacing it with a new asset C. The DM's objective is to maximize expected value, the discount rate is 0.10, the horizon is one year, taxes are not involved, and the DM is uncertain about only the operating expenses. D has no salvage value while C, which costs \$13.64 new, will be worth \$10 in one year. Let $\mathbf{X}$ be the random variable representing D's operating expense for the next year and $\mathbf{Y}$ be the random variable representing the expense for C. The DM estimates $\mathbf{X}$ to be uniformly distributed between 12 to 15.8, and $\mathbf{Y}$ to be triangularly distributed with a minimum value of 5, maximum of 16, and most likely of 6. A decision tree expressed in cost terms is shown in Figure A1. At T=1 the expected cost for D = E[$\mathbf{X}$] = 13.9, while the expected cost for C = 5 + E[$\mathbf{Y}$] = 5 + 9 = 14. Thus, if the DM wants to maximize expected value, the optimal choice is D because it has the lower expected cost.

The approach used by Lohmann [35] combines simulation and deterministic dynamic programming. For each replication of the simulation, a realization is drawn from the appropriate distribution for each random variable. The best sequence that

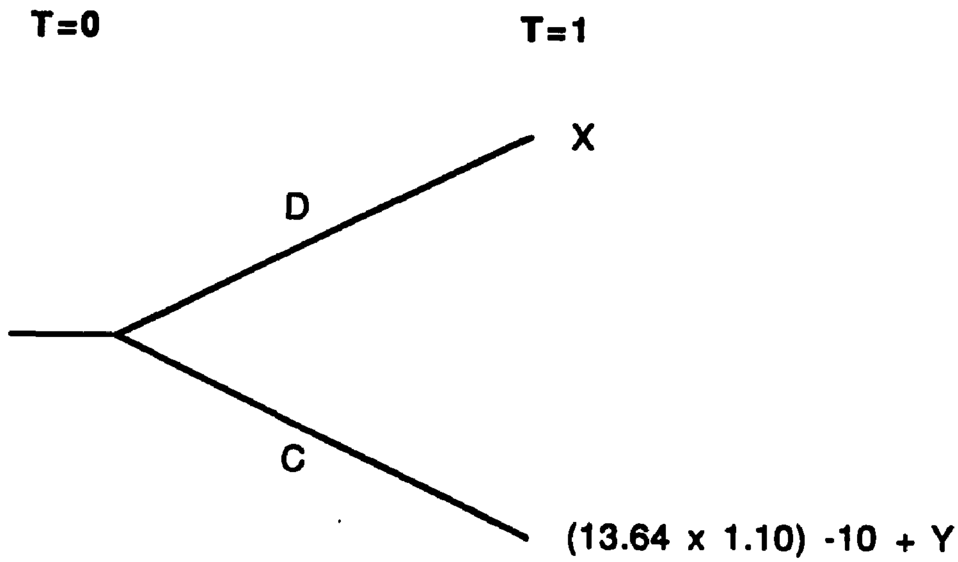

Figure A.1: Decision Tree

begins with each asset currently available is then found using deterministic dynamic programming, and the first asset's service life and the overall sequence NPV are recorded. The best sequence for each asset currently available is found so that the performance measures estimated after n replications are based on an equal number of observations. Once the best sequence has been found for each asset currently available, the overall optimal sequence for the current replication is identified and the optimal first asset recorded. After n replications, the probability of each asset currently available being the overall optimal first asset is estimated, along with the corresponding probability distributions for the service life and NPV. The DM then uses subjective judgement to interpret the results and make a decision.

For the problem above, suppose the decision rule used by the DM was to select the asset that had the highest probability of being the optimal first asset. The simulation involves generating realizations for **X** and **Y**. Dynamic programming is not required because if $5 + \mathbf{Y} < \mathbf{X}$, then C is the lower cost choice and hence the optimal first asset. Such a simulation was replicated 25,000 times and the results indicated C to be the optimal first asset 53.2 percent of the time. To check the simulation results, conditional probability was used to analytically determine that the Prob$\{5 + \mathbf{Y} < \mathbf{X}\}$ is 53.1 percent. Thus, the DM would select C rather than making the optimal decision by selecting D.

APPENDIX B

ANALYSIS OF VARIANCE TABLES FOR EX4 AND EX5

Note: In the tables in Appendix B, factors that are significant at $\alpha = 0.10$ are followed by an asterisk in the F_0 column.

Table B.1: Analysis of Variance for EX4 TRAD, Exponential Utility

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	0.526412	1	0.5264	3.878*
C	0.844650	2	0.4223	3.111*
U	0.095812	1	0.0958	0.706
H	0.491602	1	0.4916	3.621*
RC	0.059529	2	0.0298	0.219
RU	0.023300	1	0.0233	0.172
RH	0.012585	1	0.0126	0.093
CU	0.211932	2	0.1060	0.781
CH	0.076817	2	0.0384	0.283
UH	0.067303	1	0.0673	0.496
RCU	0.012321	2	0.0062	0.045
RCH	0.173469	2	0.0867	0.639
RUH	0.005228	1	0.0052	0.039
CUH	0.102222	2	0.0511	0.377
RCUH	1.908452	2	0.9542	7.029*
Error	13.03219	96	0.1358	
TOTAL	17.6438	119		

Table B.2: Analysis of Variance for EX4 TRAD, Logarithmic Utility

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	0.318917	1	0.3189	2.243
C	1.005692	2	0.5028	3.537*
U	0.079239	1	0.0792	0.557
H	0.405768	1	0.4058	2.854*
RC	0.119609	2	0.0598	0.421
RU	0.017615	1	0.0176	0.124
RH	0.003679	1	0.0037	0.026
CU	0.129350	2	0.0647	0.455
CH	0.294852	2	0.1474	1.037
UH	0.114420	1	0.1144	0.805
RCU	0.015545	2	0.0078	0.055
RCH	0.185621	2	0.0928	0.653
RUH	0.000357	1	0.0004	0.003
CUH	0.204477	2	0.1022	0.719
RCUH	1.409211	2	0.7046	4.956*
Error	13.64848	96	0.1422	
TOTAL	17.9529	119		

Table B.3: Analysis of Variance for EX4 TRAD, Power Utility

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	0.106657	1	0.1067	0.665
C	0.549621	2	0.2748	1.714
U	0.272893	1	0.2729	1.702
H	0.215387	1	0.2154	1.343
RC	0.144071	2	0.0720	0.449
RU	0.130168	1	0.1302	0.812
RH	0.012949	1	0.0129	0.081
CU	0.020044	2	0.0100	0.063
CH	0.308582	2	0.1543	0.962
UH	0.265369	1	0.2654	1.655
RCU	0.096909	2	0.0485	0.302
RCH	0.036892	2	0.0184	0.115
RUH	0.038280	1	0.0383	0.239
CUH	0.431513	2	0.2158	1.346
RCUH	1.168450	2	0.5842	3.643*
Error	15.39372	96	0.1604	
TOTAL	19.1915	119		

Table B.4: Analysis of Variance for EX4 EV, Exponential Utility

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	0.475221	1	0.4752	13.219*
C	0.335837	2	0.1679	4.671*
U	0.859031	1	0.8590	23.896*
H	0.361163	1	0.3612	10.047*
RC	0.099943	2	0.0500	1.390
RU	0.251324	1	0.2513	6.991*
RH	0.065270	1	0.0653	1.816
CU	0.318652	2	0.1593	4.432*
CH	0.400221	2	0.2001	5.567*
UH	0.172987	1	0.1730	4.812*
RCU	0.029674	2	0.0148	0.413
RCH	0.167682	2	0.0838	2.332
RUH	0.005080	1	0.0051	0.141
CUH	0.369276	2	0.1846	5.136*
RCUH	0.093437	2	0.0467	1.300
Error	3.451066	96	0.0359	
TOTAL	7.4559	119		

Table B.5: Analysis of Variance for EX4 EV, Logarithmic Utility

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	0.104984	1	0.1050	4.173*
C	0.307487	2	0.1537	6.111*
U	0.393272	1	0.3933	15.631*
H	0.325735	1	0.3257	12.947*
RC	0.103493	2	0.0517	2.057
RU	0.025376	1	0.0254	1.009
RH	0.047862	1	0.0479	1.902
CU	0.280549	2	0.1403	5.575*
CH	0.163458	2	0.0817	3.248*
UH	0.170431	1	0.1704	6.774*
RCU	0.026411	2	0.0132	0.525
RCH	0.055433	2	0.0277	1.102
RUH	0.003848	1	0.0038	0.153
CUH	0.162803	2	0.0814	3.235*
RCUH	0.011270	2	0.0056	0.224
Error	2.415343	96	0.0252	
TOTAL	4.5978	119		

Table B.6: Analysis of Variance for EX4 EV, Power Utility

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	0.054776	1	0.0548	2.430
C	0.294833	2	0.1474	6.540*
U	0.446639	1	0.4466	19.815*
H	0.305061	1	0.3051	13.534*
RC	0.070917	2	0.0355	1.573
RU	0.027648	1	0.0276	1.227
RH	0.029957	1	0.0300	1.329
CU	0.262295	2	0.1311	5.818*
CH	0.181554	2	0.0908	4.027*
UH	0.236814	1	0.2368	10.506*
RCU	0.046724	2	0.0234	1.036
RCH	0.031901	2	0.0160	0.708
RUH	0.011882	1	0.0119	0.527
CUH	0.167862	2	0.0839	3.724*
RCUH	0.018707	2	0.0094	0.415
Error	2.163887	96	0.0225	
TOTAL	4.3515	119		

Table B.7: Analysis of Variance for EX4 CME, Exponential Utility

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	0.010507	1	0.0105	1.145
C	0.013915	2	0.0070	0.758
U	0.014357	1	0.0144	1.565
H	0.017655	1	0.0177	1.924
RC	0.015731	2	0.0079	0.857
RU	0.004701	1	0.0047	0.512
RH	0.007511	1	0.0075	0.819
CU	0.014526	2	0.0073	0.792
CH	0.012341	2	0.0062	0.673
UH	0.009800	1	0.0098	1.068
RCU	0.023356	2	0.0117	1.273
RCH	0.018541	2	0.0093	1.011
RUH	0.002692	1	0.0027	0.294
CUH	0.015289	2	0.0076	0.833
RCUH	0.027975	2	0.0140	1.525
Error	0.880719	96	0.0092	
TOTAL	1.0896	119		

Table B.8: Analysis of Variance for EX4 CME, Logarithmic Utility

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	5.487145	1	5.4871	53.555*
C	0.152103	2	0.0761	0.742
U	1.802685	1	1.8027	17.594*
H	0.158773	1	0.1588	1.550
RC	0.147439	2	0.0737	0.720
RU	2.590644	1	2.5906	25.285*
RH	0.086972	1	0.0870	0.849
CU	0.034787	2	0.0174	0.170
CH	0.018783	2	0.0094	0.092
UH	0.024425	1	0.0244	0.238
RCU	0.174159	2	0.0871	0.850
RCH	0.017850	2	0.0089	0.087
RUH	0.248662	1	0.2487	2.427
CUH	0.112143	2	0.0561	0.547
RCUH	0.116856	2	0.0584	0.570
Error	9.836037	96	0.1025	
TOTAL	21.0095	119		

Table B.9: Analysis of Variance for EX4 CME, Power Utility

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	2.764369	1	2.7644	26.847*
C	0.005469	2	0.0027	0.027
U	5.087518	1	5.0875	49.409*
H	0.005877	1	0.0059	0.057
RC	0.100141	2	0.0501	0.486
RU	1.667244	1	1.6672	16.192*
RH	0.164614	1	0.1646	1.599
CU	0.212371	2	0.1062	1.031
CH	0.298324	2	0.1492	1.449
UH	0.000094	1	0.0001	0.001
RCU	0.100509	2	0.0503	0.488
RCH	0.049674	2	0.0248	0.241
RUH	0.116155	1	0.1162	1.128
CUH	0.146735	2	0.0734	0.713
RCUH	0.161837	2	0.0809	0.786
Error	9.884791	96	0.1030	
TOTAL	20.7657	119		

Table B.10: Analysis of Variance for EX5 Mean Only, (0.00,0.25)

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	0.012722	1	0.0127	0.065
C	0.123068	2	0.0615	0.315
U	0.039664	1	0.0397	0.203
H	0.103194	1	0.1032	0.528
RC	0.081395	2	0.0407	0.208
RU	0.035750	1	0.0358	0.183
RH	0.246968	1	0.2470	1.263
CU	0.431538	2	0.2158	1.103
CH	0.327803	2	0.1639	0.838
UH	0.023686	1	0.0237	0.121
RCU	0.219333	2	0.1097	0.561
RCH	0.477763	2	0.2389	1.221
RUH	0.017387	1	0.0174	0.089
CUH	0.492615	2	0.2463	1.259
RCUH	0.581285	2	0.2906	1.486
Error	18.77763	96	0.1956	
TOTAL	21.9918	119		

Table B.11: Analysis of Variance for EX5 Mean Only, (0.25,0.50)

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	0.072382	1	0.0724	0.532
C	0.065437	2	0.0327	0.241
U	0.002456	1	0.0025	0.018
H	0.241295	1	0.2413	1.774
RC	0.290534	2	0.1453	1.068
RU	0.011619	1	0.0116	0.085
RH	0.001147	1	0.0011	0.008
CU	0.264999	2	0.1325	0.974
CH	0.017389	2	0.0087	0.064
UH	0.120863	1	0.1209	0.889
RCU	1.385582	2	0.6928	5.094*
RCH	0.018787	2	0.0094	0.069
RUH	0.005548	1	0.0055	0.041
CUH	0.388991	2	0.1945	1.430
RCUH	0.393463	2	0.1967	1.447
Error	13.05639	96	0.1360	
TOTAL	16.3369	119		

Table B.12: Analysis of Variance for EX5 Mean Only, (0.50,0.75)

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	0.413355	1	0.4134	2.701
C	0.563570	2	0.2818	1.841
U	0.000310	1	0.0003	0.002
H	0.118140	1	0.1181	0.772
RC	0.386757	2	0.1934	1.264
RU	0.003389	1	0.0034	0.022
RH	0.348175	1	0.3482	2.275
CU	0.234623	2	0.1173	0.767
CH	0.369096	2	0.1845	1.206
UH	0.005292	1	0.0053	0.035
RCU	0.093747	2	0.0469	0.306
RCH	0.036208	2	0.0181	0.118
RUH	0.025354	1	0.0254	0.166
CUH	0.297816	2	0.1489	0.973
RCUH	0.112093	2	0.0560	0.366
Error	14.69182	96	0.1530	
TOTAL	17.6998	119		

Table B.13: Analysis of Variance for EX5 Mean Only, (0.75,1.00)

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	0.060113	1	0.0601	0.444
C	0.220042	2	0.1100	0.813
U	0.001262	1	0.0013	0.009
H	0.374084	1	0.3741	2.765
RC	0.063917	2	0.0320	0.236
RU	0.008413	1	0.0084	0.062
RH	0.222598	1	0.2226	1.645
CU	0.879870	2	0.4399	3.252*
CH	0.104130	2	0.0521	0.385
UH	0.018593	1	0.0186	0.137
RCU	0.190978	2	0.0955	0.706
RCH	0.398493	2	0.1992	1.473
RUH	0.111464	1	0.1115	0.824
CUH	0.033271	2	0.0166	0.123
RCUH	0.095149	2	0.0476	0.352
Error	12.98729	96	0.1353	
TOTAL	15.7697	119		

Table B.14: Analysis of Variance for EX5 Variance Only, (0.00,0.25)

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	0.139093	1	0.1391	1.497
C	0.125823	2	0.0629	0.677
U	0.016028	1	0.0160	0.173
H	0.556668	1	0.5567	5.992*
RC	0.678913	2	0.3395	3.654*
RU	0.079141	1	0.0791	0.852
RH	0.042430	1	0.0424	0.457
CU	0.534587	2	0.2673	2.877*
CH	0.253631	2	0.1268	1.365
UH	0.344069	1	0.3441	3.704*
RCU	0.689896	2	0.3449	3.713*
RCH	0.395465	2	0.1977	2.128
RUH	0.043753	1	0.0438	0.471
CUH	0.674076	2	0.3370	3.628*
RCUH	0.240826	2	0.1204	1.296
Error	8.918255	96	0.0929	
TOTAL	13.7327	119		

Table B.15: Analysis of Variance for EX5 Variance Only, (0.25,0.50)

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	0.053329	1	0.0533	0.340
C	0.003100	2	0.0016	0.010
U	0.652624	1	0.6526	4.156*
H	1.138934	1	1.1389	7.253*
RC	0.015219	2	0.0076	0.048
RU	0.000095	1	0.0001	0.001
RH	0.000895	1	0.0009	0.006
CU	0.190506	2	0.0953	0.607
CH	0.261004	2	0.1305	0.831
UH	0.258357	1	0.2584	1.645
RCU	0.480351	2	0.2402	1.529
RCH	0.219364	2	0.1097	0.698
RUH	0.018252	1	0.0183	0.116
CUH	0.177566	2	0.0888	0.565
RCUH	0.226366	2	0.1132	0.721
Error	15.07582	96	0.1570	
TOTAL	18.7718	119		

Table B.16: Analysis of Variance for EX5 Variance Only, (0.50,0.75)

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	0.453762	1	0.4538	2.831*
C	0.216407	2	0.1082	0.675
U	0.726585	1	0.7266	4.533*
H	0.029653	1	0.0297	0.185
RC	0.377332	2	0.1887	1.177
RU	0.933942	1	0.9339	5.827*
RH	0.036987	1	0.0370	0.231
CU	0.057724	2	0.0289	0.180
CH	0.316324	2	0.1582	0.987
UH	0.010428	1	0.0104	0.065
RCU	0.427281	2	0.2136	1.333
RCH	0.266483	2	0.1332	0.831
RUH	0.199247	1	0.1992	1.243
CUH	0.209763	2	0.1049	0.654
RCUH	0.062361	2	0.0312	0.195
Error	15.38650	96	0.1603	
TOTAL	19.7108	119		

Table B.17: Analysis of Variance for EX5 Variance Only, (0.75,1.00)

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	0.562612	1	0.5626	3.930*
C	0.163401	2	0.0817	0.571
U	0.827964	1	0.8280	5.783*
H	0.479212	1	0.4792	3.347*
RC	0.358865	2	0.1794	1.253
RU	1.064305	1	1.0643	7.434*
RH	0.219700	1	0.2197	1.535
CU	0.113894	2	0.0569	0.398
CH	0.336604	2	0.1683	1.176
UH	0.000450	1	0.0005	0.003
RCU	0.250615	2	0.1253	0.875
RCH	0.255006	2	0.1275	0.891
RUH	0.092174	1	0.0922	0.644
CUH	0.502296	2	0.2511	1.754
RCUH	0.088868	2	0.0444	0.310
Error	13.74464	96	0.1432	
TOTAL	19.0606	119		

Table B.18: Analysis of Variance for EX5 Mean and Variance, (0.00,0.25)

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	0.516805	1	0.5168	2.794*
C	0.945305	2	0.4727	2.555*
U	0.149145	1	0.1491	0.806
H	0.033438	1	0.0334	0.181
RC	0.086611	2	0.0433	0.234
RU	0.028411	1	0.0284	0.154
RH	0.687670	1	0.6877	3.718*
CU	0.053348	2	0.0267	0.144
CH	1.246533	2	0.6233	3.370*
UH	0.002102	1	0.0021	0.011
RCU	0.298530	2	0.1493	0.807
RCH	1.007162	2	0.5036	2.723*
RUH	0.045319	1	0.0453	0.245
CUH	0.075064	2	0.0375	0.203
RCUH	0.407952	2	0.2040	1.103
Error	17.75629	96	0.1850	
TOTAL	23.3397	119		

Table B.19: Analysis of Variance for EX5 Mean and Variance, (0.25,0.50)

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	1.115908	1	1.1159	7.584*
C	0.535148	2	0.2676	1.818
U	0.506804	1	0.5068	3.444*
H	0.448281	1	0.4483	3.047*
RC	0.001759	2	0.0009	0.006
RU	0.126360	1	0.1264	0.859
RH	0.061615	1	0.0616	0.419
CU	0.167327	2	0.0837	0.569
CH	0.065324	2	0.0327	0.222
UH	0.005125	1	0.0051	0.035
RCU	0.760220	2	0.3801	2.583*
RCH	0.645483	2	0.3227	2.193
RUH	0.010761	1	0.0108	0.073
CUH	0.048353	2	0.0242	0.164
RCUH	0.425902	2	0.2130	1.447
Error	14.12574	96	0.1471	
TOTAL	19.0501	119		

Table B.20: Analysis of Variance for EX5 Mean and Variance, (0.50,0.75)

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	0.476264	1	0.4763	3.396*
C	0.316522	2	0.1583	1.129
U	0.003621	1	0.0036	0.026
H	0.354237	1	0.3542	2.526
RC	0.112662	2	0.0563	0.402
RU	0.314605	1	0.3146	2.243
RH	0.021994	1	0.0220	0.157
CU	0.077420	2	0.0387	0.276
CH	0.266595	2	0.1333	0.951
UH	0.129985	1	0.1300	0.927
RCU	0.218564	2	0.1093	0.779
RCH	0.056552	2	0.0283	0.202
RUH	0.000440	1	0.0004	0.003
CUH	0.103020	2	0.0515	0.367
RCUH	1.370756	2	0.6854	4.887*
Error	13.46270	96	0.1402	
TOTAL	17.2859	119		

Table B.21: Analysis of Variance for EX5 Mean and Variance, (0.75,1.00)

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	0.013802	1	0.0138	0.112
C	0.363115	2	0.1816	1.469
U	0.018589	1	0.0186	0.150
H	0.145035	1	0.1450	1.173
RC	0.099717	2	0.0499	0.403
RU	0.097532	1	0.0975	0.789
RH	0.021228	1	0.0212	0.172
CU	0.204662	2	0.1023	0.828
CH	0.228251	2	0.1141	0.923
UH	0.000857	1	0.0009	0.007
RCU	0.427938	2	0.2140	1.731
RCH	0.225417	2	0.1127	0.912
RUH	0.002843	1	0.0028	0.023
CUH	0.184459	2	0.0922	0.746
RCUH	0.064160	2	0.0321	0.260
Error	11.86760	96	0.1236	
TOTAL	13.9652	119		

Table B.22: Analysis of Variance for EX5 Mean,Same Direction, (-0.50,-0.25)

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	0.147807	1	0.1478	0.816
C	0.094835	2	0.0474	0.262
U	0.132832	1	0.1328	0.734
H	0.001362	1	0.0014	0.008
RC	0.122172	2	0.0611	0.337
RU	0.201935	1	0.2019	1.115
RH	0.591231	1	0.5912	3.266*
CU	0.526680	2	0.2633	1.455
CH	1.329042	2	0.6645	3.670*
UH	0.053144	1	0.0531	0.294
RCU	0.166136	2	0.0831	0.459
RCH	0.189320	2	0.0947	0.523
RUH	0.264668	1	0.2647	1.462
CUH	0.350770	2	0.1754	0.969
RCUH	0.142854	2	0.0714	0.395
Error	17.38038	96	0.1810	
TOTAL	21.6952	119		

Table B.23: Analysis of Variance for EX5 Mean,Same Direction, (-0.25,0.00)

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	0.001187	1	0.0012	0.007
C	1.104026	2	0.5520	3.146*
U	0.000000	1	.0000	.000
H	0.015715	1	0.0157	0.090
RC	0.958876	2	0.4794	2.733*
RU	0.326965	1	0.3270	1.864
RH	0.567165	1	0.5672	3.233*
CU	0.065439	2	0.0327	0.187
CH	0.527275	2	0.2636	1.503
UH	0.010608	1	0.0106	0.060
RCU	0.170804	2	0.0854	0.487
RCH	1.104504	2	0.5523	3.148*
RUH	0.180820	1	0.1808	1.031
CUH	0.322985	2	0.1615	0.921
RCUH	0.137964	2	0.0690	0.393
Error	16.84201	96	0.1754	
TOTAL	22.3364	119		

Table B.24: Analysis of Variance for EX5 Mean,Same Direction, (0.00,0.25)

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	0.070365	1	0.0704	0.355
C	1.132349	2	0.5662	2.860*
U	0.019844	1	0.0198	0.100
H	0.109654	1	0.1097	0.554
RC	0.141809	2	0.0709	0.358
RU	0.028115	1	0.0281	0.142
RH	0.209180	1	0.2092	1.057
CU	0.889976	2	0.4450	2.248
CH	0.501413	2	0.2507	1.266
UH	0.055476	1	0.0555	0.280
RCU	0.349468	2	0.1747	0.883
RCH	1.156598	2	0.5783	2.921*
RUH	0.005428	1	0.0054	0.027
CUH	0.139145	2	0.0696	0.351
RCUH	0.282640	2	0.1413	0.714
Error	19.00621	96	0.1980	
TOTAL	24.0977	119		

Table B.25: Analysis of Variance for EX5 Mean,Same Direction, (0.25,0.50)

Source of Variance	Sum of Squares	Degrees of Freedom	Mean Square	F_0
R	0.001469	1	0.0015	0.008
C	0.781693	2	0.3908	2.140
U	0.036110	1	0.0361	0.198
H	0.005587	1	0.0056	0.031
RC	0.086744	2	0.0434	0.237
RU	0.403339	1	0.4033	2.208
RH	0.681572	1	0.6816	3.732*
CU	0.574158	2	0.2871	1.572
CH	0.148293	2	0.0741	0.406
UH	0.767466	1	0.7675	4.202*
RCU	0.808309	2	0.4042	2.213
RCH	1.368168	2	0.6841	3.745*
RUH	0.031399	1	0.0314	0.172
CUH	0.025787	2	0.0129	0.071
RCUH	0.337494	2	0.1687	0.924
Error	17.53417	96	0.1826	
TOTAL	23.5918	119		

BIBLIOGRAPHY

BIBLIOGRAPHY

[1] Baksh, M.S., *Effectiveness of Capital Budgeting Procedures for Dealing With Risk*, PhD thesis, The University of Michigan, 1986.

[2] Bard, J.F., and Bennett, J.E., "Arc Reduction and Path Preference in Stochastic Acyclic Networks," *Management Science*, vol. 37, no. 3, pp. 198–215, February 1991.

[3] Barnes, J.W., Zinn, C.D., and Eldred, B.S., "A Methodology for Obtaining the Probability Density Function of the Present Worth of Probabilistic Cash Flow Profiles," *AIIE Transactions*, vol. 10, no. 3, pp. 226–236, September 1978.

[4] Bawa, V.S., "Stochastic Dominance: A Bibliography," *Management Science*, vol. 28, pp. 698–712, 1982.

[5] Bean, J.C., Lohmann, J.R., and Smith, R.L., "A Dynamic Infinite Horizon Replacement Economy Decision Model," *The Engineering Economist*, vol. 30, no. 2, pp. 99–120, Winter 1985.

[6] Bean, J.C., Lohmann, J.R., and Smith, R.L., "Equipment Replacement Under Technological Change," Working Paper, The University of Michigan, Department of Industrial and Operations Engineering, October 1988.

[7] Benzion, U., and Yagil, J., "On the Discounting Formula for a Stream of Independent Risky Cash Flows," *The Engineering Economist*, vol. 32, no. 4, pp. 337–345, Summer 1987.

[8] Berger, J.O., *Statistical Decision Theory and Bayesian Analysis*, Springer-Verlag, Second edition, 1980.

[9] Brockett, P.L., and Golden, L.L., "A Class of Utility Functions Containing all the Common Utility Functions," *Management Science*, vol. 33, no. 8, pp. 955–964, August 1987.

[10] Currim, I.S., and Sarin, R.K., "Prospect Versus Utility," *Management Science*, vol. 35, no. 1, pp. 22–42, January 1989.

[11] Daellenbach, H.G., and De Kluyver, C.A., "Note on Multiple Objective Dynamic Programming," *Journal of the Operational Research Society*, vol. 31, pp. 591–594, 1980.

[12] Denardo, E.V., *Dynamic Programming: Models and Applications*, Prentice-Hall, 1982.

[13] Dreyfus, S.E., and Law, A.M., *The Art and Theory of Dynamic Programming*, Academic Press, 1977.

[14] Farragher, E.J., "Capital Budgeting Practices of Non-Industrial Firms," *The Engineering Economist*, vol. 31, no. 4, pp. 293–302, Summer 1986.

[15] Ganesh, K., *Serial Replacement Under Evolving Productivity*, PhD thesis, The University of Michigan, 1987.

[16] Grant, E.L., Ireson, W.G., and Leavenworth, R.S., *Principles of Engineering Economy*, John Wiley and Sons, Seventh edition, 1982.

[17] Gurnani, C., "Capital Budgeting: Theory and Practice," *The Engineering Economist*, vol. 30, no. 1, pp. 19–46, Fall 1984.

[18] Hanoch, G., and Levy, H., "The Efficiency Analysis of Choices Involving Risk," *The Review of Economic Studies*, vol. 36, pp. 335–346, 1969.

[19] Hansen, P., "Bicriterion Path Problems," in *Multiple Criteria Decision Making Theory and Application*, Randel, G., and Gal, T., editor, Springer-Verlag, 1980.

[20] Henig, M.I., "The Shortest Path Problem with Two Objective Functions," *European Journal of Operational Research*, vol. 25, pp. 281–291, 1985.

[21] Henig, M.I., "Risk Criteria in a Stochastic Knapsack Problem," *Operations Research*, vol. 38, no. 5, pp. 820–825, September-October 1990.

[22] Hillier, F.S., "The Derivation of Probabilistic Information for the Evaluation of Risky Investments," *Management Science*, vol. 9, no. 3, pp. 443–457, April 1963.

[23] Hillier, F.S., "Supplement to 'The Derivation of Probabilistic Information for the Evaluation of Risky Investments'," *Management Science*, vol. 11, no. 3, pp. 485–487, January 1965.

[24] Hillier, F.S., "A Basic Model for Capital Budgeting of Risky Interrelated Projects," *The Engineering Economist*, vol. 17, no. 1, , Fall 1971.

[25] Jaquette, S.C., "A Utility Criterion for Markov Decision Processes," *Management Science*, vol. 23, no. 1, pp. 43–49, September 1976.

[26] Kallberg, J.G., and Ziemba, W.T., "Comparison of Alternative Utility Functions in Portfolio Selection Problems," *Management Science*, vol. 29, no. 11, pp. 1257–1276, November 1983.

[27] Keeney, R.L., and Raiffa, H., *Decisions With Multiple Objectives: Preferences and Value Tradeoffs* , John Wiley, 1976.

[28] Kira, D., and Ziemba, W.T., "Equivalence Among Alternative Portfolio Selection Criteria," in *Financial Decision Making Under Uncertainty*, Levy, H., and Sarnat, M., editor, Academic Press, 1977.

[29] Laughhunn, D.J., Payne, J.W., and Crum, R., "Managerial Risk Preferences for Below-Target Returns," *Management Science*, vol. 26, no. 12, pp. 1238–1249, November 1980.

[30] Leung, L.C., and Tanchoco, J.M.A., "Multiple Machine Replacement Within an Integrated System Framework," *The Engineering Economist*, vol. 32, no. 2, pp. 89–114, Winter 1987.

[31] Levy, H., "Stochastic Dominance Rules for Truncated Normal Distributions: A Note," *The Journal of Finance*, vol. 37, no. 5, pp. 1299–1303, December 1982.

[32] Levy, H., "Two-Moment Decision Models and Expected Utility Maximization: Comment," *The American Economic Review*, vol. 79, no. 3, pp. 597–600, June 1989.

[33] Levy, H., and Paroush, J., "Multi-Period Stochastic Dominance," *Management Science*, vol. 21, no. 4, pp. 428–435, December 1974.

[34] Levy, H., and Sarnat, M., *Financial Decision Making Under Uncertainty*, Academic Press, 1977.

[35] Lohmann, J.R., "A Stochastic Replacement Economic Decision Model," *IIE Transactions*, vol. 18, no. 2, pp. 182–194, June 1986.

[36] Luxhoj, J.T., and Jones, M.S., "A Framework for Replacement Modeling Assumptions," *The Engineering Economist*, vol. 32, no. 1, pp. 39–49, Fall 1986.

[37] MacCrimmon, K.R., and Wehrung, D.A., "Characteristics of Risk Taking Executives," *Management Science*, vol. 36, no. 4, pp. 422–435, April 1990.

[38] March, J.G., and Shapira, Z., "Managerial Perspectives on Risk and Risk Taking," *Management Science*, vol. 33, no. 11, pp. 1404–1418, November 1987.

[39] Markowitz, H.M., "Portfolio Selection," *The Journal of Finance*, vol. 6, no. 1, pp. 77–91, March 1952.

[40] Moore, W.T., and Chen, S., "The Value of Perfect Information in Capital Budgeting Decisions with Unknown Cash Flow Parameters," *The Engineering Economist*, vol. 29, no. 1, pp. 41–51, Fall 1983.

[41] Oakford, R.V., Bhimjee, S.A., and Salazar, A., "The Maximum Prospective Utility Criterion, Utility Maximization in Capital Rationing Under Incomplete Information," *The Engineering Economist*, vol. 26, no. 3, pp. 181–206, Spring 1981.

[42] Oakford, R.V., Lohmann, J.R., and Salazar, A., "A Dynamic Replacement Economy Decision Model," *IIE Transactions*, vol. 16, no. 1, pp. 65–72, March 1984.

[43] Oakford, R.V., Salazar, A., and Digiulio, H.A., "Factors That Affect the Growth Rate of Equity Capital," *IIE Transactions*, vol. 17, no. 2, pp. 123–131, June 1985.

[44] Park, C.S., "Probabilistic Benefit-Cost Analysis," *The Engineering Economist*, vol. 29, no. 2, pp. 83–100, Winter 1984.

[45] Park, C.S., "The Mellin Transform in Probabilistic Cash Flow Modeling," *The Engineering Economist*, vol. 32, no. 2, pp. 115–134, Winter 1987.

[46] Park, C.S., and Sharp, G.P., *Advanced Engineering Economics*, John Wiley, 1990.

[47] Park, C.S., and Wu Yeh, M.M., "The Effectiveness of the Expected Value-Variance-Skewness Criterion in the Multi-Stage Capital Rationing Problem," *IIE Transactions*, vol. 17, no. 1, pp. 67–74, March 1985.

[48] Porter, R.B., and Gaumnitz, J.E., "Stochastic Dominance vs. Mean-Variance Portfolio Analysis: An Empirical Evaluation," *The American Economic Review*, vol. 62, no. 3, pp. 438–446, June 1972.

[49] Pratt, J.W., "Risk Aversion in the Small and in the Large," *Econometrica*, vol. 32, no. 1-2, pp. 122–136, January-April 1964.

[50] Ramis, F.J., Thuesen, G.J., and Barr, T.J., "A Dynamic Target-Wealth Criterion for Capital Investment Decisions," *The Engineering Economist*, vol. 36, no. 2, pp. 107–126, Winter 1991.

[51] Ross, S.M., *Applied Probability Models with Optimization Applications*, Holden Day, 1970.

[52] Sigal, C.E., Pritsker, A.B., and Solberg, J.J., "The Stochastic Shortest Route Problem," *Operations Research*, vol. 28, no. 5, pp. 1122–1129, September-October 1980.

[53] Swalm, R.O., "Utility Theory–Insights Into Risk Taking," *Harvard Business Review*, pp. 123–138, November-December 1966.

[54] Thompson, R.A., and Thuesen, G.J., "Dynamic Investment Criteria for Capital Budgeting Decisions," *The Engineering Economist*, vol. 31, no. 1, pp. 1–26, Fall 1985.

[55] Thompson, R.A., and Thuesen, G.J., "Applications of Dynamic Investment Criteria for Capital Budgeting Decisions," *The Engineering Economist*, vol. 33, no. 1, pp. 59–86, Fall 1987.

[56] Vander Veen, D.J., *Parallel Replacement Under Nonstationary Deterministic Demand*, PhD thesis, The University of Michigan, 1985.

[57] Wagle, B., "A Statistical Analysis of Risk in Capital Investment Projects," *Operations Research Quarterly*, vol. 18, no. 1, pp. 13–33, March 1967.

[58] Wagner, H.M., *Principles of Operations Research*, Prentice-Hall, Inc., Englewood Cliffs, Second edition, 1975. 349-356.

[59] Ziemba, W.T., and Vickson, R.G., *Stochastic Optimization Models in Finance*, Academic Press, 1975.