DECISION MAKING WITH INCOMPLETE INFORMATION

A Dissertation
Presented to
The Academic Faculty

By

Marc Christopher Canellas

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

Georgia Institute of Technology

August 2017

Approved by:

Dr. Karen M. Feigh, Advisor
School of Aerospace Engineering
Georgia Institute of Technology

Dr. Brian J. German
School of Aerospace Engineering
Georgia Institute of Technology

Dr. Amy R. Pritchett
School of Aerospace Engineering
Georgia Institute of Technology

Dr. Juan D. Rogers
School of Public Policy
Georgia Institute of Technology

Dr. Stephen E. Cross
School of Industrial and Systems Engineering
Georgia Institute of Technology

Date Approved: April 21, 2017

We will have to repent in this generation not merely

for the hateful words and actions of the bad people

but for the appalling silence of the good people.

— Martin Luther King Jr.

To Rachel, Orca, and my family

TABLE OF CONTENTS

List of Tables ix

List of Figures xiii

Summary xviii

Chapter 1: Introduction and Motivation 1

1.1 Research Question 1: Do some distributions of incomplete information result in better decision making performance than others? If so, when and why? 3

1.2 Research Question 2: Can it be determined which information to acquire or restrict without using probabilities, cue weights, or cue scores? If so, when and why? 5

1.3 Framing the Decision Task 7

1.4 Dissertation Overview 9

Chapter 2: Literature Review 11

2.1 Human Rationality 11

2.2 Incomplete Information 15

2.2.1 Time Pressure 15

2.2.2 High Cost of Information 16

2.2.3 Too Much Information 16

2.2.4 Incomplete Information in the Environment 17
2.3 Less is More 18
2.3.1 Fewer Cues 19
2.3.2 Fewer Known Cue Values 22
2.3.3 Simpler Models with Less Training 24
2.4 Information-Based Decision Support 27
2.5 Summary 31
Chapter 3: General Linear Model of Judgment and Decision Making 32
3.1 The General Linear Model 35
3.1.1 Cue Weights 39
3.1.2 Estimates of Missing Information 40
3.1.3 Utility Function 40
3.1.4 Categorizing or Choosing 44
3.2 Modeling and Simulation 47
3.2.1 Modeling Decision Making 47
3.2.2 Modeling Judgment 48
3.2.3 Simulating Judgment and Decision Making 50
3.2.4 Example: Simulating Fast-and-Frugal Trees 51
3.3 Representing Experience, Effort, and Time Pressure 53
3.3.1 Experience 54
3.3.2 Time and Effort 55
3.4 Summary 56

Chapter 4: Models, Measures, and Mediators of Decision Making Performance . 58
4.1 Models . . . 58
4.2 Measures . . . 61
4.2.1 Accuracy . . . 61
4.2.2 Effort and Time Required . . . 64
4.3 Mediators . . . 67
4.3.1 Strategy . . . 67
4.3.2 Environmental Parameters . . . 70
4.3.3 Task Parameters . . . 74
4.3.4 Incomplete Information . . . 75
4.4 Summary . . . 78

Chapter 5: Computer Simulation Study 1: Accuracy, Effort, and Distributions of Incomplete Information . . . 80
5.1 Method . . . 82
5.1.1 Scenario Generation . . . 83
5.1.2 Measuring Accuracy . . . 83
5.1.3 Measurement of Effort . . . 83
5.2 Results . . . 85
5.2.1 Strategy . . . 85
5.2.2 Total Information . . . 89
5.2.3 Option Imbalance . . . 93
5.2.4 Total Information and Option Imbalance . . . 96
5.2.5 Cue Balance and Total Information . . . 98

5.2.6 Dispersion 101

5.2.7 Dominance 102

5.3 Discussion 104

5.3.1 Heuristics in Naturalistic Decision Contexts 104

5.3.2 Variability in Effort for Heuristics 105

5.3.3 Implications for Heuristic Information Acquisition and Restriction in Decision Support 107

5.4 Conclusion 109

Chapter 6: Computer Simulation Study 2: Heuristic Information Acquisition and Restriction Rules for Decision Support 111

6.1 Background 113

6.1.1 Fitting Decision Task with Binary Cues 113

6.1.2 Heuristic Information Acquisition and Restriction Rules 114

6.2 Method 118

6.3 Results 121

6.3.1 General Effectiveness of Heuristic Information Acquisition and Restriction Rules 121

6.3.2 Effect of Strategy Components 124

6.3.3 Effect of Environmental Parameters 126

6.4 A Mathematical Explanation of Estimates and FIA as Mediators of Rule Effectiveness 129

6.5 Discussion 131

6.5.1 Rules for Heuristic Information Acquisition and Restriction 131

6.5.2 Components of Strategies: A New Level of Abstraction 132

6.5.3 Heuristic Decision Support and Decision Support for Heuristics . . 133

6.6 Conclusion 134

Chapter 7: Computer Simulation Study 3: Determinants of Decision Making with Incomplete Information 136

7.1 Background 138

7.1.1 Incomplete information 138

7.1.2 Heuristic Information Acquisition and Restriction Rules at the Aggregate-Level 139

7.2 Method 141

7.2.1 Strategies 141

7.3 Results 145

7.3.1 One-Variable Effects 145

7.3.2 Two-Variable Interaction Effects 147

7.3.3 Heuristic Information Acquisition and Restriction Rules 151

7.3.4 How Valuable is Cue Rank Information? 154

7.4 Discussion 155

7.4.1 Beyond Total Information and Cue Weights 155

7.4.2 Information Acquisition Without Cue Weights, Cue Values, or Probabilities 158

7.5 Conclusion 159

Chapter 8: Validation Human-Subjects Study 161

8.1 Background 161

8.1.1 Difficulty 162

8.1.2 Incomplete Information 162

8.1.3 Estimates of Missing Information 163

8.1.4 Interaction of Estimates of Missing Information and Distributions of Incomplete Information 164

8.2 Method 164

8.2.1 Participants 164

8.2.2 Decision Environment 165

8.2.3 Task Design 166

8.2.4 Materials and Procedure 169

8.3 Results 172

8.3.1 Did Participants Adapt Their Estimates of Missing Information? . . 172

8.3.2 Does Computational "Difficulty" Result in Human Decision Making "Difficulty"? 176

8.3.3 Did Distributions of Incomplete Information Affect Decision Making Performance? 176

8.4 Discussion 182

8.4.1 Falling into the Reality Gap 182

8.4.2 Friendly and Unfriendly Distributions of Incomplete Information . . 184

8.4.3 Generalizing the Results 185

8.5 Conclusion 186

Chapter 9: Discussion 188

9.1 Unifying Judgment and Decision Making Strategies in a Single Mathematical Model 188

9.2 Shedding the Methodological Bias Toward Examining Only "Well-Studied" Strategies 189

9.3 Revaluing the Bias of Overweighting Common Cues (Attributes) as Ecologically Rational . . . 191

9.4 Showing that Distributions of Incomplete Information Matter, Often More than the Total Amount of Information . . . 192

9.5 Defining Heuristics for Information Acquisition and Restriction . . . 194

9.6 Discovering a Potential Reality Gap Between What Makes Distributions Difficult for People Versus their Computational Representations . . . 196

Chapter 10: Summary and Future Research Directions . . . 199

10.1 Summary . . . 199

10.2 Future Research Directions . . . 201

10.2.1 Expanding the General Linear Model . . . 201

10.2.2 Judgment and Decision Making with Incomplete Information . . . 203

10.2.3 Heuristic Information Acquisition and Restriction Rules . . . 205

Chapter A: Dissertation Appendices . . . 208

A.1 General Linear Model of Judgment and Decision Making . . . 208

A.1.1 Extending Non-Compensatory Weights Beyond Strictly Binary Cues 208

A.1.2 Cue weights to categories . . . 210

A.1.3 Prior Formalization of Fast-and-Frugal Trees . . . 212

A.1.4 Dataset Description for Fast-and-Frugal Trees Example . . . 212

A.2 Computer Simulation Studies . . . 214

A.2.1 Study 2: Heuristic Information Acquisition and Restriction Rules for Decision Support . . . 214

A.2.2 Study 3: Determinants of Decision Making with Incomplete Information . . . 216

A.3 Validation Human-in-the-Loop Experiment 227
A.3.1 Institutional Review Board Documents 227
A.3.2 General Experiment Information 231
Chapter B: Publications 235
B.1 Published Articles 235
B.1.1 Journal Articles 235
B.1.2 Conference Papers with Podium Presentations 235
B.1.3 Podium Presentations (Only) 236
B.1.4 Poster Presentations (Only) 236
B.2 Planned 236
B.2.1 Journal Articles 236
B.2.2 Magazine Articles 237
References 237
Vita 256

LIST OF TABLES

1.1 Dissertation organization. 9

2.1 Comparison of normative and heuristic decision support methods. 28

3.1 Definitions of judgment and decision making. 33

3.2 Definitions of the parameters of the general linear model and the alternative names used in judgment and decision making literature. 37

3.3 Mathematical representations of the components of the general linear model. 39

3.4 The four variants of the binary utility function along two dimensions: positive or negative cue directions, and prior or relative cutoff values. 43

3.5 Count of unique categories possible for both non-compensatory and equal cue weights, for a given number of cues. 47

3.6 Parametrization of well-studied judgment and decision making strategies with respect to the general linear model. 49

3.7 List of matrices for the matrix form of the general linear model of judgment and decision making. 50

3.8 Parameter values for two exemplar fast-and-frugal trees. 53

3.9 Performance measures for both example FFT's showing identical results between the online tool (ATO) and the general linear model (GLM). 53

4.1 List of elementary information processes (EIPs) with time estimates (Bettman et al., 1990; Payne et al., 1990). 66

4.2 Description of the five selected decision making strategies using the general linear model of decision making in Chap. 3. 68

5.1 Mapping of elementary information processes (EIPs), which approximate effort, to the individual decision making strategies. 84

5.2 Average accuracy of, and effort required when using, decision strategies for each level of each measure of incomplete information. 86

5.3 Average accuracy of, and effort required when using, decision strategies for each level of each task parameter: dispersion and dominance. 87

5.4 Analysis of effect of strategy, total information, dominance, and disperson on effort required to perform TTB. 88

5.5 Statistical significance of the effect of option balance on the accuracy of each decision strategy within each paired level of total information. 96

5.6 Accuracy and EIP count for each decision making strategy for all 14 combinations of total information and cue balance values. 99

6.1 Environmental parameter values for the 5-cue decision environments. 120

6.2 Positive accuracy count (PAC) and average accuracy change (AAC) for each decision strategy when using heuristic information acquisition and restrictions rules averaged across all decision environments. 122

6.3 Significance of strategy components on the PAC and AAC values for each heuristic information acquisition and restriction rule. 125

6.4 Significance of environmental parameters on the AAC values for each combination of decision making strategy and heuristic information acquisition and restriction rule using linear models. 127

6.5 Resulting decision equations and accuracy as a function of estimates of missing information for a two-option, one-cue decision task with incomplete information. 130

7.1 Characteristics of the seven datasets used in Study 3. 142

7.2 Information acquisition and restriction types for the 2-option decision tasks. 143

7.3 Strategy codes for the 18 variations of the components of the strategies based on the general linear model of decision making. 146

7.4 Accuracy of decision making strategies with incomplete information. Cell colors are conditionally formatted with a blue-white-red transition from low to high accuracy. 150

8.1 Options and criterion scores. 167

8.2 13 distributions of incomplete information examined in the experiment. . . 169

8.3 Number of tasks in which strategies using different mechanisms for treating missing information make different predictions. 173

8.4 The average accuracy and time required for each of the 13 distributions of incomplete information measured, sorted by accuracy. 177

9.1 Important contributions with further discussion in this chapter. 188

A.1 Potential geometric series for cue weights (w_j) for various common ratios (r): $w_j = r^{j-1}$. 209

A.2 Criterion values for all combinations of 2-state binary cue scores for three cues. Non-compensatory cue weights generate 8 unique criterion values while equal cue weights generate 4 unique criterion values. 210

A.3 Criterion values for all combinations of 3-state binary cue scores for three cues. Non-compensatory cue weights generate 27 unique criterion values while equal cue weights generate 7 unique criterion values. 211

A.4 Environmental parameters for the 5-cue decision environments. 214

A.5 Environmental parameters for the 4-cue decision environments. 215

A.6 Environmental parameters for the 3-cue decision environments. 215

A.7 The average accuracy change of the information acquisition and restriction types. 225

A.8 The positive accuracy count of the information acquisition and restriction types. 226

A.9 Incomplete information combinations. 231

A.10 Transformations of experimental cue scores to presented cue scores for the participants. 232

A.11 Description of the decision task "Type" in Tables A.12 and A.13 232

A.12 Description the three blocks of easy difficulty tasks: incomplete information, type, and decision task ID. 233

A.13 Description the three blocks of hard difficulty tasks: incomplete information, type, and decision task ID. 234

LIST OF FIGURES

1.1 Decision scenario with two distributions of incomplete information. 4

1.2 Information acquisition and restriction from an initial decision task. 6

1.3 Decision making for this dissertation will examine how analytic and heuristic strategies perform in static, single decision tasks with incomplete information. 8

3.1 Example of how missing cue score information results in estimated cue scores. The estimate of missing information, e, is set to 0.5. 41

3.2 Binary utility functions for negative and positive cue directions accounting for cue values (a^v), cutoff value (c), and just noticeable differences (Δ). . . 42

3.3 Categorization tree representation of the general linear model for judgment with non-compensatory cue weights and either two cue score states $\{0, 1\}$ or three cue score states $\{0, 0.5, 1\}$. 46

3.4 Two exemplar fast-and-frugal trees for categorization. 52

4.1 Model of contextual determinants of decision making accuracy based on prior literature. 59

4.2 Model of contextual determinants of decision making accuracy based on the research presented in this dissertation. 59

4.3 Lens Model with labeled statistical parameters (Cooksey, 1996; Bass, 2002; Rothrock and Kirlik, 2003). 62

4.4 Description of accuracy and achievement. 65

4.5 Measures of distributions of incomplete information. 76

5.1 Model of the simulation engine for decision making with incomplete information used in Computer Study 1. 82

5.2 Effect of total information on the distribution of effort requirements for decision strategies. The markers indicate the mean EIP count. 91

5.3 Effect of total information, dispersion, and dominance on the effort requirements of TTB. 92

5.4 Effect of option imbalance on the accuracy and information bias of decision strategies. Accuracy data is indicated by the white markers whereas information bias data is indicated by the gray markers. Since this study examines a two-option decision task, an information bias of 50% indicates that a strategy has no preference toward selecting options with more information. 95

5.5 Heatmap of the effect of combinations of total information (x-axis) and option imbalance (y-axis) on the accuracy (intensity of gray shading). Empty boxes indicate that the combination of option imbalance and total information is not possible with this simulation. For example, for a total information of two, there are only two levels of option imbalance possible: zero (both cue values are known about one option) or two (one cue value is known about each option). 97

5.6 Accuracy of each decision making strategy for each combination of total information and cue balance. 100

5.7 Effect of dispersion level on the difference between the WADD-REAL option scores and the accuracy of decision strategies. The box plots describe the distribution of the difference between the two option scores when the dispersion level is low, medium, or high. The gray markers indicate the accuracy of the decision strategies for each level of dispersion. 101

5.8 Effect of restriction of cue values by a decision support system (DSS) to reduce total information and option imbalance on the accuracy and effort of decision making strategies. The accuracy, and effort required, for each decision strategy are the averages for all combinations of cue values. 109

6.1 Examples of the four heuristic information acquisition and restriction rules at the individual decision task level for the exemplar driving scenario. The specific distribution of known and unknown information is denoted by white K's and black ?'s, respectively. 117

6.2 Example of option imbalance acquisition (OI-A) and restriction (OI-R) rules, and cue balance acquisition (CB-A) and restriction (CB-R) rules for TTB strategy for the House dataset with 3 cues. The comparisons of accuracy for acquisition are between the two endpoints of the arrows and the comparisons of accuracy for restriction are shown by the direction of the arrows. The change in accuracy is shown in bold-italic font. 119

6.3 Average accuracy change (AAC) and positive accuracy count (PAC) when using each of the heuristic information acquisition and restriction rules for each decision making strategy: option imbalance acquisition (OI-A) and restriction (OI-R), and cue balance acquisition (CB-A) and restriction (CB-R). AAC measures the average change in accuracy of a strategy when using the rule and PAC measures how often the change in accuracy is positive. Each individual box plot has 45 data points for PAC, 420 data points for OB-R for AAC, and 285 data points for OB-A, CB-A, and CB-R for AAC. The box plots indicate the 25^{th} percentile, median, and 75^{th} percentile. Data points (denoted by parentheses) are outliers if they lie outside the whiskers corresponding to approximately 2.7σ if the data was normally distributed. 123

7.1 Accuracy of decision making strategies as a function of total information, option imbalance, and cue balance. 148

7.2 Exemplar two-way interactions between the three measures of incomplete information. The left column showing the results of the PNN strategy represents strategies whose estimates of missing information did not match the environment. The right column showing the results of the PMN strategy represents strategies whose estimates of missing information did match the environment. 149

7.3 Average accuracy change (AAC) for each information acquisition and restriction type. Parentheses denote the change in cue balance and option imbalance, respectively. 153

7.4 Average accuracy change (AAC) for each information acquisition and restriction type by cue rank. Parentheses denote the change in cue balance and option imbalance, respectively. 156

8.1 Visualization of the decision environment with reference to the two exemplar targets: a missile with a high level of danger and a transport with a low level of danger. 166

8.2 Designing decision tasks to elicit estimates of missing information, information bias, and accuracy. 168

8.3 Example of the interface used in Experiment 1. 170

8.4 Effect of block order on estimate used. For a participant using average estimates, the average estimates would predict decisions in 100% of the tasks, and positive and negative estimates would predict decisions in 69% of the tasks. Error bars represent one standard error. 174

8.5 Effect of block order on accuracy and time required. Error bars represent one standard error. 175

8.6 Effect of total information, option imbalance, and cue balance on participants' decision making accuracy and time. Error bars represent one standard error. 178

8.7 Two-variable interactions between the three measures of incomplete information with performance measured in accuracy (left column) and time required (right column). 180

8.8 Interaction of the distribution of incomplete information (total information, option imbalance, and cue balance) on participants' decision making accuracy and time required, as a function of total information. Exemplar decision tasks with visual distributions of incomplete information are provided for each data point. The dark blocks indicate known cue scores and white blocks indicate unknown cue scores. Error bars represent one standard error. 181

A.1 Two-way interaction between total information and option imbalance for strategies with prior cutoff values and each combination of estimates and weights. 218

A.2 Two-way interaction between total information and option imbalance for strategies with relative cutoff values and each combination of estimates and weights. 219

A.3 Two-way interaction between total information and cue balance for strategies with prior cutoff values and each combination of estimates and weights. 220

A.4 Two-way interaction between total information and cue balance for strategies with relative cutoff values and each combination of estimates and weights. 221

A.5 Two-way interaction between cue balance and option imbalance for strategies with prior cutoff values and each combination of estimates and weights. 222

A.6 Two-way interaction between cue balance and option imbalance for strategies with relative cutoff values and each combination of estimates and weights. 223

A.7 Georgia Tech Institute Review Board approval. 228

A.8 Human-subjects study consent form approved by Georgia Tech Institute Review Board - page 1. 229

A.9 Human-subjects study consent form approved by Georgia Tech Institute Review Board - page 2. 230

A.10 Criterion scores of the 16 options. 232

SUMMARY

Decision makers are continuously required to make choices in environments with incomplete information. This dissertation seeks to understand and, ultimately, support the wide range of decision making strategies used in environments with incomplete information. The results show that the standard measure of incomplete information, i.e., total information, is insufficient for understanding and supporting decision makers faced with incomplete information. Instead, the distribution of information is often a more important determinant of decision making performance. Two new measures of the distribution of incomplete information were introduced (option imbalance and cue balance, Chap. 4) and tested across three computer simulations of 18 variations of decision making strategies within hundreds of environments and millions of decision tasks with incomplete information (Chap. 5-7). The simulations were powered by a new general linear model of decision making which can efficiently and transparently model a wide range of strategies beyond the traditional set in the literature (Chap. 3). Of the many potential mediators of the relationship between the distributions of incomplete information and performance, only the strategies' estimates of missing information were significant in the computational studies (Chap. 5-7). For strategies with accurate estimates, only total information determined accuracy, while for strategies with inaccurate estimates, accuracy was maximized when option imbalance was low and cue balance was high. The simulation results were partially contradicted by a human-subjects study (Chap. 8) in which decision makers with accurate estimates were affected by option imbalance and cue balance in the same manner as inaccurate estimates – suggesting that some distributions might simply be difficult regardless of the estimates (Chap. 8). These results support the argument that decision support should modify the presentation of information away from difficult distributions. This argument was codified into heuristic information acquisition and restriction rules which, when tested, increased strategy accuracy without probability and cue weight information (Chap. 6-7).

CHAPTER 1

INTRODUCTION AND MOTIVATION

Decision support systems are intended to improve decision making effectiveness and efficiency (Todd and Benbasat, 1999). Designing decision support systems depends upon a thorough understanding of the factors that affect decision making performance. The extensive research on decision making within the fields of economics, psychology, and human factors has contributed to defining and characterizing the effect of various contexts on decision making performance. However, one fundamental context of decision making has been left understudied, the impact of incomplete information.

Time pressure, high information acquisition costs, information overload, or ill-structured environments have all been shown to cause people to make decisions with incomplete information. Engineers, doctors, pilots, and warfighters are only a few of the myriad of professions for whom success depends upon operating in some or all of these environmental conditions.

One of the primary ways that a decision support system developer can assist these decision makers within environments with incomplete information is to have the system dynamically change what information it presents to the user (referred to as modification of content in Feigh et al., 2012). To determine what information should or should not be presented requires a fundamental understanding of how distributions of incomplete information affect decision making effectiveness and efficiency. However, prior to this dissertation, the three most relevant research domains for on decision making with incomplete information (fast-and-frugal heuristics, consumer and marketing research, and probabilistic evidence acquisition) were limited in either their ability to characterize what information should be presented or not presented (specificity), explain the effect of incomplete information in a wide range of environments (comprehensiveness), or support decision makers in their real-world environments (practicality).

The fast-and-frugal heuristics research program uses mathematical, computational, and human-subjects studies to show that heuristics, simple rules that use little information, can often be just as accurate or even more accurate than analytic, so-called ‘optimal,’ statistical models, particularly in decision tasks with low amounts of information (e.g. Gigerenzer et al., 1999; Todd et al., 2012). With respect to characterizing various distributions of incomplete information, total information has been the only measure (Martignon and Hoffrage, 2002; Garcia-Retamero and Rieskamp, 2008; Hogarth and Karelaia, 2005a). This means that there are no descriptions for how various distributions of the same amount of information might affect the effectiveness and efficiency of decision making. So, while the results of the fast-and-frugal heuristics research program are comprehensive and practical, they lack specificity for this problem.

In addition to total amount of information, consumer and marketing research has developed some specific measures of incomplete information: common attributes (values are available for all considered options) and unique attributes (values are available for one option but not others) (Slovic and MacPhillamy, 1974; Kivetz and Simonson, 2000). These descriptions of distributions of incomplete information are specific, but they have not been systematically tested (computationally or mathematically) on a diverse set of environments so that they can be incorporated into more formal decision theory. Therefore, these measures of incomplete information lack the comprehensiveness necessary to guide decision support system design.

Lastly, probabilistic frameworks for evidence acquisition describe many methods for determining what piece of information should be acquired (or presented) to encourage better decisions (for a review, see Nelson, 2005). However, these methods require reliable assessments of probabilities which are almost impossible to achieve in real-world environments and are often too opaque or complex to be confidently used by decision makers (Katsikopoulos and Fasolo, 2006; Katsikopoulos et al., 2008). These probabilistic frameworks have been and can be used to specifically identify what information should be acquired; however, they are essentially unusable in many real-world environments.

Given the limitations individually affecting each of these diverse research domains, this dissertation integrates them into a general understanding of decision making with incomplete information which has the ability to characterize what types of information should be presented or not presented, explain the effect of incomplete information in a wide range of environments, and support decision makers in their real-world environments. The path toward this general understanding is guided by two research questions.

1.1 Research Question 1: Do some distributions of incomplete information result in better decision making performance than others? If so, when and why?

Decision making is formally described as the act of choosing between two or more *options* to maximize some measure of *utility* through the assessment of *cues* which characterize the *options*. To achieve the goal of choosing the option with the maximum criterion, decision makers employ strategies which are defined as the overall methods, or sequences of options, for searching through a decision problem space (Payne et al., 1993). Incomplete information occurs when one or more cues (or cue scores) are unknown to the decision maker. In this dissertation, an individual cue can either be complete unknown or known.

The two decision tasks in Fig. 1.1 show two decision tasks that are identical except for having different distributions of incomplete information. For both Decision Task A and Decision Task B, there are three known cue scores and three unknown cue scores (denoted by the question marks). The goal is to select the fastest route between two points – representing the fundamental decision task of many path planning applications.

Based on previous research, which only measured decision making performance as a function of total information, no inference could be made about which distribution of incomplete information in Fig. 1.1 would result in higher accuracy for the decision maker. Essentially, the two decision tasks could be considered equivalent. Examining the two

Decision Task A

	Distance Cue 1	Traffic Cue 2	Road Type Cue 3
Option 1	5	Heavy	?
Option 2	15	?	?

Decision Task B

	Distance Cue 1	Traffic Cue 2	Road Type Cue 3
Option 1	5	Heavy	Highway
Option 2	?	?	?

Figure 1.1: Two decision tasks of the same scenario with different distributions of incomplete information. Question marks refer to unknown cue scores.

decision tasks, however, it seems obvious that they should not be considered equivalent. Decision Task A at least has one cue score known for each option whereas Decision Task B only has information about Option A.

Leveraging the consumer and marketing research, these two options *would* be considered different. The options in Decision Task A has one common cue (Distance) and one unique cue (Traffic, Option 1) whereas in Decision Task B, all the cues are uniquely known for Option 1. However, it is not known how these different distributions of common and unique cues affect decision making effectiveness or efficiency across a wide range of decision environments.

To address RQ1, a computer simulation engine was constructed that could evaluate the accuracy of a wide-range of decision making strategies for any decision scenario[1] and any combination of incomplete information. Using this simulation engine, three computer simulation studies were conducted using the same computer techniques and datasets (containing decision options) that founded the recent advances in decision making research (e.g. Payne et al., 1990; Czerlinski et al., 1999; Hogarth and Karelaia, 2006; Katsikopoulos et al., 2010). The important concentration for this dissertation was the effort to define measures of incomplete information and directly measure their effect on decision making accuracy, including their interactions with components of the decision making strategies and other well-studied environmental and task parameters (characteristics of the datasets

[1] This research focuses on decision tasks with two options and three-to-five cues, but the simulation engine is capable of processing any number of options and cues.

and scenarios that have been shown to affect performance).

A human-in-the-loop (HITL) experiment was then used to validate the impact of these measures on decision making accuracy. As with the computer simulation studies, the set-up of the HITL experiment was based on other researchers' HITL experiments whose participants made decisions with incomplete information (Garcia-Retamero and Rieskamp, 2009; Rieskamp and Hoffrage, 2008). The HITL experiment in this dissertation is the first to focus on how specific distributions of incomplete information affect the effectiveness (accuracy) and efficiency (decision time) of decision makers.

1.2 Research Question 2: Can it be determined which information to acquire or restrict without using probabilities, cue weights, or cue scores? If so, when and why?

RQ2 is a question of the value or usefulness of a piece of information to a decision maker. From decision analysis, the common analytical methods of measuring value of information require reliable assessments of probabilities (the probability that an option is the correct decision given a specific value for one of its cues), cue weights (the relative importance of cues), and cue scores (the value for an option's cue as used by the decision maker) (Nelson, 2005, 2008). However, there are multiple environmental issues with using these analytic methods in decision support (Katsikopoulos and Fasolo, 2006; Katsikopoulos et al., 2008). First, real-world problems can exhibit statistical dependencies which create computational issues in calculating probabilities and cue weights. Second, an operator may not be able to provide accurate estimates of probabilities and cue weights to the DSS. Third, the analytic process may seem too cumbersome, too complex, and too opaque to the operator causing the operator to be reluctant to use the DSS or accept its suggestions. Lastly, as these methods have focused on acquisition, they do not provide suggestions as to what information could be restricted from the operator, to alleviate the growing issue of information overload (Fasolo et al., 2007).

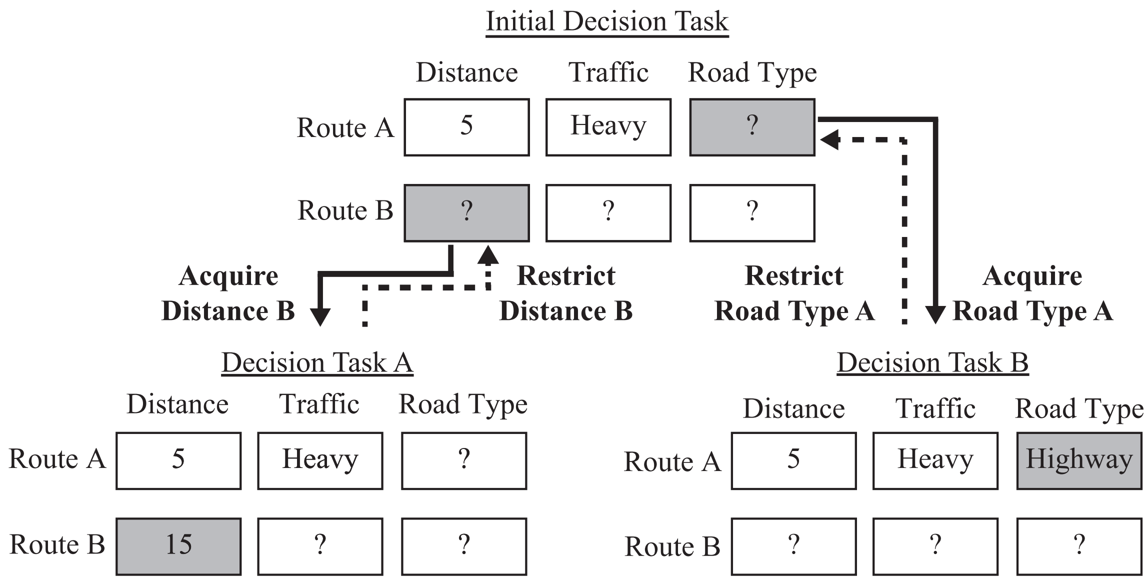

Figure 1.2: Information acquisition and restriction from an initial decision task.

Given these issues, this dissertation leverages the conclusions of RQ1 to develop heuristic information acquisition and restriction rules that are simple, transparent, and do not rely on probabilities, cue weights, or cue scores. Once preferences between distributions of incomplete information can be made (the conclusions of RQ1), then it follows that preferences can be made in terms of acquiring and restricting information so that certain distributions of incomplete information are achieved or avoided (see Fig. 1.2). Since the preferences between distributions of incomplete information do not rely on probabilities, cue weights, or cue scores, the resulting information acquisition or restriction rules do not rely on them either.

To determine whether these heuristic information acquisition and restriction rules are useful, the simulation engine was used. First, general preferences between types of distributions of incomplete information were codified into simple rules for information acquisition and restriction which were shown to be effective at an aggregate level. Then to understand the effectiveness of the rules at the individual decision task level, the computer simulation engine was modified to measure the change in accuracy across a sequence of distributions of incomplete information. The results enabled the understanding of how components of the decision making strategies and the environmental parameters affect the performance of the heuristic information acquisition and restriction rules.

1.3 Framing the Decision Task

Almost all behavioral decision making research using computer simulations utilize a decision task in which each option is encountered simultaneously (non-sequential or single) in an environment which does not change with time or actions (static). Modeling this simpler form of decision making was the original representation of decision tasks (Raiffa, 1968) and has been the starting point for many studies. Most notably, many of the foundational studies supporting the adaptive decision maker program (Payne et al., 1993) and the fast-and-frugal heuristics program (Czerlinski et al., 1999; Katsikopoulos, 2010; Todd et al., 2012) have been based on these single, static tasks.

This dissertation follows these foundational studies and will focus solely on how analytic and heuristic strategies perform in static, single decision tasks with incomplete information (see Fig. 1.3). A single, static decision task can be represented as a table of options, cues, and cue values as shown in Figs. 1.1 and 1.2.

Since these types of tasks are the building blocks of sequential, dynamic tasks, the results of studying these simpler tasks can still have implications for more complicated sequential and dynamic tasks. Future work will take the results of this dissertation and apply it to model and study tasks in which options are encountered at different times based on the environment or actions (sequential) in an environment which does change with time (dynamic).

Sequential decision tasks are those in which options are encountered asynchronously (most commonly, one at a time) resulting in a potential cost for searching through or waiting for the next options (Dudey and Todd, 2001; Todd, 2007). Alternately, following the driving example used in Fig. 1.1, sequential decision making is a down-selection process found in many path planning activities. There are typically multiple routes from home to work,

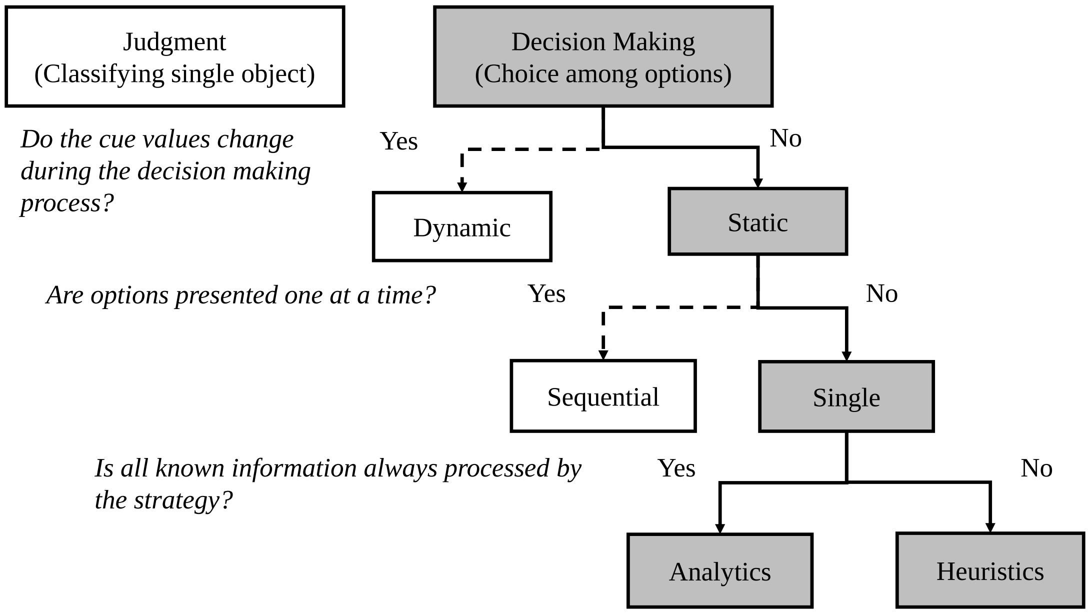

Figure 1.3: Decision making for this dissertation will examine how analytic and heuristic strategies perform in static, single decision tasks with incomplete information.

but as you drive, you are consistently eliminating routes from being able to be chosen. Therefore, the task could be broken down into many single, static decisions in which the goal is to determine which decision options to eliminate from the task or when to eliminate options, rather than select one specific option initially.

Dynamic decision tasks are sequential decision tasks with added difficulty of 1) a dynamic environment that changes as a function of the decision maker's actions or the nature of the environment (Dudey and Todd, 2001; Gonzalez, 2005), and 2) time-delayed responses of the environment to actions made by the decision maker (Gonzalez, 2005). More advanced models incorporate multiple aspects of cognition such as the interaction between memory, learning, and environment such as ACT-R (Anderson, 2007). Given that dynamic decision tasks are sequential, they can also be broken down into single, static decisions. However, the challenge for these tasks is that the action-feedback loop can have very specific impacts on performance that are not easily generalized (Bertuccelli et al., 2009; Gonzalez, 2005).

Table 1.1: Dissertation organization.

Chapter	Title
Chap. 1	Introduction and Motivation
Chap. 2	Literature Review
Chap. 3	General Linear Model of Judgment and Decision Making
Chap. 4	Models, Measures, and Mediators of Decision Making Performance
Chap. 5	Accuracy, Effort, and Distributions of Incomplete Information
Chap. 6	Heuristic Information Acquisition and Restriction Rules for Decision Support
Chap. 7	Determinants of Decision Making with Incomplete Information
Chap. 8	Validation Human-Subjects Study
Chap. 9	Discussion
Chap. 10	Summary and Future Research Directions

1.4 Dissertation Overview

There are 9 chapters in this dissertation in addition to this introduction (Table 1.1). Chapter 2 provides the necessary background to understand the research questions and methodology used in this dissertation. Topics covered include the historical developments of understanding human rationality, why people typically make decisions with incomplete information, and the established measures and mediators of decision making performance. Chapter 3 introduces a general linear model of decision making with incomplete information that was used for computer simulations and related mathematical results. The general linear model also shows how various contextual features of decision making relate to decision making accuracy. Chapter 4 introduces the measures of incomplete information and places them within a full contextual model of the determinants of decision making accuracy. Chapters 5-7 comprise the three computer simulation studies of decision making with incomplete information. The computer simulations provided the initial results showing the importance of option imbalance and cue balance (Chaps. 5 and 7), and the potential for heuristic information acquisition and restriction rules (Chaps. 6 and 7). The human-subjects study in Chapter 8 showed that distributions of incomplete information had an effect on participants with accurate estimates, unlike the computer simulation studies. Chapter 9 discusses the major contributions and their relationship to the general research of decision making and decision support. Lastly, Chapter 10 summarizes the entire dissertation and provides many directions for future research.

CHAPTER 2
LITERATURE REVIEW

This dissertation builds upon a large history of economics, psychology, and decision support system literature to construct a more complete understanding of decision making with incomplete information. The foundation starts in Sec. 2.1, which explains how bounded rationality split from the traditional research of unbounded rationality. Bounded rationality advocates for research into how humans make judgments and decisions when information, computation speed or power, are limited. Section 2.2 examines the specific situations that have been shown to cause humans to make decisions with incomplete information. Section 2.3 addresses the question of whether people actually need to process complete and abundant information in the first place. Rather than viewing people's use of less-than-all-available information as evidence of irrationality, increasing amounts of research has shown that people can and do use fewer cues, fewer known cue values, simpler models, and less training while still making accurate decisions. In comparison, it is clear that most research has focused on how various components of people's decision making strategies are affected by incomplete information, leaving the direct study of incomplete information in the decision task relatively unstudied. Section 2.4 describes normative and heuristic decision support methods for modifying incomplete information and shows how the lack of study of incomplete information has left only normative methods (requiring accurate knowledge of cue weights, cue scores, and probabilities) to guide information acquisition and restriction methods.

2.1 Human Rationality

Competing views of human rationality attempt to explain how people process the information available to them when making judgments and decisions. The division between unbounded and bounded rationality is the most fundamental split in decision making theory.[1] As defined by Todd (2007, pp. 1318, emphasis added), “unbounded rationality says that decisions *should* be made by gathering and processing all available information, without concern for the human mind’s computation speed or power,” whereas bounded rationality says that decisions *are* made within constraints such as the human mind’s limited computation speed or power, or limited information.

Unbounded rationality suggests that decision makers should be provided with as much information as possible so that they can make the decision by processing as much information as possible. There are two problems with this perspective (see Sec. 2.2). The first is descriptive: people do not generally make decisions by processing as much information as possible. The second problem is prescriptive: providing decision makers with as much information as possible does not provide a good basis for building decision support systems (Katsikopoulos and Fasolo, 2006).

Simon (1955, p. 99), famously and influentially, responded to the unbounded rationality perspective by challenging researchers to “replace the global rationality of economic man with a kind of rational behavior that is compatible with the access to information and the computational capacities that are actually possessed by organisms, including man, in the kinds of environments in which such organisms exist.” This counterpoint to unbounded rationality is now known as bounded rationality.

Three additional works in the 1950’s similarly influenced the study of judgment and decision making research (Edwards, 1954; Meehl, 1954; Hammond, 1955). Edwards (1954) integrated the psychological ideas of human behavior with the economic decision theory on the topics of riskless choices, risky choices, transitivity in decision making, games, and statistical decision functions. Hammond (1955) developed statistical representations of Brunswick’s (1943) Lens Model, providing the first quantitative measurements of the rela-

[1]For more recent and more comprehensive reviews of the judgment and decision making literature see the following: Kleinmuntz (1990); Goldstein and Hogarth (1997); Katsikopoulos et al. (2008); Mosier and Fischer (2010); Katsikopoulos (2011).

tionship between the environment and the behavior of judges in the environment. Meehl (1954) published a comparison of clinical (human intuition) versus actuarial (statistical or mechanical) judgment that was the first comprehensive test of human judgment. The results showing that clinicians were usually outperformed by actuarial methods were extremely provocative and provided the model for many comparisons today between normative-rational and heuristic decision making methods (Katsikopoulos et al., 2008).

All four works (Simon, 1955; Edwards, 1954; Meehl, 1954; Hammond, 1955) guided researchers toward studying how decision making occurs in the real world, leading to the emergence of two different research paradigms based on bounded rationality. The first paradigm acquiesced to the limitations of the human mind but searched for the limitations in order to overcome them. This paradigm is summarized as "we would, and should, all be unboundedly rational, if only we could" (Todd, 2007, p. 1319). For example, Tversky and Kahneman's "heuristics-and-biases" research found that people to not necessarily judge uncertainty according to the rules of probability and statistics as they are influenced by framing and other cognitive biases (Tversky and Kahneman, 1974; Kahneman et al., 1982; Kahneman and Tversky, 1984). This perspective of human cognitive suboptimality has resulted in searches for methods to debias people, in other words, to *fix* their behavior (Kleinmuntz, 1990; Todd, 2007).

The second and contrasting paradigm emerged more recently showing that decision makers can make good, and potentially optimal, decisions with simple rules or heuristics that use little information and process it quickly. This work is embodied by three research frameworks. First, within the adaptive decision maker framework, people trade a strategy's cost (the effort required to use the strategy) against the benefits (the perceived accuracy of the strategy) and select a strategy with the best effort-accuracy trade-off in a given environment (Payne et al., 1988, 1990, 1993; Marewski and Link, 2014). Second, within the naturalistic decision making framework, people faced with time pressure, uncertainty, ill-defined goals, and other complexities in familiar and meaningful environments, have been shown use various strategies such as matching strategies to situations, experience-based knowledge modeling, and implementable-based strategy selection (Klein et al., 1988; Orasanu and Connolly, 1993; Klein and Calderwood, 1996; Klein, 2008; Lipshitz et al., 2001). Third, within the fast-and-frugal heuristics framework, people use the bounds on rationality of simplicity, speed, and frugality as a mechanism for simple, robust, and accurate strategies that adapt to the environment and ecology (Gigerenzer and Goldstein, 1996; Gigerenzer et al., 1999; Gigerenzer and Gaissmaier, 2011; Todd et al., 2012). Notably, Keller et al. (2010) argued that the independent frameworks of naturalistic decision making and fast and frugal heuristics can ultimately be synthesized into a "naturalistic heuristics" framework.

To directly show how this new paradigm of heuristic decision making can justify the previously considered cognitive suboptimiality of human decision making, Gigerenzer (2004, p. 66, Table 4.1) provided examples of "phenomena that were first interpreted as 'cognitive illusions' but later revalued as reasonable judgments given the environmental structure." Given the current need for studying decision making with incomplete information, there is potential for this dissertation to revalue further cognitive illusions as reasonable. For example, decision makers in paired comparison tasks tend to overweight 'common' attributes (cues that have information known for both options) and underweight 'unique' attributes (cues that have information known for only one option) (Slovic and MacPhillamy, 1974; Kivetz and Simonson, 2000). This tendency was robust to various debiasing techniques (Slovic and MacPhillamy, 1974): (1) measuring all dimensions on the same scale, (2) prewarning subjects about the bias in favor of common attributes, (3) providing feedback (after each judgment) that promotes equal weighting of dimensions, (4) providing monetary rewards for equal weighting of dimensions, and (5) providing detailed information about the distributions of the values of the attributes. By studying incomplete information directly, the mathematical, computational, and human-subjects studies in this dissertation should be able to determine if or when the focus on common attributes is rational.

2.2 Incomplete Information

The adaptive decision maker, naturalistic decision making, and fast and frugal heuristics frameworks all support the bounded rationality perspective that people make decisions *without* gathering or processing all possible information. This section explains the four major reasons why people do not or cannot make decisions with incomplete information: time pressure, high information acquisition costs, too much information, and incomplete information in the environment. While they are separated for discussion, they are often coincident.

2.2.1 Time Pressure

Dudey and Todd (2001, pp. 195–196) established an evolutionary argument for time pressure as the reason why humans use heuristics:

> "Humans and other animals frequently must make decisions in as rapid a manner as possible. To hesitate is often to be lost, whether this means losing an opportunity for a meal or a mate to a competitor, or losing one's life or limb to a predator or otherwise hostile environment. Organisms seeking to make choices and take action as quickly as possible can either speed up the input side of the decision process, reducing the time needed to gather information on which to base those choices, or speed up the processing/output side, reducing delays in processing the information and converting it into behavior. While evolution has developed faster processing mechanisms over the eons (e.g. myelinization to increase nerve impulse travel speed), the greatest time advantage is likely to come from reducing the amount of information sought before making a decision (Todd and Gigerenzer, 2000)."

As described above, there are two general causes of time pressure: opportunity cost and individual decision limits. Rieskamp and Hoffrage (2008) empirically identified that when time pressure was high for both types, people used heuristics. Previous studies supported this conclusion that under time stress people tend to focus their attention on fewer, more important pieces of information, and thus use a more selective information search (Payne et al., 1988; Maule, 1994; Rieskamp and Hoffrage, 1999).

2.2.2 High Cost of Information

Information acquisition can cost time, as in time pressure. However, what if there is no time pressure but the information costs another resource? For example, a stock broker may need to know which company's stock to buy, but each piece of information about a company costs money and the stock broker only has a fixed amount of initial money (see Experiment 3 in Bröder, 2000), or an engineering designer may need to run costly experiments to explore a design space (Heller, 2013; Jones et al., 1998). In studies like these, people have been show to prefer heuristics when the cost of information is high (Bröder, 2000; Newell and Shanks, 2003; Newell et al., 2003). The explanation of this result is identical to time pressure: when acquiring or processing information costs more time, money, or other commodity than people want to spend, then they use heuristics to focus on fewer, more important pieces of information, in order to save time, money, etc. while retaining as much accuracy as possible.

2.2.3 Too Much Information

People also tend to restrict their information search when too much information is provided (Kivetz and Simonson, 2000; Fasolo et al., 2007). This perspective was introduced as "information (over)load" in consumer research in the 1970's with a self-described controversy (Malhotra et al., 1982) as to whether more information was better (Russo, 1974; Summers, 1974; Wilkie, 1974) or worse (Jacoby et al., 1974a,b; Malhotra, 1982).

By now it is commonly accepted that more information has the potential to be worse and this perspective is often referred to as 'the tyranny of choice' (Fasolo et al., 2007, in reference to Schwartz, 2003). The tyranny of choice is increasingly relevant as technological advancements have introduced online shopping, big data, and geonomic medicine, among others. Note also that, although the discussion regarding too much information has focused on too many options, researchers have also investigated too may cues with similar results (Fasolo et al., 2007; Malhotra, 1982). The abundance of available information for decision makers can make the decision difficult and dissatisfying (e.g., Beattie et al., 1994; Schwartz, 2003). One example is regret aversion, where anticipated regret promotes regret-minimizing decisions rather than risk-minimizing decisions (Zeelenberg et al., 1996; Zeelenberg and Beattie, 1997). Beyond dissatisfaction or difficulty of decision making with too much information, investigation of the accuracy-effort trade-off showed that the amount of effort required to make a decision can lead people to use heuristics because normative-rational strategies require processing of information about all options and all cues (Payne et al., 1988, 1990, 1993).

2.2.4 Incomplete Information in the Environment

The most salient cause of people making decisions with incomplete information is that not all information was available, either because it was difficult to discern, not provided, or other environmental causes. Two separate sets of literature have directly examined this type of incomplete information: naturalistic decision making and consumer marketing.

Naturalistic decision making "typically takes place in a world of incomplete and imperfect information. The decision maker has information about some part of the problem but not others" (Orasanu and Connolly, 1993, p. 8). The study of naturalistic decision settings provides a varied source of potential environmental reasons for decision making with incomplete information (Orasanu and Connolly, 1993; Lipshitz et al., 2001; Klein et al., 1988; Klein, 1993, 2008). Of the eight naturalistic decision settings described by Orasanu and Connolly (1993, p. 7), two are particularly relevant to decision making with incomplete information: ill-structured problems and uncertain dynamic environments. Ill-structured problems occur when significant work needs to be done to develop appropriate options and determine the cue scores which might be related by complex causal links. Uncertain dynamic environments occur when the environment can change quickly, i.e., within the time frame of the required decision, and the information may be ambiguous or of poor quality.

In consumer marketing environments, marketers have significant control over the information provided, or not provided, to consumers and the manner in which this information is presented (Kivetz and Simonson, 2000). In relation to the information overload discussions, Kivetz and Simonson (2000, p. 427) have studied the opposite case for consumer decision making information:

> "Even when complete information is potentially available, obtaining attribute [scores] for all options and making comparisons is typically much easier for some attributes (e.g. price) than for others (e.g. reliability). Accordingly, a common problem consumers face is making choices with complete and easily compared information on some attributes but only partial (or difficult to compare) information on other attributes."

Because of this difficulty, some have argued that in actual consumer decision making, consumers rarely have complete information (Dick et al., 1990; Simmons and Lynch, 1991; Ross and Creyer, 1992).

2.3 Less is More

The previous section described the myriad of reasons why people make decisions with incomplete information: time pressure, high information acquisition costs, too much information, and incomplete information in the environment. But, how much information do people need to make good decisions? The push of modern technology argues that gathering and analyzing all available information with all available tools will result in more accurate decisions. In other words, more resources means more accuracy. Yet, beyond the advertising and normative-rational perspectives, recent computational and human-subject studies suggest that, in many environments, less is actually more (Todd, 2007). Rather than viewing people's use of less-than-all-available information as evidence of irrationality, increasing numbers of studies are showing that people do use or could use fewer cues, fewer known cue values, simpler models, or less training while still making accurate decisions (for reviews, see Gigerenzer et al., 1999; Todd et al., 2012).

2.3.1 Fewer Cues

The hallmark of both the naturalistic decision making and the fast-and-frugal heuristics program is that people tend to use heuristics to focus on few, important cues, even in environments with abundant information. In many comparative tests among analytic strategies and heuristic strategies for judgment and decision making the heuristic strategies which rely on fewer cues have shown to be just as, or more, accurate and have been shown to often better match the strategies people use.

There are conditions to this broad statement about the usefulness of using fewer cues. It is true that focusing on the most important cues typically results in lower proportion of value loss when ignoring cues (Fasolo et al., 2007). This is especially true when cues are positively correlated and validities unequal. However, equally important cues with negative correlations can create 'unfriendly' environments for heuristic strategies using fewer than all cues (Shanteau and Thomas, 2000).

One group of well-studied strategies that use fewer cues is the one-reason decision heuristics. Gigerenzer and Gaissmaier (2011) categorized these one-reason decision heuristics into two classes. The first class is the "one-clever-cue" heuristics which rely on a single cue to make decisions. For example, animals do not analytically compute trajectories in three-dimensional space when intercepting incoming prey or catching objects. Instead, they use a gaze heuristic with maintains a constant optical angle between their target and themselves in order to position themselves (Gigerenzer, 2007; McLeod and Dienes, 1996). Even beyond the instinctual example of the gaze heuristic, single-cues can be useful in the business settings of predicting consumer behavior. For example, how do commercial retailers distinguish between active customers who will purchase again versus those who are inactive and will not purchase again? The hiatus heuristic states that customers who have not purchased within a certain number of months (the hiatus) are inactive, and otherwise, are active (Wbben and v. Wangenheim, 2008). The hiatus heuristic was compared to a state-of-the-art Pareto negative binomial distribution model (Pareto/NBD) which used 40 weeks of customer data of apparel, airline, and music purchases. While the hiatus heuristic only required a single value threshold from experts, it correctly classified 83%, 77%, and 77% of customers in each respective domain, compared to the Pareto/NBD model's 75%, 74%, and 77%. More generally, one-clever-cue heuristics have been shown to perform well in environments where there is high variability in relative importance of cues, high cue redundancy, and small sample sizes (Hogarth and Karelaia, 2005a, 2007; Katsikopoulos et al., 2010).

The second class of one-reason heuristics expands the one-clever-cue to sequential search through multiple cues, referred here as multiple-clever-cue heuristics. The exemplar strategy of multiple-clever-cue heuristics is take-the-best (TTB) which searches through cues in rank-order, stops when one cue discriminates, then selects the option with the positive cue score (Gigerenzer and Goldstein, 1996). For 2-option decision tasks, an individual cue discriminates when one option has a positive cue score and the other option have a negative or unknown cue score. Due to the discrimination rule, TTB will tend to stop its information search before evaluating all available information from all available cues. The ability for TTB to evaluate multi-cue decision tasks, has enabled many of the surprising results of one-clever-cue heuristics to be replicated in broader mathematical and simulation studies: high performance in environments with moderate to high variability in cue weights, moderate to high cue redundancy, and environments with dominant options (Hogarth and Karelaia, 2005a,b, 2006; Katsikopoulos et al., 2010; Katsikopoulos, 2013; Şimşek, 2013).

Human-subjects studies have expanded this perspective to show that people often use heuristics instead of analytic strategies which evaluate all available cues in each task (Gigerenzer and Gaissmaier, 2011). In consumer choice, people tend to use heuristics like TTB to eliminate most items from further consideration, identifying the final group of items to more thoroughly evaluate (Kohli and Jedidi, 2007). In residential burglary, expert groups of experienced burglars and police officers, were shown to be best predicted by TTB whereas a novice group of graduate students were shown to be best predicted by an analytic weighted-additive strategy (Garcia-Retamero and Dhami, 2009). In situations of high individual choice time pressure and opportunity costs, participants were shown to be better represented by a lexicographic strategy like TTB rather than the weighted-additive strategy (Rieskamp and Hoffrage, 1999, 2008).

Viewed through the perspective of judgment theory, one-reason decision strategies become fast-and-frugal trees (FFTs) which have also been shown to have similar accuracy to analytic models in simulations, while requiring fewer cues for categorization. FFTs are classification trees that have "$n + 1$ exits, with one exit for each of the first $n - 1$ cues and two exits for the last cue" (Luan et al., 2011, p. 320). They have become a popularly studied tool for decision support and have been shown to be quick to use, easy to remember, and accurate across many domains: medical (Green and Mehr, 1997; Fischer et al., 2002; Katsikopoulos et al., 2008; Jenny et al., 2013, 2015; Jenny, 2016), military (Keller et al., 2014; Keller and Katsikopoulos, 2016), and financial (Aikman et al., 2014). Their transparency and consistency have aided their adoption among professionals (Katsikopoulos et al., 2008).

In a comprehensive comparative study by Jenny and her colleagues (2015; 2016), leveraged the simpler cue structure of FFTs as a useful decision aid for emergency medicine doctors to categorize patients presenting with nonspecific complaints as either low or high morbidity. These patients present particular issues for emergency departments because they are difficult to accurately triage and diagnose, yet delayed decisions can result in adverse outcomes. To support better decisions, Jenny and her colleagues compared 18 statistical and machine learning algorithms, including the most complex machine learning algorithm, random forest, to the FFT (referred to as STEL) and the physicians' "looking ill" (LI) cue. In terms of area under the ROC curve, the random forest and the FFT achieved near-identical performance, 0.82 and 0.81, whereas LI achieved 0.65. In terms of required cues, however, they could not be more different. The random forest classification tree had 14 cues whereas the FFT had only 4.

2.3.2 Fewer Known Cue Values

The direct result of time pressure, high information acquisition costs, too much information, and incomplete information in the environment, is that decision makers face tasks with missing cue values and must make decisions with fewer known cue values. Despite this direct relationship, few studies have examined how accuracy is affected by incomplete information in the decision task. Those that have studied incomplete information, have only measured total information, the total amount of known cues values, largely ignoring how that information is distributed. In sum, these studies have shown that more information is better, but that heuristic strategies are more robust to environments with few known cue values than analytic strategies.

In decision tasks with binary cues, Martignon and Hoffrage (2002) showed that TTB had higher accuracy than the analytic strategy, Dawes' Rule, when the total information was low or scarce. As total information increased toward high or abundant information, the accuracy of Dawes' Rule increased while the accuracy of TTB stayed about the same. The additional information does not increase the accuracy of TTB because the decision is likely to be made during the first few cues. However, since Dawes' Rule (like most analytic strategies) allows lower-ranked cues to compensate for higher-ranked cues, the increased information enabled the strategy to compensate for any mistakes that would have been made with less information.

Payne et al. (1990) also showed that the accuracy of TTB was less affected by decreasing total information than analytic strategies in a comparison of multiple decision making strategies. Time pressure was used to stop decision making strategies in the middle of their information search which forced strategies to make decisions with missing cue values. Their computational study showed that severe time pressure reduced the accuracy of WADD by over 80%, and while LEX (similar to TTB) also experienced reduced accuracy, the reduction was only approximately 20%.

Why is TTB so robust to decision making environments with fewer known cue values? Many studies with complete information had shown the effectiveness of TTB and other one-reason heuristics as a result of fit between the environment and the strategy, particularly variability in cue weights and cue redundancy (Hogarth and Karelaia, 2005a,b, 2006; Katsikopoulos et al., 2010; Katsikopoulos, 2013; Şimşek, 2013). However, it was Garcia-Retamero and Rieskamp (2008, 2009) who showed that, for decision making tasks with incomplete information, how strategies estimate missing information is essential to achieving high accuracy in decision tasks with missing cue values. In fact, the mechanism of estimating missing information has a stronger impact on accuracy than the mechanism of integrating the known information. In general there are three potential estimates: positive (replace the missing cue value with a high cue value, assuming the missing information is positively correlated with the criterion), negative (replace the missing cue value with a low cue value, assuming the missing information is negatively correlated with the criterion), or average (replace the missing cue value with the average value of the cue, assuming the missing information is not correlated with the criterion). For both heuristic and analytic strategies, matching the assumption about the relationship between missing information and the criterion value is the key to achieving high accuracy.

In summary, TTB and other heuristics are more robust than analytic methods in environments with few known cue scores because of their reliance on only a few, important cues, and their accurate assumptions about the correlation between missing information and the environment.

2.3.3 Simpler Models with Less Training

Some researchers have argued that, while heuristics look simple on the surface, by using simple models with fewer cues, they actually require complex calculations to apply, such as ordering the cues or determining the relationship between the cues and the environment (Juslin and Persson, 2002; Newell, 2005; see Hilbig, 2010, for a broader critique). It is accepted that linear regression models and advanced machine learning algorithms are complex models requiring large amounts of training data for good performance, but how much complexity is inherent in the naturalistic decision making strategies and fast-and-frugal heuristics? The two research programs have diverse but equally important responses to this question.

Naturalistic Decision Making

The naturalistic decision making program would respond that yes, their strategies are based on experience and expertise, but the difficulty is matched to their decision makers as the strategies are simply formalizations of the cognitive processes people are already using. For example, the simple-looking questions of "will [the plan] work?" and "are the cues relevant?" within the models of recognition-primed decision making (Klein, 1993) and recognition/meta-recognition (Cohen et al., 1996) certainly belie the complexity of the answers. Expertise and experience are essential to forecasting potential results of implementation (will the plan work?) and knowing the relevance of the information provided by different cues (are the cues relevant?). However, the need for experience and expertise is not an exceptional requirement to perform the strategy effectively; they are the foundation of the strategies themselves. The entire set of naturalistic decision making strategies is founded on the understanding of "how people make decisions in real-world contexts that are meaningful and familiar to them" (p. 332, Lipshitz et al., 2001).

This does not mean that only experts with large amounts of prior data can use naturalistic decision making strategies. In fact, one of the main successes of the naturalistic decision making program has been to implement effective training programs to communicate these strategies in order to help novices perform like experts. The most striking example comes from a pair of studies which compared a standard military model of tactical decision making to a proposed naturalistic decision making model (Ross et al., 2004; Thunholm, 2003). In the standard model, three courses of actions are generated and evaluated, each are compared, then the best is selected. In the naturalistic decision making model, the first course of action generated is evaluated and used if satisfactory. If the course of action is not satisfactory, then reevaluate the situation and generate a new course of action. In tests with military commanders in the United States (Ross et al., 2004) and Sweden (Thunholm, 2003), the naturalistic decision making model was described by the commanders as more natural and comfortable than the military standard, reduced planning time by 30% and 20%, respectively, and resulted in better courses of action. Why was the naturalistic decision making strategy more effective and efficient, even though it was more ambiguous (and potentially more complex) than the military standard? Because it was structured around the inherent capabilities of experts now known to researchers: theses decision makers do not generate and compare option sets, they use prior experience to rapidly categorize situations; they rely on a synthesis of their experiences; and they do not await outcomes of gambles, they actively shape events (Klein, 2008; Rasmussen, 1983; Hammond et al., 1987).

Fast-and-Frugal Heuristics

The fast-and-frugal heuristics program would agree with the naturalistic decision making program's response about the importance of experience and expertise, but the mathematical and computational formalizations of fast-and-frugal heuristics have contributed direct results showing that heuristics are fairly robust to errors resulting from smaller training data sizes (Czerlinski et al., 1999; Hogarth and Karelaia, 2006; Martignon et al., 2008; Brighton and Gigerenzer, 2012).

Katsikopoulos et al. (2010) directly compared two variants of TTB (with continuous or binary cues) to the benchmark strategies of multiple linear regression (MR) and naïve Bayes (NB), in a range of training data sizes across 19 datasets. For every amount of training data, ranging from 3 objects to 50% of the training dataset, TTB with continuous cues outperformed NB and MR by an average of 5%. In fact, NB and MR with continuous cues performed at the same predictive accuracy as TTB with binary cues. Furthermore, with only seven objects, TTB with continuous cues achieved a performance close to the maximum achieved by any model with 50% of the training dataset.

Further analysis by Katsikopoulos et al. (2010) showed that the success of TTB was a result of the "robust beauty" of cue directions, the correlation between the cue value and the criterion. Analysis of their data showed that cue order actually did not matter that much. For small training data sizes of "two to eight objects, it [was] more likely that the three most valid cues were ranked incorrectly than at least one of them [was] ranked correctly" (p. 1261, Katsikopoulos et al., 2010). In total, the mean correlation between the accuracy of TTB and the accuracy of the cue order, only achieved around 0.20 across all training data sizes. Conversely, the mean correlation between the accuracy of TTB and the accuracy of the cue directions, was over 0.60 for 2-object training datasets decreasing linearly to approximately 0.35 for 10-object training datasets. They concluded that minute samples do not represent the true cue order but do provide assistance with cue order and the more informative cue direction.

Through comparison testing, the fast-and-frugal heuristics research program has essentially concluded that although the complexity and flexibility of analytic models lead to better fit of known data than heuristics, the models do not necessary lead to better predictions of unknown data (p. 464, Gigerenzer and Gaissmaier, 2011). Martignon et al. (2008)

showed that fast-and-frugal trees (FFTs) had similar performance in prediction tasks across 30 real-world datasets as compared to classification and regression trees (CART) and logistic regression (LR). The FFTs were shown to be 14% and 6% less accurate than CART and LR in fitting tasks but only 2% and 4% less acurate in prediction tasks with 50% of the training data. Brighton and Gigerenzer (2012) performed a similar comparison of TTB to CART, a nearest-neighbor classifier, and Quinlan's decision tree induction algorithm, across 20 datasets. Within each dataset, TTB had consistently higher accuracies than other strategies in prediction tasks with training data sizes ranging from just 5 objects to nearly the complete training dataset.

2.4 Information-Based Decision Support

Decision support systems (DSSs) and decision tools assist humans in making decisions that more closely resemble the decision one would make with complete and perfect information. There are many aspects of DSS design and evaluation (e.g. Silver, 1988, 1990; Todd and Benbasat, 1999; Feigh et al., 2012) and types of decision tools (e.g. Kleinmuntz and Schkade, 1993; Edwards and Fasolo, 2001; Katsikopoulos and Fasolo, 2006). Therefore, just as judgment and decision making research can be founded in the overcoming or embracing of bounded rationality perspectives (Sec. 2.1), so can the prescriptions for how to acquire or restrict information, referred to as 'information-based decision support' in this dissertation.

The two types of decision support, normative and heuristic,[2] mirror the two paradigms of unbounded and bounded rationality research as shown in Table 2.1. Normative decision support uses DSSs and decision tools to assist the implementation of normative decision making strategies, debiasing the decision wherever possible. Conversely, heuristic decision support uses DSSs and decision tools to extend the capabilities already present within

[2]This dichotomy is an updated and recontextualized version of the two 'schools of thought' described by (Stabell, 1987) and reviewed by (Todd and Benbasat, 1999). The dichotomy is also mirrored in the new behavioral economics and policy perspectives of nudging (Thaler and Sunstein, 2009) versus boosting (Grüne-Yanoff and Hertwig, 2016).

Table 2.1: Comparison of normative and heuristic decision support methods.

	Normative Support	**Heuristic Support**
Examples	Decision analysis: Bayes' rule, expected utility, subjective expected utility (Edwards and Fasolo, 2001); Optimal experimental design (Nelson, 2005; Meder and Nelson, 2012; Nelson et al., 2010); Nudging (Thaler and Sunstein, 2009)	Natural frequency formats; frugal multiattribute models; fast-and-frugal trees (Katsikopoulos and Martignon, 2006); Boosting (Grüne-Yanoff and Hertwig, 2016)
Pros	Used when decision makers have to justify decisions (Tetlock, 1983; Tetlock and Kim, 1987); Analyzing large datasets and decision tasks (e.g. 3000 options) (Edwards and Fasolo, 2001)	Easier to communicate, easy to apply (Elwyn et al., 2001; Fischer et al., 2002; Jenny, 2016) Usable with limited time, information, and computational resources (Katsikopoulos et al., 2008, p. 458)
Cons	Necessary information often not available (Green and Mehr, 1997; Elwyn et al., 2001); Requires reliable assessments of probabilities, attributes weights, and values which are difficult to obtain (Katsikopoulos and Fasolo, 2006)	Lack of accountability (Tetlock, 1983; Tetlock and Kim, 1987); Incomplete feedback to refine attribute rankings, weights, and judgments (Katsikopoulos et al., 2008, p. 447 and 458)

decision makers, particularly during the implementation of heuristic decision strategies. Importantly, as Katsikopoulos and Fasolo (2006) argue, these two perspectives should not be seen as competitive but as synergistic, because they are each valid for different environments and decision makers.

Normative Decision Support

Normative decision support is exemplified by the three steps decision analysts perform (Katsikopoulos and Fasolo, 2006, p. 960): "1) they help the decision maker assess the probabilities of all possible outcomes of each available option; 2) they help assess the weights of all attributes for each option; and 3) they help combine these judgments and then suggest the option with the highest overall value." Decision analysis extensively uses

Bayes' rule, expected utility theory, and subjective expected utility theory (Edwards and Fasolo, 2001).

Most relevant to this dissertation, the formal information acquisition methods for judgment and decision making have focused on normative Bayesian optimal experimental design (OED) models which provide a framework for evaluating the utility of queries (questions about which information to acquire). OED models focus on expected utility maximization (Savage, 1954) where utility is defined according to some quantification of the value of information.[3] Therefore, to identify which piece of information should be acquired, the OED methods require reliable assessments of probabilities, cue weights, and cue values.

The study of OED models have two main results. First, the most comprehensive competitive test of various OED models by Nelson et al. (2010) showed that people's choices about which questions to ask prior to categorizing an object best matched probability gain, whose utility function measures the improvement in the expected probability of making the correct choice by acquiring that information (Baron, 1985). Nelson et al. (2010) were careful to note that although they were the first to competitively test OED models, much more work needs to be done to generalize the results to other contexts such as medical diagnosis or scientific-hypothesis testing. Second, the OED models have found wide adoption in statistical modeling and designs of experiments. The measures of the value of information have been leveraged to sequentially sample a design space in order to fit a model to data, or to iteratively search for the optimum value of a design space (Jones et al., 1998).

Heuristic Decision Support

Heuristic decision support has been developed to address situations where the normative decision analysis tools are not practical. These heuristic tools are exemplified by natural frequency formats, frugal multi-attribute models, and fast-and-frugal trees. Specifically,

[3] Nelson (2005, 2008) provides a review of various OED methods while Meder and Nelson (2012, Appendix, Table A1) provide a more comprehensive table of OED methods.

each of these three tools, in order, address the psychological and environmental limitations in each of the three tools required for normative decision tools (Katsikopoulos and Fasolo, 2006, pp. 961–962): matching probability elicitation required for Bayes' theorem to better match human judgment; reducing effort required for calculating multi-attribute utility; and making quick, transparent, and accurate decisions instead of maximizing subjective expected utility, especially when the computation is intractable or the inputs cannot be obtained.

Heuristic decision making strategies (and their judgment equivalents, Katsikopoulos, 2013) are versions of information acquisition methods, but only when two conditions are satisfied: (1) the distribution of available information to be acquired is fixed and (2) the decision maker or judge is starting the task without any previously considered information. For example, take-the-best searches for the highest-ranked cue and selects the option with the highest positive score on the cue (Gigerenzer and Goldstein, 1996). If that cue does not discriminate, then the cue with the second-highest validity is evaluated to determine discrimination, and so on.

However, due to these restrictive conditions, there are no formal heuristic information acquisition methods relevant for DSS design, even given the large number of mathematical, computational, and human-subjects studies. DSSs are capable of violating the first condition because they can restrict or guide information available to a decision maker – in other words, the distribution of available information can be modified. The second condition can be violated when some information is already known to the decision maker for some reason, e.g., experience or learning. Therefore, there is a need for understanding how information acquisition and restriction can can be performed in a heuristic manner: modifying distributions of incomplete information to increase the performance of decision making strategies without the need for probabilities, cue weights, and cue values.

2.5 Summary

Synthesizing the literature related to decision theory and decision support has shown that there are open questions related to how people make decisions with incomplete information. It is clear that people, either by choice or by environmental causes, do not make decisions using all possible or available information, and often heuristics can be quite accurate in these environments. However, studies of decision making have only partially answered how the incomplete information affects decision making accuracy: focusing on total information alone, ignoring the varied distribution of information between options and between cues. As a result, the only clear method for information-based decision support (acquisition or restriction) is based on normative methods requiring accurate measures of cue weights, probabilities, and cue values. Therefore, answering the two research questions of this dissertation will address significant gaps in the decision theory and decision support related to incomplete information: 1) how do various distributions of incomplete information affect decision making performance, and 2) how can information acquisition and restriction methods be designed which do not require probabilities, cue weights, or cue values?

CHAPTER 3

GENERAL LINEAR MODEL OF JUDGMENT AND DECISION MAKING

Judgment and decision making are two separate perspectives within the paradigm of bounded rationality. They are "motivated by different questions, forming two different 'foci' of the emerging field... [proceeding] rather independently until relatively recently" (Goldstein and Hogarth, 1997, p. 4). As defined by the research community (Table 3.1), judgment and decision making have related but distinctly different research interests. Judgment is the process of perceiving and understanding the environment (or the state of the world) whereas decision making is the process of choosing a course of action from a set of options. One of the most significant open questions within the domain of judgment and decision making research is how to integrate the two (Katsikopoulos, 2013).

A survey of developments in judgment and decision making research shows that this integration has already begun, but remains unfinished. In a review of the literature, Katsikopoulos (2011) showed that lexicographic strategies (e.g. Take-the-Best, TTB) and equal-weight strategies (e.g. Tallying) can be used for both characterizing a single option (judgment) and choosing between options (decision making). For example, Tallying has been referred to as both a judgment (Gigerenzer and Gaissmaier, 2011) and decision making strategy (Payne et al., 1990) in various publications. The two heuristic strategies, TTB and Tallying, have even been formally transformed into categorization trees and referred to as fast-and-frugal trees (Martignon et al., 2008).

This chapter presents a general linear model that further conceptually and mathematically integrates judgment and decision making to make the argument that, from the perspective of an information process model, they are equivalent. Both judgment and decision making require descriptions of cues that describe options (cue values) which can be known or unknown (incomplete information) and therefore missing cue values must be estimated

Table 3.1: Definitions of judgment and decision making.

	Judgment	Decision Making
Similar terms	Classification, categorization, generalization, perception	Preferential choice, inference, judgment[c]
Research Questions[a]	How do people integrate multiple, probabilistic, potentially conflicting cues to arrive at an *understanding* of a situation, a *judgment*? How accurate are peoples judgments? Does judgment improve with training and experience? How does human judgment compare with actuarial prediction? How do people identify relevant cues and the proper weights to assign to them? How does the nature of the task environment affect learning and performance?	How do people decide on a course of *action*? How do people *choose* what to do next, especially in the face of uncertain consequences and conflicting goals? Do people make these decisions rationally? If not, by what psychological processes do people make decisions? Can decision making be improved?
Normative Measures of Performance[b]	Linear and logistic regression, classification and regression trees, and Bayesian networks	Linear utility function

[a] Goldstein and Hogarth (1997, p. 4)
[b] Katsikopoulos (2013, p. 337)
[c] The language of judgment and decision making research can sometimes be confusing when presented as single names without definitions. Katsikopoulos (2011) defines judgment as the answer to: is one option's criterion value higher than the other? He then defines categorization as the answer to: does a single option belong to one category or another? Therefore, what he refers to as judgment and categorization, I refer to as decision making and judgment, respectively. The confusion hints at the potential for equivalency between the two perspectives.

(estimates of missing information). The cue values are integrated based the sign of the correlation between the cues and the criteria (cue direction) and the relative importance of the cues (cue weights). The only difference between judgment and decision making within this model is that judgment examines one option (or aspect of the environment) at a time and outputs a single value describing that option whereas decision making compares judgments of multiple options and outputs a choice of one option.[1]

[1]This perspective follows the work of Hogarth and Karelaia who used essentially the same probabilistic

The general linear model can represent almost any combination of components for both judgment and decision making: cue weights, utility functions, and estimates of missing information, for m-options, n-cues, and any distribution of incomplete information. Section 3.1 describes the general linear model and all of its variables and components – expanding and reformulating the first publication of the model (Canellas and Feigh, 2016b). Section 3.2 shows how specific variants of judgment and decision making strategies can be represented using the general linear model. The section concludes the chapter by showing how the algebraic general linear model can be formulated as a matrix equation form for fast computational simulations along with a validation of two models of fast-and-frugal trees.

The contribution of this chapter is to provide a single, unifying framework for representing, modeling, and simulating a wide range of judgment and decision making strategies. This framework is important because it provides three major benefits to the research community. First, the general linear model allows for specificity in model selection and description. It is common to refer to strategies using proper names (Take-the-Best, Tallying, etc.) which can leave ambiguity in how to appropriately represent and model the strategy, and limit intuition as to how two strategies may or may not be different. This results in further ambiguity as to whether various studies are even comparable to each other. Second, for the engineers, designers, and decision analysts, this model enables the simulation of a wide range of judgment and decision making strategies quickly and easily – and transparently – in a single equation. Rather than implementing an algorithmic tool or specialty code, it is an equation that can be looked-up and implemented. It can be used to perform sensitivity studies on the robustness of overall designs or the interaction of components across ranges of judgment and decision making strategies. Third, the simple, representative form can link together human-subjects, computational, and mathematical studies completed by various research programs such as fast-and-frugal heuristics and naturalistic decision making (Canellas and Feigh, 2016b).

models to study binary choice (2005a) and judgment between to two states (2007).

3.1 The General Linear Model

The general linear model of judgment and decision making is:[2]

$$C_i = \sum_{j=1}^{n} w_j \cdot U_j\Big(e_j + (a^v_{i,j} - e_j)z_{i,j}\Big) \tag{3.1}$$

Judgments and decisions are made within an environment which is represented by an m-by-n matrix, A, of options and cues, respectively. Each entry of A is a cue value ($a^v_{i,j}$) representing the state of option $i \,\epsilon \{1, ..., m\}$ with respect to cue $j \,\epsilon \{1, ..., n\}$ in the cue's natural or original sensed units. The first distinction of the general linear model is made between whether the task is to make a judgment or a decision. The goal of judgment is to evaluate information about a single option ($i = 1$) characterized by n-cues to categorize the option using the criterion, C. The goal of decision making is to evaluate information about multiple options ($i \geq 2$) characterized by n-cues to select the option that maximizes the criterion, C.

Individual cue values can be known ($z_{i,j} = 1$) or unknown ($z_{i,j} = 0$). If known, the operator uses the cue value ($a^v_{i,j}$). If unknown, the operator must estimate the missing cue values (e_j). The cue values (either known or estimated) are then converted to cue scores ($a^s_{i,j}$, which have units useful for the formal judgment and decision making process) through utility functions ($U_j : (a^v_{i,j},\ e_j) \rightarrow a^s_{i,j}$). In some cases, the utility functions convert the cue values to binary cue scores through cue directions (d_j) and comparing cue values to cutoff values (c_j) with thresholds (Δ_j). Once the cue scores have been determined, they are combined with cue weights (w_j) to determine the option's criterion, C_i.

The general form in Eq. 3.1 can model almost any combination of the components for

[2] An alternate presentation of Eq. 3.1 uses the notation of Katsikopoulos et al. (2014), where a general option is referred to as A instead of through a subscript i:

$$C_A = \sum_{j=1}^{n} w_j \cdot U_j\Big(e_j + (a^v_j(A) - e_j)z_j(A)\Big)$$

both judgment and decision making shown in Table 3.2: cue weights, estimates of missing information, and utility functions, for m-options, n-cues and any distribution of incomplete information. The subscripts of the variables in Eq. 3.1 are particularly important. The cue weights (w_j), estimates of missing information (e_j), and utility functions (U_j) are defined with respect to cues only, whereas there is a unique information state ($z_{i,j}$) for each cue value ($a^v_{i,j}$).

The general linear model of judgment and decision making builds from the foundational and ubiquitous mathematical model of judgment and decision making: the weighted utility function (also referred to as the multi-attribute utility fuction, Raiffa, 1968). As shown in Eq. 3.2, the weighted utility function calculates the criterion of a specific option, i, by multiplying each cue's weight, w_j, by the utility of each cue value, $U_j(a^v_{i,j})$. This form generally assumes that all cue values are known to the decision maker and leaves the utility function undefined.

$$C_i = \sum_{j=1}^{n} w_j \cdot U_j(a^v_{i,j}) \tag{3.2}$$

Table 3.2: Definitions of the parameters of the general linear model and the alternative names used in judgment and decision making literature.

Component (Alternate Name)	**Parameter**	**Definition**
Goal		Objective of decision making
Criterion (option scores)	C	Utility of the option with respect to the goal
Options (Alternatives, Objects)	i	Separate, discrete courses of action or states of the world
Cues (Attributes, Predictors)	j	Descriptive dimensions of the options
Cue values	$a^v_{i,j}$	State of the option with respect to the cue in its natural units
Cues scores	$a^s_{i,j}$	State of the option with respect to the cue in units useful to the decision maker
Cue weights (Validity)	w_j	Relative quality of the cue to predicting the criterion
Utility function	U_j	Converts the cue values to cue scores.
Information state	$z_{i,j}$	Indicates whether the corresponding cue value is known (1) or unknown (0)
Estimate of missing information	e_j	The estimated cue value when the cue value is unknown
Cue direction (Polarity)	d_j	Sign of the correlation between the cue scores and criterion
Cutoff value	c	The value of a cue used by utility functions to covert cue values to binary cue scores.
Thresholds	Δ	Defines what difference between two cue values is large enough to be meaningfully different.

To account for the possibility of unknown information, the argument of the utility function is modified in Eq. 3.3 to include 1) a binary variable, $z_{i,j}$, accounting for whether the corresponding cue value is known ($z_{i,j} = 1$) or unknown ($z_{i,j} = 0$), and 2) a variable that specifies the cue value estimated by the decision maker when actual cue value is uknown,

e_j (see Eqs. 3.4 and 3.5, respectively).

$$C_i = \sum_{j=1}^{n} w_j \cdot U_j\Big(e_j + (a_{i,j}^v - e_j)z_{i,j}\Big) \tag{3.3}$$

$$\begin{aligned} U_j(z_{i,j} = 1) &= U_j(a_{i,j}^v) && (3.4) \\ U_j(z_{i,j} = 0) &= U_j(e_j) && (3.5) \end{aligned}$$

To account for the two most common variants of utility fuctions, the general linear model in Eq. 3.3 is split into two forms. First, an 'exact' utility function assumes that the cue values are already in usable form with the correct magnitudes and cue directions ($U_j(x) = x$, Eq. 3.6). The 'binary' utility function in Eq. 3.7 converts the cue values to binary cue scores represented by 1 and 0 (e.g. yes and no, or positive and negative). The binary utility function accounts for: 1) cue directions (d), the sign of correlation between the cue values and the criterion, 2) the cutoff values (c) that distinguish between the two binary states, and 3) the threshold values (Δ), the minimum difference between two cue values to be meaningfully different.

$$\begin{aligned} C_i &= \sum_{j=1}^{n} w_j \cdot \Big(e_j + (a_{i,j}^v - e_j)z_{i,j}\Big) && (3.6) \\ C_i &= \sum_{j=1}^{n} w_j \cdot H\Big[d_j\Big(e_j + (a_{i,j}^v - e_j)z_{i,j}\Big) - (d_j \cdot c) + \Delta_j\Big] && (3.7) \end{aligned}$$

The following subsections provide more detail about the cue weights, the estimates of missing information, and the utility functions.

Table 3.3: Mathematical representations of the components of the general linear model.

Type	**Mathematical Representation**
Cue weights: the relative importance of cues (w_j)	
Magnitude, ordered	v_j
Ordered, non-compensatory	$w_j > \sum_{k>j} w_k$
None or equal	1
Estimates: Estimating missing information (e_j)	
Positive	$\max a_j^s$
Average	$\bar{a}_j^s$
Median	$\tilde{a}_j^s$
Negative	$\min a_j^s$
Utility function: maps the state of the option in its natural units to a general utility (U_j)	
Exact	1
Binary, pre-specified cutoff, positive cue direction	$H[a_{i,j}^v - (c_j - \Delta)]$
Binary, pre-specified cutoff, negative cue direction	$H[-a_{i,j}^v + (c_j + \Delta)]$
Binary, relative cutoff, positive cue direction	$H[a_{i,j}^v - (\max a_j^v - \Delta)]$
Binary, relative cutoff, negative cue direction	$H[-a_{i,j}^v + (\min a_j^v + \Delta)]$

3.1.1 Cue Weights

The cue weights, w_j, within a task measure the relative importance of each cue for predicting the criterion. There are three aspects of cue weights (magnitude, order, and compensatoriness) and different judgment and decision making strategies use and ignore different aspects. Magnitude is the numerical value of the importance of the cue (v_j), typically on a ratio scale of 0 to 1, although other scales are possible. There are many methods for getting the magnitude of cue weights, from eliciting weights from human-subjects (e.g., Edwards and Fasolo, 2001) to calculating weights from information about the cue values and criterion scores (e.g., regression weights, ecological validity, success validity, or conditional validity, as described in Martignon and Hoffrage, 2002). Ordered cues are ranked and evaluated from most to least important. Whereas order can be determined from magnitude, some strategies are only concerned with order, not magnitude. Compensatoriness is the measure of whether a high score on a lower-ranked cue can compensate for a low score on a higher-ranked cue (Martignon and Hoffrage, 2002; Hogarth and Karelaia, 2006). If cues can compensate, the weights are defined as compensatory whereas if cues cannot compensate, the weights are defined as non-compensatory.

3.1.2 Estimates of Missing Information

Garcia-Retamero and Rieskamp (2008, 2009) suggest that there are four ways that individuals estimate missing information: positive ($\max a_j^s$), negative ($\min a_j^s$), using the average of past observations for replacement ($\bar{a}_j^s$), or using the most frequent observation of the available information as a placeholder ($\tilde{a}_j^s$). Figure 3.1 shows how the utility function maps cue values to cue scores, but for scores that are unknown, assumes some standard cue score. For strategies that use binary utility functions, positive and negative estimates are typically represented as 1 and 0, respectively, because those are the upper and lower bounds of the possible cue scores ($a_{i,j}^s$). Furthermore, to represent the use of binary cues with three states of cue scores, $\{0, 0.5, 1\}$, the average or median estimate (which is preferred to 0, but less preferred than 1) should be set equivalent to the cutoff value, $e_j = c_j$. Appendix A.1.1 describes how a specific form of non-compensatory cue weights ($w_j = 4^{1-j}$) allows strategies to use binary utility functions with three-state cue scores $\{0, .5, 1\}$.

3.1.3 Utility Function

Though the general form in Eq. 3.1 allows any utility function to be used, commonly studied strategies typically have one of two forms of the utility function: exact or binary. Exact utility functions use cue values as the cue scores because the cue values were already in units useful to the operator and the cue directions are already accurate in relationship to the cue weights ($U_j(x) = x$). Binary utility functions transform the cue values into binary cue scores such as 1 or 0, yes or no, and positive or negative.[3]

The binary utility functions use the Heaviside function, $H(x)$ in Eq. 3.8, such that if

[3] Binary cue scores are commonly used within fast-and-frugal heuristics framework and heuristics have been shown to perform particularly well in binary cue environments (Hogarth and Karelaia, 2005a,b; Katsikopoulos, 2013).

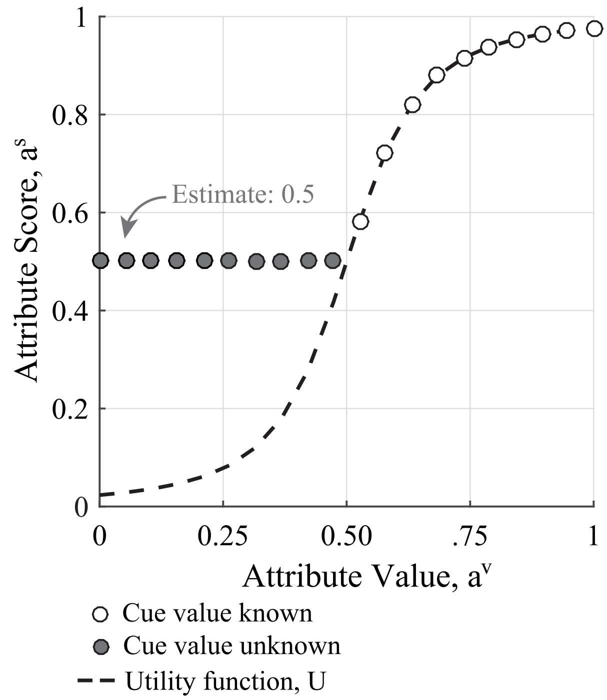

Figure 3.1: Example of how missing cue score information results in estimated cue scores. The estimate of missing information, e, is set to 0.5.

the argument, x, is greater than 0, the result is 1; if the argument is less than 0 the result is 0; and, if the argument is equal to 0, the result is 0.5. By modifying the argument as shown in Eq. 3.9, and simplified in Eq. 3.10, the Heaviside function can account for 1) the sign of correlation between the criterion and the cue (d_j, referred to as cue directions, Katsikopoulos et al., 2010, or cue polarity, von Helversen et al., 2013); 2) the cutoff value (c) that distinguishes between the two binary states; and, 3) the threshold applied to determine if the difference between two options is large enough to tell them apart meaningfully (Δ_j, Luan et al., 2014; also referred to as just noticeable differences, JND, Payne et al., 1990). A visualization of this binary utility function in Fig. 3.2 shows how the three parameters

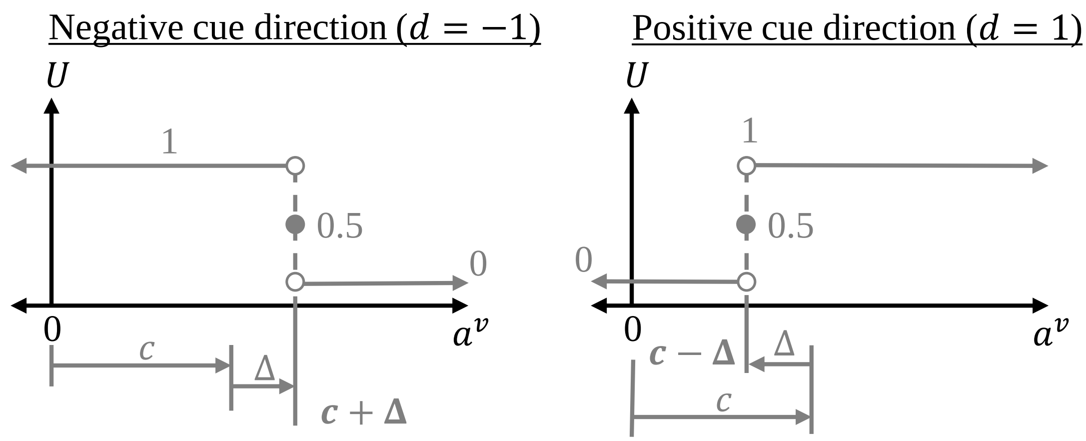

Figure 3.2: Binary utility functions for negative and positive cue directions accounting for cue values (a^v), cutoff value (c), and just noticeable differences (Δ).

affect the relationship between the cue values and their utility (or cue score).

$$H[x - c] = \begin{cases} 0, & \text{for } x < c \\ 0.5, & \text{for } x = c \\ 1, & \text{for } x > c \end{cases} \tag{3.8}$$

$$U_j(x) = H\Big[d_j x - d_j(c - d_j \Delta_j)\Big] \tag{3.9}$$

$$U_j(x) = H\Big[d_j x - d_j c + \Delta_j\Big] \tag{3.10}$$

The cue directions split the binary utility function into two equations based on whether the cue values positively correlate with the criterion score $(d = 1)$ or whether the cue values negatively correlate with the criterion score $(d = -1)$, as shown in Table 3.4. Knowledge of cue directions represents familiarity or experience with a task, such as a driver knowing that increased traffic tends to correlate with increased time to destination – a positive cue direction. In a simulation study, Katsikopoulos et al. (2010) showed that the ordinary information of correct cue directions in small samples can increase the accuracy of TTB more

Table 3.4: The four variants of the binary utility function along two dimensions: positive or negative cue directions, and prior or relative cutoff values.

	Prior ($c = c_j$)	Relative ($\max \setminus \min(a_j^v)$)
Positive ($d = 1$)	$H[a_{i,j}^v - (c_j - \Delta_j)]$	$H[a_{i,j}^v - (\max a_j^v - \Delta_j)]$
Negative ($d = -1$)	$H[-a_{i,j}^v + (c_j + \Delta_j)]$	$H[-a_{i,j}^v + (\min a_j^v + \Delta_j)]$

than having the correct rank-order of cues, even though TTB assumes the cues are searched in rank-order. In a human-subjects study of judgment and decision making, von Helversen et al. (2013) showed that participants with knowledge of cue directions mastered an environment faster (fewer learning tasks required, better learning performance, and better task performance) than participants without knowledge of the cue directions. Given these recent results, it is, perhaps, no surprise that cue directions can even be used as the foundation for a linear model for judgment and decision making by setting the weights equivalent to the cue directions ($w_j = 1$ or -1). These unit-weight models (e.g., Dawes's Rule and Tallying) have been shown to be efficient and effective in many environments (Gigerenzer et al., 1999; Martignon and Hoffrage, 2002).

The cutoff value, c, can be specified prior to the decision task or relative to other cue values in the decision task. When defining a cutoff value prior to the decision task (c_j), the value should be set by an analysis of the environment. However, there are methods for analytically defining the cutoff value from cue value data (Slegers et al., 2000): the median-split model which uses the median of the cue values ($\tilde{a}_j^v$, and is by far the most commonly used method), and the δ model which calculates the cutoff value by algorithmically determining when a cue discriminates.

Alternatively, the cutoff value can be set relative to the other cue values being considered so that cues values are mapped to a cue score of 1 only when they are the best, either the maximum or minimum cue values depending on the cue direction. These relative-cutoff binary utility functions have been largely studied with respect to lexicographic strategies (e.g., Tversky, 1969; Czerlinski et al., 1999; Katsikopoulos et al., 2010; Luan et al., 2014).

One of the most interesting results showed that, across 19 environments, TTB with relative cutoff values significantly outperformed TTB with prior-specified (median-split) cutoff values, and even multiple regression and naïve Bayes models with continuous cues (Katsikopoulos, 2010).

The threshold value, Δ, is used to describe what difference between two cue values is large enough to tell them apart meaningfully. For example, within a driving task, a difference in distance of one mile or less between two options may not be meaningful ($\Delta = 1$), or a difference of five minutes in travel time may not be meaningful ($\Delta = 5$). A threshold value is a fundamental part of satisficing strategies known as lexicographic semiorders, which search in rank order and only stop to make a decision when the difference between the options exceeds the threshold (e.g., Tversky, 1969; Bettman et al., 1990; Kohli and Jedidi, 2007; Luan et al., 2014). The most common threshold value used in simulation studies has been zero – often because when ignored in modeling and simulation, the threshold is effectively zero. But in the first systematic analysis of thresholds, across 39 real-world task environments, Luan et al. (2014) showed that a threshold of zero actually resulted in the best inferences for relative-cutoff, binary utility strategies. The implication of a zero threshold in relative cutoff strategies is that the strategies truly follow the lexicographic strategy generally described as take the best and ignore the rest.

3.1.4 Categorizing or Choosing

While the calculation of the criterion makes no distinction between judgment and decision making, the two task types diverge from each other once the criterion is known. Decision making strategies calculate the criterion of two or more options ($i \geq 2$) and then select the option with the maximum criterion. The first version of the decision making goal function, G_{DM} in Eq. 3.12, is the simplest form: identify the option i which maximizes the criterion, C_i. The second version of the decision making goal function accounts for the potential requirement that in order to distinguish between two options, a threshold, Δ', must be exceeded (e.g. for lexicographic semi-orders described in Sec. 3.1.3). If the goal function returns multiple elements, typically one option is selected randomly from within the set of elements or the threshold is disregarded such that the option with the absolute maximum criterion is selected.

$$G_{DM} = argmax_i(C_i) \tag{3.11}$$

$$G_{DM} = argmax_i\Big(H\big[C_i - \big(\max(C_i) - \Delta'\big)\big]\Big) \tag{3.12}$$

Judgment strategies, on the other hand, use the goal function, G, to covert the criterion of a single option ($i = 1$) to a category that describes the state of the environment. For judgment, G is related to the utility functions described in Sec. 3.1.3. If an exact utility function is used, there is no general method for mapping the criterion value to a category and possibly the criterion can be used as the value itself. However, for binary utility functions, G can leverage an understanding of the cue weights to assign a unique category to each unique criterion value. The Heaviside function causes the criterion for any option to be a sum of the cue weights for two-state cue scores $\{0, 1\}$ and a sum of the cue weights and halves of cue weights for three-state cue scores $\{0, 0.5, 1\}$. Therefore, the number of unique criterion values generated by the judgment strategy is equivalent to the maximum possible categories.

Two major types of goal functions can be specified based on the type of cue weights for binary utility functions. Judgment strategies with non-compensatory cue weights can categorize every combination of cue scores for both two-state and three-state cue scores, 2^n and 3^n categories, respectively. Judgment strategies with equal cue weights can categorize $n + 1$ and $2n + 1$ for two-state and three-state cue scores, respectively.

Figure 3.3 provides categorization trees for non-compensatory cue weights with two cues and shows how there are unique criterion values and therefore, unique categories, for

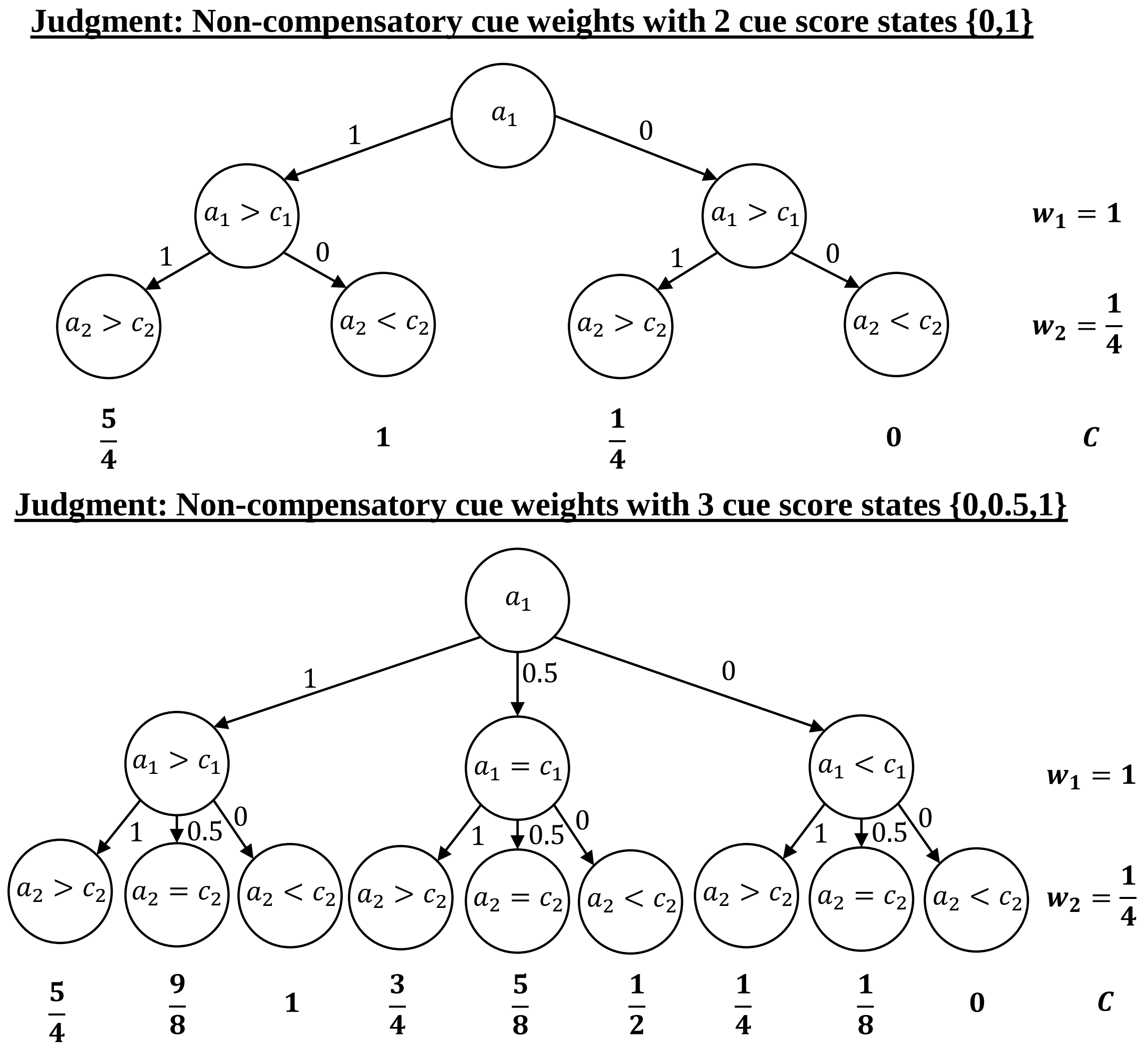

Figure 3.3: Categorization tree representation of the general linear model for judgment with non-compensatory cue weights and either two cue score states $\{0, 1\}$ or three cue score states $\{0, 0.5, 1\}$.

each combination of cue scores. The exponential relationship between cues and categories is a result of the geometric series used to calculate cue weights. Every cue weight is greater than the sum of all lower-ranked cues (except at the limit of infinity). Therefore, as the categorization progresses down through each cue, the potential criterion values is incrementally limited to non-overlapping ranges. The count of unique categories possible for both non-compensatory and equal cue weights for given number of cues are provided in Table 3.5 and a complete description of categories for three cues is provided in Appendix A.1.2.

Table 3.5: Count of unique categories possible for both non-compensatory and equal cue weights, for a given number of cues.

Cues (n)	2-state binary function: $\{0,1\}$ Non-Compensatory	Equal	3-state binary function: $\{0,0.5,1\}$ Non-Compensatory	Equal
$\mathbf{n}$	$\mathbf{2^n}$	$\mathbf{n+1}$	$\mathbf{3^n}$	$\mathbf{2n+1}$
1	2	2	3	3
2	4	3	9	5
3	8	4	27	7
4	16	5	81	9
5	32	6	243	11
6	64	7	729	13
7	128	8	2187	15
8	256	9	6561	17
9	512	10	19 683	19
10	1024	11	59 049	21

3.2 Modeling and Simulation

3.2.1 Modeling Decision Making

Given the multitude of decision making strategies that have been studied, there is typically confusion as to what ways various strategies are the same or different (e.g. Tallying, Katsikopoulos et al., 2010, versus Equal-Weighting, Payne et al., 1988). In other cases, when reading it can be difficult to know exactly how to model the strategy being described as authors may not specify what estimates or thresholds are being used and just reference a strategy by name, e.g. Take-the-Best. The general linear model provided in Sec. 3.1 has 4 parameters that, when defined, can fully specify a judgment or decision making strategy in addition to the goal function.[4] Table 3.6 contains the specific parameters for six well-studied decision making strategies: Take-the-Best (TTB), Minimalist, Tallying (TAL), Weighted-Additive (WADD), Multiple Regression (MR), and Equal-Weighting (EW). However, fol-

[4]The cue values, (a^v), and incomplete information (z) are given by the environment. The cue directions can either be assumed to be correct or generated from a sample of the environment (Katsikopoulos et al., 2010), but in either case are not a formal part of specifying a strategy. Similarly, the threshold value (Δ) is typically not specified and assumed to be zero (Luan et al., 2014).

lowing the footnotes of the table shows that there is still ambiguity in the definitions. The following equations are the mathematical form of each of the strategies:

$$C_{TTB} = \sum_{j=1}^{n} 4^{1-j} \cdot H\Big[d_j \cdot a^v_{i,j} \cdot z_{i,j} - d_j \cdot \tilde{a}_j + \Delta_j\Big] \tag{3.13}$$

$$C_{MIN} = \sum_{j=1}^{n} 4^{1-rand(j)} \cdot H\Big[d_j \cdot a^v_{i,j} \cdot z_{i,j} - d_j \cdot \tilde{a}_j + \Delta_j\Big] \tag{3.14}$$

$$C_{TAL} = \sum_{j=1}^{n} H\Big[d_j \cdot a^v_{i,j} \cdot z_{i,j} - d_j \cdot \tilde{a}_j + \Delta_j\Big] \tag{3.15}$$

$$C_{WADD} = \sum_{j=1}^{n} v_j \cdot H\Big[d_j\Big(\bar{a}_j + (a^v_{i,j} - \bar{a}_j)z_{i,j}\Big) - d_j \cdot \tilde{a}_j + \Delta_j\Big] \tag{3.16}$$

$$C_{MR} = \sum_{j=1}^{n} \beta_j \cdot \Big(\bar{a}_j + (a^v_{i,j} - \bar{a}_j)z_{i,j}\Big) \tag{3.17}$$

$$C_{EW} = \sum_{j=1}^{n} \Big(\bar{a}_j + (a^v_{i,j} - \bar{a}_j)z_{i,j}\Big) \tag{3.18}$$

3.2.2 Modeling Judgment

Through the use of specific cue weights and goal functions, the general linear model is a linear classifier capable of calculating a score, or criterion, based on many components (cue weights, utility functions, incomplete information, estimates of missing information, cue directions, thresholds, and cutoff values). Specifically, the general linear model can represent classification trees with up to 3 cue values (or edges) per cue (or node) as indicated in Sec. 3.1.4. While this general linear model is not a method for *generating* classification trees, it is capable of *representing* a vast amount of classification trees. Therefore, if a judgment problem can be formulated as a classification tree, then this general linear model can likely handle it.

These types of judgments do not explicitly account for the more complex methods of judgments as described in domains such as naturalistic decision making theory (Klein,

Table 3.6: Parametrization of well-studied judgment and decision making strategies with respect to the general linear model.

Strategy	**Take-the-Best(TTB)**[a], **Determ. Elim. by Aspects (DEBA)**[b]	**Minimalist**[c]	**Tallying (TAL)** [d]
Utility function (U)	Binary	Binary	Binary
Cue weights (w)	Non-CF, ordered[e]: 4^{1-j}	Non-CF, random: 4^{1-j}	Equal: 1
Estimates[g] (e)	Negative: 0	Negative: 0	Negative: 0
Cutoff values (c)	Prior-median[f]: $\tilde{a}_j$	Prior-median[f]: $\tilde{a}_j$	Prior-median[f]: $\tilde{a}_j$

Strategy	**Weighted-Additive (WADD)**[h]	**Multiple Regression (MR)**[i]	**Equal Weighting (EW)**[d]
Utility function (U)	Binary	Exact[a]	Exact
Cue weights (w)	Ecological validity[e]: v	Regression	Equal: 1
Estimates[g] (e)	Average: $\bar{a}_j$	Average: $\bar{a}_j$	Average: $\bar{a}_j$
Cutoff values (c)	Prior-median[f]: $\tilde{a}_j$	N/A	Prior-median[f]: $\tilde{a}_j$

[a] TTB was introduced by Gigerenzer and Goldstein (1996) and has been one of the most commonly studied strategies (e.g. Katsikopoulos et al., 2010; Mata et al., 2007; Dieckmann and Rieskamp, 2007; Hogarth and Karelaia, 2006; Martignon and Hoffrage, 2002).

[b] Deterministic Elimination by Aspects (DEBA) was introduced by Hogarth and Karelaia (2005b) and is equivalent to TTB (Katsikopoulos, 2013).

[c] Minimalist was introduced by Gigerenzer and Goldstein (1996) as a form of TTB in which the cues are ordered randomly (e.g. Katsikopoulos et al., 2010; Katsikopoulos and Martignon, 2006; Hertwig and Todd, 2003; Martignon and Hoffrage, 2002).

[d] Introduced by Dawes and Corrigan (1974) and Dawes (1979) as a single idea that cue weights could be equal and effective. TAL and EW have evolved to be sometimes interchangeable and sometimes not (e.g. Canellas et al., 2015; Katsikopoulos and Martignon, 2006; Katsikopoulos et al., 2010; Gigerenzer and Goldstein, 1996; Payne et al., 1990). In this table are split by their use of binary or exact utility functions. Exact still typically converts all cue scores to the same range to make the equal-weighting meaningful.

[e] Ecological validity is commonly used for ordering TTB cues and as the cue weights for WADD weighted additive, though others can be used (Martignon and Hoffrage, 2002).

[f] Medians are the standard cutoff value for binary functions – sometimes referred to as 'dichotomized by median-split' (e.g. Gigerenzer and Goldstein, 1996; Czerlinski et al., 1999; Hogarth and Karelaia, 2006) but relative cutoff values have also been used (Katsikopoulos et al., 2010).

[g] Garcia-Retamero and Rieskamp (2008, 2009) examined both WADD and TTB with all four variants of estimates of missing information (positive, negative, average, and ignore).

[i] Katsikopoulos (2010) used multiple regresssion with both binary and exact (continuous) cue scores.

Table 3.7: List of matrices for the matrix form of the general linear model of judgment and decision making.

Variables	Elements	Rows	Columns	Description
A	$a_{i,j}^{v}$	m	n	Cue values
W	w_j	1	n	Cue weights
E	e_j	1	n	Estimates of missing information
Z	$z_{i,j}$	m	n	Information states
D	d_j	1	n	Cue directions
C	c_j	1	n	Cutoff values
Δ	Δ_j	1	n	Threshold values
$\mathcal{X}$	1	m	1	Transformation matrix

2008). For example, the well-regarded recognition-primed decision making model in which options are generated sequentially, then judged as to their likely effectiveness, will require the integration of the judgment and decision making models into a cascading decision model – an important area of future work described in Sec. 10.2.2.

3.2.3 Simulating Judgment and Decision Making

The major benefit of the algebraic form of the general linear model is that it provides transparency with regards to what strategy and environment is being modeled. However, just as important, if not more important, is that the algebraic form can be converted into matrices so that the criterion scores can be calculated for any scale of m-options and n-cues using matrix multiplication. Table 3.7 describes the seven matrices of parameters that characterize judgment and decision making strategies. The eighth matrix, $\mathcal{X}$, is used solely to transform the matrices of size $1 \times n$ into matrices of size $m \times n$.

Two matrix equations represent the general linear model: Eq. 3.19 for strategies with exact utility functions from Eq. 3.6, and Eq. 3.20 for strategies with binary utility functions from Eq. 3.7. Two notes about the mathematical notation: 1) the $\odot$-symbol represents the Hadamard product which is the element-wise multiplication of two matrices, and 2) the H in Eq. 3.20 represents the Heaviside function as described in Sec. 3.1.3.

$$C = [\mathcal{X}E + (A - \mathcal{X}E) \odot Z]W^T \tag{3.19}$$

$$C = \Big[H\Big[\mathcal{X}D \odot \big(\mathcal{X}E + (A - \mathcal{X}E) \circ Z\big) - \mathcal{X}(D \odot C) + \mathcal{X}\Delta\Big]\Big]W^T \tag{3.20}$$

3.2.4 Example: Simulating Fast-and-Frugal Trees

To show the ability of the general linear model to model and simulate strategies, this subsection simulates two fast-and-frugal trees (FFTs, a type of judgment and classification tree) and shows that the statistical results are equivalent to an online benchmarking tool, the Adaptive Toolbox Online.[5] FFTs were selected as a representative example for three reasons. First, the Adaptive Toolbox Online is a freely available tool maintained by the Center for Adaptive Behavior and Cognition who helped popularize FFTs. Therefore, readers themselves can validate the results of the general linear model. Secondly, FFTs have become a popularly studied tool for decision support (e.g. Jenny et al., 2013; Luan et al., 2011). These simple classification trees have been shown to be quick to use, easy to remember, and accurate across many domains: medical (Green and Mehr, 1997; Fischer et al., 2002; Katsikopoulos et al., 2008), military (Keller and Katsikopoulos, 2016), and financial (Aikman et al., 2014). Lastly, as of this dissertation, there has been no general mathematical model of FFTs[6] – so this is the first presentation of a complete mathematical form for these increasingly popular decision support tools.

Formally, an FFT is "a decision tree that has $n + 1$ exits, with one exit for each of the first $n - 1$ cues and two exits for the last cue" (Luan et al., 2011, p. 320). The two exemplar

[5] The Adaptive Toolbox Online is maintained by the Center for Adaptive Behavior and Cognition (ABC) at the Max Planck Institute for Human Development. The website (http://www.dotwebresearch.net/AdaptiveToolboxOnline) provides a graphical user interface to easily build fast-and-frugal trees and analyze their performance.

[6] Appendix A.1.3 has a description of the the two-parameter formalization of FFT's by Martignon et al. (2008). Their formalization is limited in that they do not allow cue directions and categorizations to be defined independently, restrict the number of categories to two, and do not provide a way to account for incomplete information.

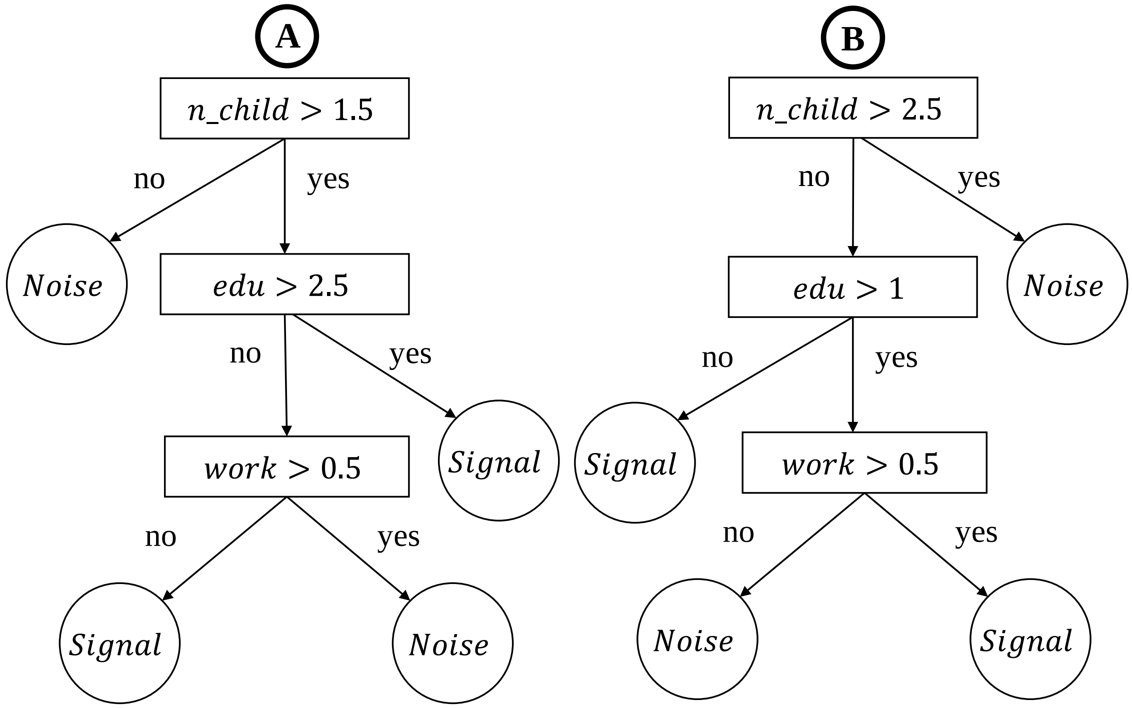

Figure 3.4: Two exemplar fast-and-frugal trees for categorization.

FFTs that will be simulated are provided in Fig. 3.4. The task is to categorize a sample of 1,473 cases of married women as either using (1, or signal) or not using (0, or noise) contraceptives by examining three cues: the number of children ever born (numerical, from 0 to 12), the wife's education (categorical, from 1-low to 4-high), and whether the wife was now working (binary, 0-yes, 1-no).[7] The cues are searched in the same order in both FFTs (A and B), but other aspects change: the cutoff values, whether no's or yes's lead to exits, whether no's or yes's lead to noise or signal. The general linear model can fully specify both FFTs as shown in Table 3.8 with no modifications or assumptions about the relationship between the cue scores and the cue directions.

Both FFT's were modeled using Adaptive Toolbox Online and using the general linear model to show that their performances were equivalent (as shown in Table 3.9); thus validating a significant portion of the general linear model. FFT performance is characterized using measures from signal detection theory (Luan et al., 2011; Martignon et al., 2008;

[7]More information about the dataset is provided in Appendix A.1.4.

Table 3.8: Parameter values for two exemplar fast-and-frugal trees.

Variables	**Example A**	**Example B**
A	[n_child, edu, work]	[n_child, edu, work]
W	[1, 0.25, 0.0625]	[1, 0.25, 0.0625]
E	[1, 1, 1]	[1, 1, 1]
Z	[1, 1, 1; 1, 1, 1]	[1, 1, 1; 1, 1, 1]
$\mathcal{X}$	[1; 1; 1]	[1; 1; 1]
D	[-1, 1, -1]	[1, -1, 1]
C	[1.5, 2.5, 0.5]	[2.5, 1, 0.5]
Δ	[0, 0, 0]	[0, 0, 0]
S	[0, 1, 1]	[0, 1, 1]

Table 3.9: Performance measures for both example FFT's showing identical results between the online tool (ATO) and the general linear model (GLM).

	Example A		**Example B**	
Statistics	**ATO**	**GLM**	**ATO**	**GLM**
True Positive	673	673	87	87
True Negative	280	280	757	757
False Positive	349	349	132	132
False Negative	171	171	497	497
Frugality	2.003	2.003	1.849	1.849

Jenny et al., 2013). The first four statistical measures represent the error matrix (or confusion matrix) which count the number of cases for which the FFT's correctly or incorrectly identified the signal or noise. From these counts, essentially all other performance metrics of FFT's can be calculated. The last measure, frugality, is not a measure of signal detection theory, but was introduced to measure how 'fast-and-frugal' the FFT would be in practice (Martignon et al., 2008). The frugality of an FFT is the mean number of cues, across all options, that are used to make a judgment.

3.3 Representing Experience, Effort, and Time Pressure

The previous sections have developed representative mathematical models of many judgment and decision making strategies, however, no component was explicitly labeled experience, effort, or time – three very important characteristics of judges and decision makers. Nevertheless, there are techniques for modeling and representing these contexts, as described in the following subsections.

3.3.1 Experience

There are two ways to represent experience in the general linear model: first, when selecting the components or strategies used to represent the operator, and second, when generating values for components. One very general synthesis of the FFH literature on strategy selection is that experts or decision makers with high time pressure, high information acquisition costs, information overload, or ill-structured environments tend to make decisions matching heuristic strategies (relying on fewer, more important, cues). Novices or decision makers with low time pressure and low information acquisition costs tend to make decisions matching normative strategies (relying on many cues, e.g. a linear weighted additive model). For more thorough summaries of the FFH literature see Katsikopoulos (2011), Gigerenzer and Gaissmaier (2011), and Todd et al. (2012).

Beyond strategy selection, models and simulations approximate levels of experience and expertise by using different proportions of datasets, which represent the environment, when training the models. A novice may have very little experience with the environment and thus will base their judgments or decisions on only a few examples, whereas an expert will use information from many prior examples to make judgments and decisions. The relative level of expertise (or the proportion of the dataset used to train models) affects the types and amount of error in defining components of the models: the cue weights (w_j, e.g. Gigerenzer and Goldstein, 1996; Czerlinski et al., 1999), cue directions (d_j, e.g. Katsikopoulos et al., 2010; von Helversen et al., 2013), estimates of missing information (e_j), cutoff values (c_j), and thresholds (Δ_j). Therefore, experts are usually assumed to be "properly calibrated" in a modeling sense (Katsikopoulos et al., 2014, p. 155):

> "The decision maker... knows how to code the attributes [cues] a_j so that

the weights w_j are positive and is able to order the w_j according to their magnitude. These assumptions are particularly plausible when the decision maker is familiar with the choice."

3.3.2 Time and Effort

The literature of judgment and decision making support the bounded rationality perspective that people make decisions without gathering or processing all possible information. As described in Sec. 2.2, the most salient cause for making decisions with incomplete information is that not all information was available either because it was difficult to discern or not provided. These difficulties can create high information acquisition costs that are limited by time, money, cognition (effort), or some other resource. Typically, the limitations on time and effort are often used as proxies for each other in empirical and computational studies.

Limits on time and effort has been shown to affect the types of components and strategies used by judges and decision makers. Rieskamp and Hoffrage (2008) empirically identified that when time pressure was high in terms of both opportunity cost and individual decision limits, people used heuristics as opposed to normative decision strategies. This confirmed previous findings that people focus their attention on more important information, and thus use a more selective information search (e.g. Payne et al., 1988; Maule, 1994; Rieskamp and Hoffrage, 1999). Therefore, limits on time and effort, in addition to expertise, tends to result in the use of heuristic strategies that focus on a few, important cues.

Direct computational simulations of strategy performance with time or effort constraints have typically counted the number of elementary information processes to perform a strategy or its associated time (Payne et al., 1990). For example, through human-subjects studies, Bettman et al. (1990) estimated that it required 2.23 seconds to multiply a single cue value by a cue weight and therefore a lexicographic strategy (e.g. TTB) would require an average of 8.7 seconds to analyze a 2-option, 2-cue decision task while a normative weighted additive strategy would require 18.4 seconds. Related computational studies tracked the amount of time being used by strategies as they progressed through each decision and stopped the strategy when the time limit was reached (Payne et al., 1993). While this dissertation does not present a method for incorporating measures of elementary information processes into the general mathematical form, it is an important avenue of future work.

Outside of using direct time and effort estimations, mathematical models can approximate limits on time and effort as their functional equivalent: incomplete information. As the available time and effort decreases, less information can be processed. Similar to studies of high time pressure and limits on effort, studies have shown that heuristics can perform well in environments with low total information (Martignon and Hoffrage, 2002; Garcia-Retamero and Rieskamp, 2009). These synchronous results support the general thesis of the FFH program that people use the bounds on rationality such as lack of time and lack of information as a mechanism for simple, robust, and accurate strategies that adapt to the environment and ecology (Gigerenzer and Goldstein, 1996; Gigerenzer et al., 1999; Gigerenzer and Gaissmaier, 2011; Todd et al., 2012).

3.4 Summary

The general linear model presented in this chapter overcomes many of the limitations inherent to previous studies of naturalistic decision making, heuristic decision making, and weighted-utility models of decision making. The general linear model is transparent as it allows for specificity in model selection and environments at the component-level. The general linear model is representative as it can approximate expertise and various heuristic and naturalistic strategies. The general linear model is simple as it can be looked-up and used without specialty codes. Lastly, the general linear model supports applied work because it can model unique combinations of components without being restricted to previously-studied or established models of decision making.

The general linear model is not without its limitations. The primary limitation is caused by the same mechanism that makes the general linear model fast and transparent: each option is evaluated independently, even in decision making tasks. Therefore, in decision tasks with multiple options it is not clear how to incorporate the measures of effort or time into the model, nor how to account for certain types of decision strategies that rely on comparisons between options (satisficing and elimination by aspects). Additionally, there is no set of monotonically decreasing weights that mimic the compensatory strategy of Take Two (Dieckmann and Rieskamp, 2007) or two-cue confirmation (Karelaia, 2006). Future work will focus on how to address these limitations and how to develop a method of statistically fitting decision makers in real-time.

CHAPTER 4

MODELS, MEASURES, AND MEDIATORS OF DECISION MAKING PERFORMANCE

Decision makers are situated within an environment that shapes both the decision task and their decision making strategy. Based on an information process model of decision making, this chapter introduces a new theoretical model in Sec. 4.1 of how context (characteristics of the environment and strategies) determines decision making accuracy in decision tasks with incomplete information. The rest of the chapter follows the structure of the model which serves as a reference list of all the measures (Sec. 4.2) and mediators (Sec. 4.3) of decision making accuracy that will be measured in the following computational (Chap. 5-7) and human-subjects studies (Chap 8).

4.1 Models

A synthesis of the decision theory and behavioral decision making literature enables the construction of a basic model of contextual determinants of decision making accuracy as shown in Fig. 4.1. Environmental and task parameters, the decision making strategy, and incomplete information all affect decision making accuracy. An extensive amount of literature has been focused on how strategy and environment interact to affect accuracy in decision tasks with complete information in the decision task (Czerlinski et al., 1999; Hogarth and Karelaia, 2005b, 2007; Katsikopoulos et al., 2010; Hogarth and Karelaia, 2005a, 2006; Katsikopoulos, 2011; Gigerenzer and Gaissmaier, 2011; Todd et al., 2012). A much smaller number of studies have focused on decision making with incomplete information, but they are also limited by the singular focus on amount of information, not the distribution of incomplete information (see Sec. 2.3.2, Martignon and Hoffrage, 2002; Payne et al., 1990; Garcia-Retamero and Rieskamp, 2008).

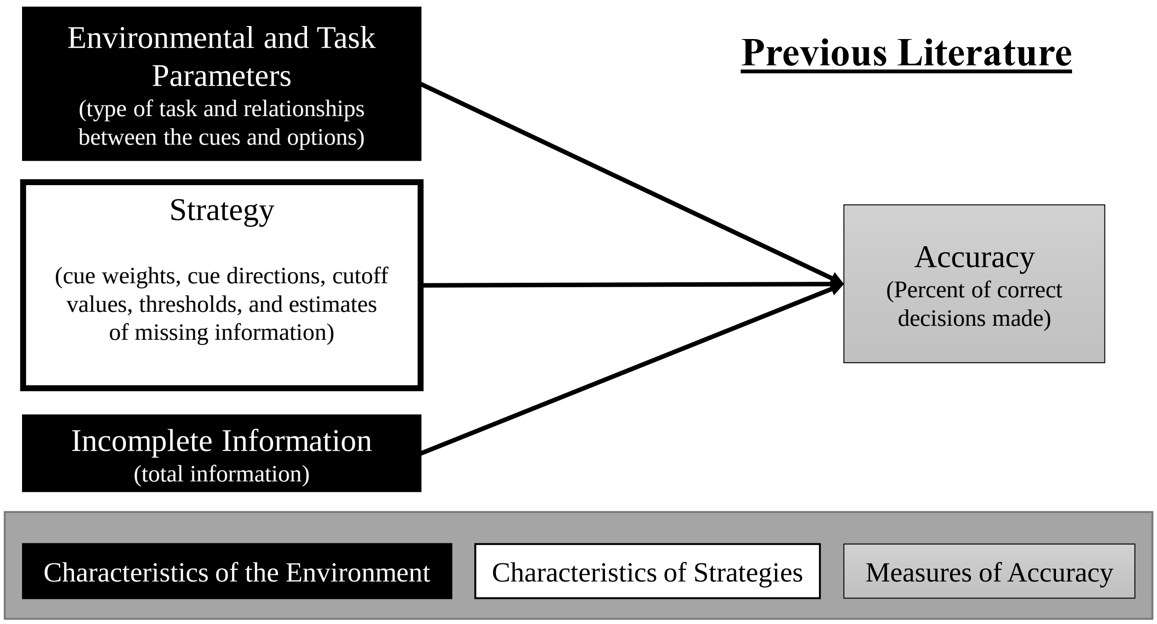

Figure 4.1: Model of contextual determinants of decision making accuracy based on prior literature.

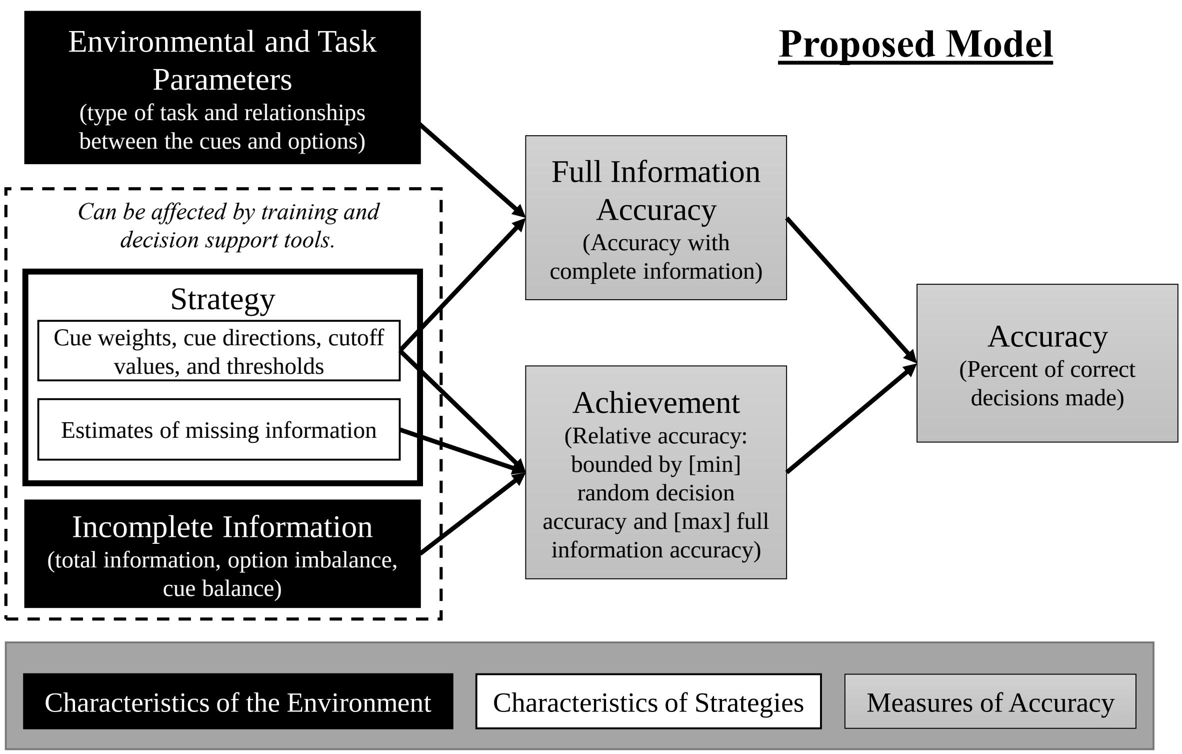

Figure 4.2: Model of contextual determinants of decision making accuracy based on the research presented in this dissertation.

The prior literature provided useful methods of how future studies could be designed to analyze distributions of incomplete information but their results do not have direct implications for two research questions governing this dissertation: how does incomplete information affect decision making performance, and how can decision support tools increase accuracy by acquiring and restricting information without probabilities, cue weights, and cue values?

The new model of contextual determinants of decision making accuracy in Fig. 4.2 was constructed to better explain the interaction of incomplete information and decision support on the performance of decision makers than prior literature. By comparing the two models, one can see the incorporation of new measures and mediators. The new model – and its supporting research in this dissertation – makes three main contributions to current decision theory.

First, the model introduces two new measures of decision making accuracy in Sec. 4.2: full information accuracy (FIA) and achievement. Full information accuracy is the accuracy of the strategy with full information. Achievement is the relative accuracy of the strategy with incomplete information scaled to be bounded by the minimum accuracy of random decisions and the maximum accuracy of full information accuracy. Achievement was introduced to separate the maximum achievable accuracy controlled by the fit of the environment and the decision making strategy (the focus of most prior studies), away from the effects of incomplete information on performance.

Second, the model separates the context variables that are innate to the environment and task (environmental and task parameters) from those which are under the control of the decision maker or can be addressed by decision support methods or training (strategy and incomplete information). This separation is mostly used for framing purposes rather than functional, but it is an important perspective. For decision makers in uncertain environments, they may not have control over, or even understand the nuanced characteristics of their environment and task. What they do have control over is what strategy they use and what information they want to integrate to come to a decision. These are the two aspects in which training and decision support can augment the skills of the decision makers: strategy selection and assessing incomplete information.

Third, the model introduces two new measures in Sec. 4.3 to more completely characterize distributions of incomplete information within a decision task: option imbalance and cue balance. Option imbalance measures the difference in amount of information between the most known option and the least known option. Cue balance counts the number of cues which have known values for each option.

4.2 Measures

Decision making performance can be viewed through the perspective of cost and benefit: the benefit is making the decision that maximizes the criterion value whereas the cost is the effort, time, money, or other resource expended to make the decision. Accuracy is clearly displayed in the contextual model of decision making in Fig. 4.2, however, this section will also focus on how effort and time are measured in computational and empirical studies.

4.2.1 Accuracy

The two new measures of decision making accuracy enable us to decompose the overall accuracy of a strategy into measures of accuracy based on how well the strategy fits the environment (full information accuracy, F) and how well the strategy performs within the limits of the fit (achievement, G). These two additional terms are inspired by the underlying logic of the lens model of judgment in Fig. 4.3 introduced by Brunswik (1943) and quantified by Hammond and his colleagues (Hammond, 1955; Hammond et al., 1964). The lens model attempts to describe the relationship between the environment and the behavior of organisms in the environment. The lens model uses a linear regression between the data used by the judge and the actual criterion values to compute a correlation matrix of all the cues and the two dependent variables. In decision making, the overall accuracy is equiva-

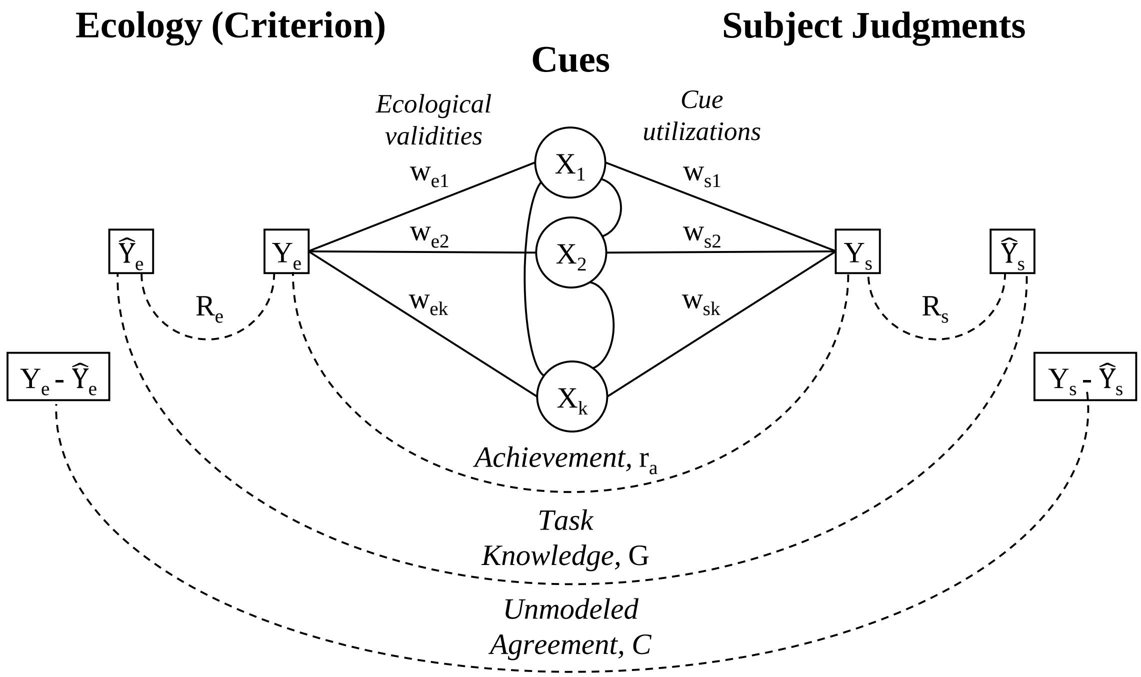

Figure 4.3: Lens Model with labeled statistical parameters (Cooksey, 1996; Bass, 2002; Rothrock and Kirlik, 2003).

lent to the lens model's achievement index, r, which is the correlation between the actual criterion and the judge's predicted criterion – thus, r has a maximum value of 1. The lens model equation for calculating the achievement index is presented in Eq. 4.1:

$$r_a = GR_eR_s + C\sqrt{1 - R_e^2}\sqrt{1 - R_s^2} \tag{4.1}$$

Just as a focus on overall accuracy within decision making limits the explanation power of decision analysis, the lens model accounts for the fact that achievement index alone does not explain why a judge's performance may change in different contexts. The achievement index is partially a product of how consistent the judge is in executing his or her strategy (cognitive control, R_s), how predictable the environment is with a linear model (environmental predictability, R_e), and how much linear knowledge the judge has about the environment (task knowledge, G). The following example from Bass (2002, p. 27) helps explain the terms' interactions:

> "To take an example, a particular human judge may demonstrate an achievement of (only) 0.75 on a traffic conflict prediction task. As the relationship above indicates, it could be that the judge's knowledge of the environment is limited ($G = 0.75$) but the environment is fully predictable ($R_e = 1.0$) and the judge is perfectly consistent in executing his or her judgment strategy ($R_s = 1.0$). On the other hand, the judge might have perfect linear knowledge and cognitive control but be working in an environment that is less than fully predictable ($R_e = 0.75$). Finally it could be the case that the environment is fully predictable and the judge has perfect knowledge but makes judgments inconsistently ($R_s = 0.75$)."

Just as the achievement index in the lens model is governed by the interaction of environmental predictability, cognitive control, and task knowledge, so too decision making accuracy with incomplete information is governed by the analogous components of full information accuracy (F), random decision accuracy (R), and achievement (G) as shown in Fig. 4.2 and Eq. 4.2. The accuracy of a strategy is generally bound by its full information accuracy as an upper limit and by random decision accuracy (R) as a lower limit. Full information accuracy measures accuracy of the decision making strategy with complete information (or knowledge) of the decision task. Full information accuracy approximates the fit of the strategy to the environment. Random decision accuracy measures the degenerate strategy of selecting at random. This is analogous to cognitive control as it is the default action when a decision maker cannot actively make a decision.

The relative values of F and R can be grouped into three pairings based on their implications for decision support and incomplete information: First, if F is higher than R, then more information should generally increase accuracy and a decision support system should help the decision maker execute the strategy appropriately. Second, if F is very close to R, then than no matter how much more information is provided or how well the decision maker executes the strategy, there will be little gain in accuracy compared to just selecting at random. Third, if F is lower than R, then the strategy is mismatched to the environment and helping the decision maker execute the strategy better will only serve to decrease overall accuracy.

$$G = \frac{A - R}{F - R} \tag{4.2}$$

Achievement (G), measures the decision making accuracy between the minimum accuracy of random selection (R) and maximum accuracy of full information accuracy (F). The measure describes how much of the maximum possible accuracy was achieved by the decision maker in the task with incomplete information. Achievement is a useful construct because it accounts for the limitations on performance due the fit of the strategy to environment (F) and the lack of the ability to make a decision (R).

The example in Fig. 4.4 shows how achievement can help differentiate between the effects of the strategy fit to the environment and the effects of incomplete information. Assume that there are two strategies (X and Y) used within a two-option decision task ($R = 50\%$), with two different full information accuracies such that strategy X ($F(X) = 95\%$) is better fit to the environment than strategy Y ($F(Y) = 65\%$). These two different strategies are applied to two specific distributions of incomplete information (1 and 2). The difference in accuracy between the two distributions for both strategies is 5%, which may make it seem like neither strategy would prefer one distribution to the other. However, through the perspective of achievement, it is clear that Strategy Y strongly prefers Distribution 1 (67%) to Distribution 2 (33%).

4.2.2 Effort and Time Required

Effort and time measure the total use of cognitive resources or time required to complete the decision making or judgment task (Payne et al., 1990). Both can be a measure and a mediator of decision making performance. As a measure of performance, judges and decision makers often have the goal of minimizing the amount of time and effort required

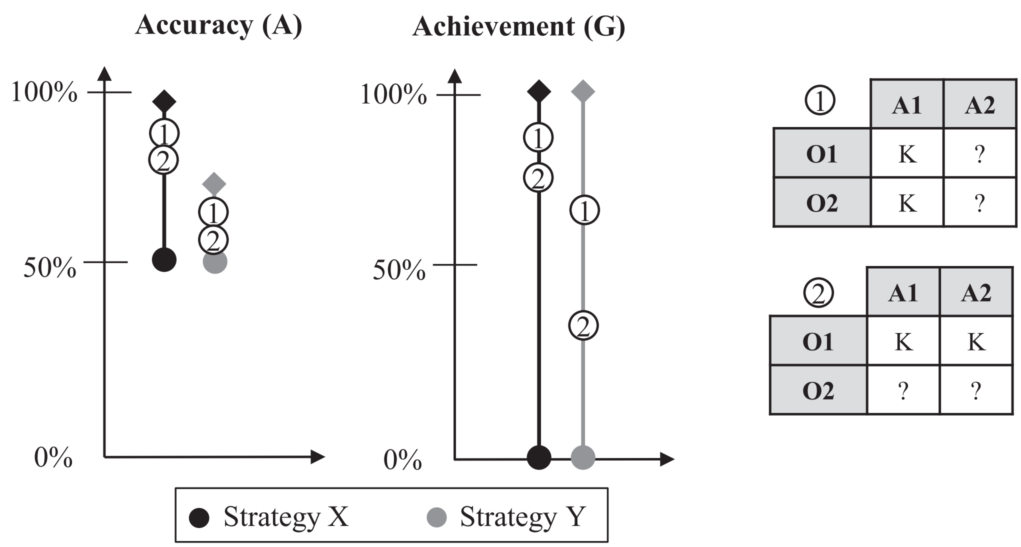

	Strategy X (F = 95%)		Strategy Y (F = 65%)	
	Accuracy (A)	**Achievement (G)**	**Accuracy (A)**	**Achievement (G)**
Distribution 1	90%	89%	60%	67%
Distribution 2	85%	78%	55%	33%

Figure 4.4: Description of accuracy and achievement.

to make a judgment or decision. As a mediator of performance, perception of how much effort or time is required to perform a strategy is used to feed back into the strategy selection process (Todd and Benbasat, 1999).

In human-subjects and computational studies, time and effort are used as proxies for each other because time is more easily measured in human-subjects studies whereas effort is more easily measured in computational studies. In human-subjects studies, measuring time required is as simple as measuring the time from the presentation of the task to the decision. Computational studies have used elementary information processes (EIPs) to estimate the mental effort required to perform various strategies (Payne et al., 1988, 1990; Mata et al., 2007; Payne et al., 1996). Based on estimates of time required for each EIP (as shown in Table 4.1), one can then calculate the total time required to perform a strategy (Bettman et al., 1990). However, even though EIPs had been used in these referenced studies, none explicitly stated how EIPs were mapped to the strategies.

Table 4.1: List of elementary information processes (EIPs) with time estimates (Bettman et al., 1990; Payne et al., 1990).

Time (s)	Elementary Information Process Description
1.19	`Read` a single decision option's value for an cue into working memory. Value is defined as the measurable quantity of cue in its standard units.
	`Estimate` a decision option's score on an cue. Score is defined as the estimated utility (desirability) of the quantified cue value.
0.84	`Add` two cue scores together.
0.32*	Calculate the size of the `difference` of two decision options on an cue; equivalent to the number of subtraction signs in any calculation.
0.09*	`Compare` two numbers (cue scores, decision option scores, etc.) to determine greater than, less than, or equal; equivalent to the mathematical operators $<$, $>$, and $=$.
2.23	`Weigh` one cue score by its cue weight; equivalent to the number of multiplication signs required in any weighting calculation.
1.80	`Eliminate` a decision option from consideration.
1.80	`Choose` the preferred decision option and end the process.
	`Move` to the next element or cue of the task environment.

*Not significantly different from zero at $p < 0.05$ (Bettman et al., 1990)

The measure of effort for a strategy within a scenario is the total count of EIPs required to choose a decision option. If necessary, it can be assumed that each EIP requires the same amount of time or mental effort and a decision maker does not place greater importance on one EIP versus another. This assumption that all EIPs require the same amount of time or mental effort should not greatly affect the conclusions because simulation studies have shown that assuming all EIPs require equal time and effort leads to almost identical conclusions as studies without the assumption (Payne et al., 1988, 1996).

4.3 Mediators

4.3.1 Strategy

To achieve the goal of choosing the option with the maximum criterion, decision makers employ strategies which are defined as the overall methods, or sequences of options, for searching through a decision problem space (Payne et al., 1993). Due to the variety and number of published strategies for decision making, this dissertation examined a subset of well-studied strategies to span the range from analytic to heuristic decision making: Weighted Additive (WADD), Equal-Weighting (EW), Tallying, Take Two, and Take-the-Best (TTB).

To characterize this broad range of decision making strategies, we leverage the component-strategy general linear model (GLM) of decision making strategies from Chap. 3 as shown in Table 4.2. The GLM constructs strategies by combining variations of three components: utility functions, cue weights, and estimates of missing information. Utility functions describe how the strategy converts cue values describing the environment into cue scores with units useful for decision making. Cue weights describe how strategies measure the relative importance of each cue for predicting the criterion. Estimates of missing information describe how strategies estimate the cue value when it is unknown. As the process of each strategy is described below, the variations of the components will be specified.

WADD makes decisions by first calculating the criterion of each option by multiplying each of the cue weights by their respective cue scores and then selecting the option with the highest criterion. WADD uses compensatory cue weights meaning that it does not have any restrictions on what the cue weights should be, allowing high scores on lower-weighted cues to potentially compensate for low scores on higher-weighted cues. When cue values are unknown, WADD estimates the cue score as the average of typical cue scores. Since we assume that the strategies do not have any prior information about the distribution of the cue scores, the estimate is 0.5 – the expected value of the binary attribute scores. If the

Table 4.2: Description of the five selected decision making strategies using the general linear model of decision making in Chap. 3.

Strategy (Operationalization)	**Weighted Additive (WADD)** (Payne et al., 1990)	**Equal-Weighting (EW)** (Payne et al., 1990)	**Tallying** (Gigerenzer and Goldstein, 1996)	**Take Two** (Karelaia, 2006; Dieckmann and Rieskamp, 2007)	**Take-the-Best (TTB)** (Gigerenzer and Goldstein, 1996)
What type of cue scores are used?	Binary	Binary	Binary	Binary	Binary
How to estimate missing information?	Average (0.5)	Average (0.5)	Negative (0)	Negative (0)	Negative (0)
What types of cue weights?	Compensatory	Equal	Equal	Compensatory	Non-Compensatory
Describes what type of decision maker?	Design decisions and multi-attribute decision making (Park, 2004); Complex underlying processes (Hogarth and Karelaia, 2007; Einhorn et al., 1979); Novice decision makers (Garcia-Retamero and Dhami, 2009)	When cue weights are difficult to estimate, predict, or be agreed upon (Dawes, 1979); Cues are from “from diverse and incomparable sources” (Dawes, 1979, p. 574)	Similar to EW; Examples include eye exams for detecting stroke (Kattah et al., 2009) and avoiding avalanche accidents (McCammon and Hägeli, 2007)	Young adults (Mata et al., 2007); Low information redundancy (Dieckmann and Rieskamp, 2007)	Time stressed decision making (Rieskamp and Hoffrage, 2008); Consumer choice (Kohli and Jedidi, 2007); Expert decision making (Garcia-Retamero and Dhami, 2009)

criterion scores are equal, then WADD randomly selects between the equivalent options.

EW uses a process identical to WADD except that it does not use cue weights, thus the cue weights are all equivalent (Dawes and Corrigan, 1974; Dawes, 1979). The criterion of each option is calculated by adding all of the cue scores together and selecting the option with the highest criterion score.

Unlike WADD and EW, the heuristic strategies of Tallying, Take Two, and TTB categorize the binary cue scores as positive if 1, negative if 0, and unknown cue values are estimated as negative (0). The positive or negative categorization can be conceptualized as a judgment of that cue score as being good or bad, respectively.

Tallying makes decisions by counting the number of positive cues and selecting the option with the most positive cues (Gigerenzer and Goldstein, 1996). If two or more options are equivalent in the number of positive cues, Tallying randomly selects between the tied options.

TTB searches through cues in rank-order, stops when one cue discriminates, then selects the option with the positive cue score (Gigerenzer and Goldstein, 1996). For 2-option decision tasks, an individual cue discriminates when one option has a positive cue score and the other option have a negative or unknown cue score. The non-compensatory cue weights represent the fact that once a cue discriminates, the decision will not change regardless of the information processed for lower-ranked cues (Martignon et al., 2008). Ultimately, if no cue discriminates, then TTB selects randomly between the options.

Take Two uses an identical process as TTB with a modification requiring that two cues discriminate in favor of the selected option (Karelaia, 2006; Dieckmann and Rieskamp, 2007). Though this may seem to be a small difference, it is significant as it allows for lower-ranked cues to compensate while TTB is strictly non-compensatory. In both instances where Take Two processes were introduced (Karelaia, 2006; Dieckmann and Rieskamp, 2007), they were "developed to overcome the descriptive shortcomings of TTB," (Hogarth and Karelaia, 2007, p. 735), i.e., the experimental results showing that, in some cases, people seek information beyond the first discriminating cue.

4.3.2 Environmental Parameters

Environmental parameters are commonly studied as mediators of decision making, both in simulation and empirical studies.[1] There is no particular definition of what is and is not an environmental parameter. Therefore, in this dissertation, environmental parameters are defined as intensive characteristics of the decision task: properties of the environment that do not change based on the knowledge or lack of knowledge of the cue scores within an individual decision task. For example, redundancy is an environmental parameter that measures the amount of correlation between the cues. Neither the number of cues known within a specific decision task, nor the values of the cues within a specific decision task, affect the value of redundancy for a decision making environment.

Environmental parameters are of particular interest to the study of judgment and decision making in the perspective of bounded rationality. Todd et al. (2012) introduced ecological rationality to explain that within the specific ecologies (environments) that judges and decision makers are placed in, they are often rational - though they may not follow the tenants of normative, unbounded rationality. This focus on environmental parameters has led to the identification of many parameters that affect judgment and decision making performance; for example, the number of cues (e.g. Payne et al., 1990; Hogarth and Karelaia, 2006), predictability (e.g. Karelaia and Hogarth, 2008), redundancy (e.g. Dieckmann and Rieskamp, 2007; Hogarth and Karelaia, 2007), variability (e.g. Payne et al., 1990; Hogarth and Karelaia, 2005a; Gigerenzer and Gaissmaier, 2011), and distribution (e.g. Martignon and Hoffrage, 2002; Hogarth and Karelaia, 2005a, 2006, 2007).

So much research has been completed that the environmental parameters listed in this section can be considered well studied. The following subsections define each of the environmental parameters studied in this dissertation, along with a summary of how they have

[1] This dissertation uses the term environmental parameters (Hogarth and Karelaia, 2007) while other works have used the term environmental structures (Gigerenzer and Gaissmaier, 2011; Todd et al., 2012).

been shown to affect decision making performance of the five strategies from Sec. 4.3.1. (For reviews of environmental parameters, see Karelaia and Hogarth, 2008; Hogarth and Karelaia, 2007; Todd et al., 2012.)

Cues

The number of cues describes the size of the decision task, in conjunction with the number of options. Isolating the specific effect of number of cues on the performance of decision making strategies is difficult due to the correlation with other environmental parameters. For example, mathematical analysis shows that if one assumes a random distribution of cue weights, as the number of cues increases, the likelihood that the cue weights are non-compensatory decreases (matching the cue weight structure of heuristics like take-the-best) and the likelihood that the cue weights are compensatory increases (matching the cue weight structure of analytic strategies like equal weighting or weighted additive) (Hogarth and Karelaia, 2006). As discussed in Sec. 4.3.2, matching the distribution of cue weights to the strategy is a strong positive influence on decision making accuracy. In a meta-analysis comparing the effects of various environmental parameters on the accuracy of judgment models in 249 environments from 86 articles, the number of cues had the least explanatory power as compared to redundancy, cue weights (variability and distribution), functional form in the ecology (predictability), expertise, and learning (Karelaia and Hogarth, 2008). So although many simulation studies evaluate a range of numbers of cues, they rightfully caveat their interpretations of decreased or increased accuracy by discussing distribution directly (Hogarth and Karelaia, 2006) or other related parameters such as variability (Payne et al., 1990, 1993).

With respect to the studies in this dissertation, the decision tasks will only include 3 to 5 cues. This range is used in order to reduce analytical complexity and because it is seen as “sufficient to understand what people can actually do within limited information processing constraints” (Hogarth and Karelaia, 2006, p. 246).

Distribution

Distribution of cue weights is categorized into three groups (Hogarth and Karelaia, 2007): non-compensatory (non-CF), when the cue weights are ordered in magnitude the weight of each attribute exceeds the sum of those smaller than it (Martignon and Hoffrage, 2002); equal-weighted (equal-CF), all cue weights are equivalent; and compensatory (CF), all other combinations of cue weights. Distribution of cue weights has been extensively studied so the conclusion is fairly straightforward: when the distribution of cue weights in the environment matches the representative cue weights of the strategy, the strategy fits the environment, resulting in high accuracy (Hogarth and Karelaia, 2007; Payne et al., 1993; Karelaia and Hogarth, 2008; Hogarth and Karelaia, 2005a, 2006). Thus, TTB and Take Two have the best accuracy in non-compensatory environments, EW and Tallying have the best accuracy equal-weighted environments, and WADD has the best accuracy in compensatory environments.

Redundancy

Redundancy is the average of the absolute values of correlation coefficients between the cues (Hogarth and Karelaia, 2005a; Karelaia and Hogarth, 2008). High levels of redundancy (0.50) indicate that the cues are strongly correlated, facilitating the interchangeability of cues which could contribute to improving the reliability of heuristic decisions because information can be reduced without significant reductions in accuracy (Karelaia and Hogarth, 2008; Dieckmann and Rieskamp, 2007). Low levels of redundancy (0.10) are also possible. Simulations by Dieckmann and Rieskamp (2007) showed that TTB, Take Two, and NB (similar to WADD) had higher accuracy in environments with low redundancy than in environments with high redundancy. More specifically, Hogarth and Karelaia (2007) used simulations to show that in compensatory (CF) environments TTB performs best when redundancy is high; however, in non-compensatory (non-CF) environments, TTB performs best when redundancy is low.

Note that this measure of redundancy, due to the absolute values, does not account for whether cues are positively or negatively correlated. These respective cue correlations have been termed 'friendly' and 'unfriendly' environments for heuristics because positively correlated cues facilitate the replacement of cues whereas negatively correlated cues (Shanteau and Thomas, 2000). Though friendly and unfriendly environments will not be studied in this dissertation, the general consensus is that environments can be considered unfriendly to heuristics using fewer cues only if there are negative correlations among equally weighted attributes (Fasolo et al., 2007).

Variability

Variability of cue weights is the difference between the maximum and minimum cue weights (Hogarth and Karelaia, 2005a; Gigerenzer and Gaissmaier, 2011; and referred to as dispersion by Payne et al., 1990). For normalized cue weights, low dispersion is typically around 0.10 while high dispersion is 0.50 or higher.[2] Hogarth and Karelaia (2005a) showed that increasing variability increases the accuracy of a single cue heuristic (similar to TTB), equal weighting, and multiple linear regression. However, as stated in the description of cue weight distributions, the effects of variability are overtaken by the match of the strategy's cue weights to the environment's cue weights.

Predictability

Environmental predictability (R_e, hereafter referred to as predictability), from the lens model, measures the linear relationship between the cues and the criterion: how predictable is the criterion with an optimal linear model? Specifically, predictability is calculated as the R^2 value of the linear least-squares regression model for the dataset with the criterion as the dependent variable and the cue scores as the independent variables. Most of the

[2] While trends are easily comparable across different studies, when comparing specific values of variability is it important to account for the types of cue weights being used. Ecological validity weights are bounded from 0 to 1, whereas the Goodman-Kruskal rank correlation weights are a rescaling of ecological validity weights to bounds of -1 to 1 (Martignon and Hoffrage, 2002).

research using predictability comes from the study of judgment with the general consensus that accuracy across all strategies is lower in nonlinear environments (Karelaia and Hogarth, 2008).

4.3.3 Task Parameters

As opposed to environmental parameters, task parameters characterize individual decision tasks without respect to the overall environment or intensive features. Two important parameters will be examined in the computational studies: dispersion and dominance.

Dispersion

Dispersion describes the difference in cue values between options. The simulation uses the mean absolute deviation about the mean (MADM) to calculate this context feature in accordance with Payne et al. (1990) who calculated the dispersion between varying cue weights. Though Payne et al. (1990) did not mention the exact method for calculating dispersion, MADM is a standard measure of statistical dispersion and can scale to any number of options and cues, as long as there ratio scale cue values.

Dispersion was only used the computer study in Chap. 5 with a cue value scale between 0 and 100. Therefore, three levels of dispersion were categorized based on the value of MADM: 1) Low dispersion occurs when two decision options are, on average, very similar across all four decision cues (MADM $\leq$ 5); 2) Medium (5 $<$ MADM $<$ 15); and 3) High (15 $\geq$ MADM) dispersion occur when two options are noticeably different. The low and high dispersion levels corresponded to the exemplars used in Payne et al. (1990).

Dominance

Dominance is defined as whether one option has a greater cue value for every cue than the other option. There are three categories of dominance for a two-option decision: 1) Strong dominance occurs when one option has a better cue value for every cue than the other option; 2) Weak dominance occurs when one option has a better cue value for at least cue and never has a worse cue value than the other option; and 3) Non-dominance occurs when each option has at least one cue value which is better than the other option. In contrast to dispersion, which is a holistic comparison between two decision options, dominance compares the individual cue values between options.

4.3.4 Incomplete Information

Decision making with incomplete information refers to unknown cue values within the decision task (Martignon and Hoffrage, 2002; Garcia-Retamero and Rieskamp, 2008). Incomplete information is often caused by time pressure, high information acquisition costs, too much information, or incomplete information in the environment (see Sec. 2.2). Prior to the research presented in this dissertation, the only direct measure of incomplete information in traditional decision theory was total amount of information (Martignon and Hoffrage, 2002; Garcia-Retamero and Rieskamp, 2008; for an exception in consumer and marketing research see Kivetz and Simonson, 2000). To better describe the distribution of incomplete information two new measures have been introduced: option imbalance (OI) and cue balance (CB).[3] The two measures of information balance, option imbalance and cue balance, measure whether there is equal distribution of information between options and within cues, respectively.

As described in Chap. 3, each decision task is made up m options, each with n cues. An individual option is represented as i and cue as j. The matrix of known and unknown cue scores is denoted Z such that if option i's j^{th} cue score is known to the decision maker $z_{i,j} = 1$ whereas $z_{i,j} = 0$ indicates that option i's j^{th} cue score is unknown to the decision maker. The counts for the three measures (total information, option imbalance, and cue balance) are calculated within the brackets of Eqs. 4.3-4.8 and shown in Fig. 4.5. Due to the variation in maximum values for each measure resulting from the changing number of

[3] Option imbalance is new name for information imbalance (Canellas et al., 2014, 2015; Canellas and Feigh, 2017). Cue balance is a new name for complete attribute pairs (Canellas and Feigh, 2014, 2017).

Decision Tasks

A

n' = 1	Cue 1	Cue 2
Option 1	K	?
Option 2	?	?

B

n' = 1	Cue 1	Cue 2
Option 1	K	?
Option 2	K	?

C

n' = 2	Cue 1	Cue 2
Option 1	K	K
Option 2	?	?

D

n' = 2	Cue 1	Cue 2
Option 1	K	K
Option 2	K	?

E

n' = 2	Cue 1	Cue 2
Option 1	K	K
Option 2	K	K

Measures of Incomplete Information

Decision Task	Total Information		Option Imbalance		Cue Balance	
	Fixed	**Relative**	**Fixed**	**Relative**	**Fixed**	**Relative**
A	25%	50%	50%	100%	0%	0%
B	50%	100%	0%	0%	50%	100%
C	50%	50%	100%	100%	0%	0%
D	75%	75%	50%	50%	50%	50%
E	100%	100%	0%	0%	100%	100%

Figure 4.5: Measures of distributions of incomplete information.

cues, the measures are calculated as percents of the maximum values to enable comparison.

There are two versions of each of the three measures, with the versions separated by the number of cues used to convert the count into a ratio. Fixed versions of the measures divide the counts by the total number of cues in the task, n. Relative versions of the measures divide the counts by the number of cues for which there is a known cue for any option, n'. The relative measures are introduced to give a more specific understanding of the distribution of incomplete information. For example, in a decision task with one cue, two options, and all information known, the fixed and relative cue balance is 100%. If you add one more cue with no known information, the fixed cue balance decreases to 50% while the relative cue balance stays at 100% (see decision task B in Fig. 4.5).

Total Information

Total information (TI), as calculated in Eq. 4.3, measures a scale of how much information, or data, the decision maker has regarding the decision event. Total information is varied by changing the number of cue values known for each decision option. 100% total information indicates that the decision task is fully known by the decision maker whereas 0% total information indicates that the decision maker has no information about the decision task. Relative total information (TI') only considers the total information relative to the cues that have some known information. Previous work is in agreement that more information generally increases accuracy of decision making strategies though the amount of increase is dependent on the strategy (Martignon and Hoffrage, 2002; Garcia-Retamero and Rieskamp, 2008).

$$TI = \left[\sum_{i=1}^{m} \left(\sum_{j=1}^{n} z_{i,j} \right) \right] \cdot \frac{100}{m \cdot n} \tag{4.3}$$

$$TI' = \left[\sum_{i=1}^{m} \left(\sum_{j=1}^{n} z_{i,j} \right) \right] \cdot \frac{100}{m \cdot n'} \tag{4.4}$$

Option Imbalance

Option imbalance (OI), as calculated in Eq. 4.5, measures the difference in total information between the option with the most information and the option with the least information.[4] Option imbalance is similar to total information but captures the distribution of information for situations in which more information is known about one option than another (e.g. one is a default option). 100% option imbalance indicates that the decision maker has full information about one option and no information about the other; whereas 0% option imbalance indicates that the decision maker has equal information about both

[4]In publications prior to this dissertation, option imbalance was referred to as information imbalance (Canellas et al., 2015; Canellas and Feigh, 2017). Option imbalance is now the preferred name because it specifies the term as a measure of information imbalance between options.

options. Relative option imbalance (OI', Eq. 4.6) measures the option imbalance only with respect to the cues that have some information known.

$$OI = \left[max_i \left(\sum_{j}^{n} z_{i,j} \right) - min_i \left(\sum_{j}^{n} z_{i,j} \right) \right] \cdot \frac{100}{n} \tag{4.5}$$

$$OI' = \left[max_i \left(\sum_{j}^{n} z_{i,j} \right) - min_i \left(\sum_{j}^{n} z_{i,j} \right) \right] \cdot \frac{100}{n'} \tag{4.6}$$

Cue Balance

Cue balance (CB), as calculated in Eq. 4.7, measures how many cues have all values known for all options.[5] 100% cue balance indicates that every cue has a known cue score for every option, whereas, 0% cue balance indicates that no cues have a known cue score for every option. Due to the cue-wise information search that many strategies use (e.g. TTB and Take Two), cue balance measures how many fully informed comparisons of options these strategies are able to make.

$$CB = \left[\sum_{j=1}^{n} \mathbb{1}_n \left(\sum_{i=1}^{m} z_{i,j} \right) \right] \cdot \frac{100}{n} \tag{4.7}$$

$$CB' = \left[\sum_{j=1}^{n} \mathbb{1}_n \left(\sum_{i=1}^{m} z_{i,j} \right) \right] \cdot \frac{100}{n'} \tag{4.8}$$

4.4 Summary

This chapter introduced a conceptual model of how context influences the decision making accuracy of decision tasks with incomplete information. The model separates the context variables that are innate to the environment (environmental parameters) from those which

[5] In publications prior to this dissertation, cue balance was referred to as complete attribute pairs (Canellas and Feigh, 2014, 2017). Cue balance is now the preferred name as it compliments option imbalance and because the term can scale to more than two options ('pair' was a specific reference to two-option decision tasks).

are under control of the decision maker or which can be addressed by decision support tools or training (strategies and incomplete information). The conceptual model guides the rest of this dissertation by listing measures and mediators of accuracy that will be examined in the following computational (Chap. 5-7) and human-subjects studies (Chap. 8). The introduction of two new measures of accuracy (full information accuracy and achievement) and the two new measures of incomplete information (option imbalance and cue balance) will help the following studies better explain how incomplete information affects decision making performance and what heuristic information acquisition and restriction methods can be defined to support better performance.

CHAPTER 5

COMPUTER SIMULATION STUDY 1: ACCURACY, EFFORT, AND DISTRIBUTIONS OF INCOMPLETE INFORMATION

The results in this section are an integrated version of the the following articles: Canellas et al., 2014; Canellas and Feigh, 2014; and Canellas et al., 2015.

How does incomplete information affect decision making accuracy? How are strategies differentially affected by changing distributions of incomplete information? Can information acquisition and restriction methods which only rely on knowledge of the distribution of incomplete information increase decision making accuracy? These are the motivating questions of this dissertation. The three studies discussed in this chapter and the next two (Chap. 6 and Chap. 7) address these questions through computer simulations which represent the performance of decision making strategies in a variety of environments with incomplete information while measuring accuracy, effort, and task and environmental parameters (for a comprehensive list of measures and mediators, see Chap. 4). The results raised further questions that could only be addressed with a human-subject study which is reported in Chapter 8.

The purpose of this computational study is to document the required effort and obtainable accuracy when using different decision strategies and to understand the influence of the distribution of incomplete information on the two measures. The goal is to apply this understanding in the design of decision support systems (DSSs) that are capable of adapting to the differential needs of each decision strategy. As stated by Todd and Benbasat (1999, pp. 370), “it is clear that ‘engineering’ effort considerations as a conscious part of DSS design is central to the notion of guiding the decision maker towards the objective of ‘working smarter’ [yet the question remains as to] how can such guidance be provided?”

In particular, this study aims to improve understanding of how DSSs can provide guidance for information acquisition or restriction.

Decision makers are known to commonly operate under scenarios of incomplete information due to lack of available information or limited resources (Sec. 2.2), yet simulation studies (Payne et al., 1990) and empirical studies (Rieskamp and Hoffrage, 2008) comparing or identifying decision strategies have focused explicitly on time pressure. Studying incomplete information outside of time pressure enables the identification of previously unstudied relationships between different measures of incomplete information and the accuracy of, and amount of effort required when using, selected decision making strategies. To this end, a simulation engine has been developed that allows for evaluation of specific decision strategies as independent modules (see Fig. 5.1). The simulation measures accuracy and effort in accordance with cost-benefit (effort-accuracy) framework simulations (Payne et al., 1993), while also incorporating two new measures of distribution of incomplete information within a decision task: option imbalance and cue balance (see Sec. 4.3.4).

The purpose of identifying these new relationships is to use them to suggest novel DSS information acquisition and restriction measures for decision makers with incomplete information. Building upon the literature describing how DSSs may acquire or restrict the information search of decision makers (Todd and Benbasat, 1999; Silver, 1990; Nelson, 2005), this simulation engine has the capability to simulate information acquisition and restriction by approximating information search patterns. Following a series of distributions of incomplete information with increasing information approximates information search patterns or potential methods for acquisition, and following a series of distributions with decreasing information approximates methods for restricting information. Previous work studying time pressure cannot be leveraged in DSS design in this way, because DSSs cannot easily alter many of the factors that produce time pressure.

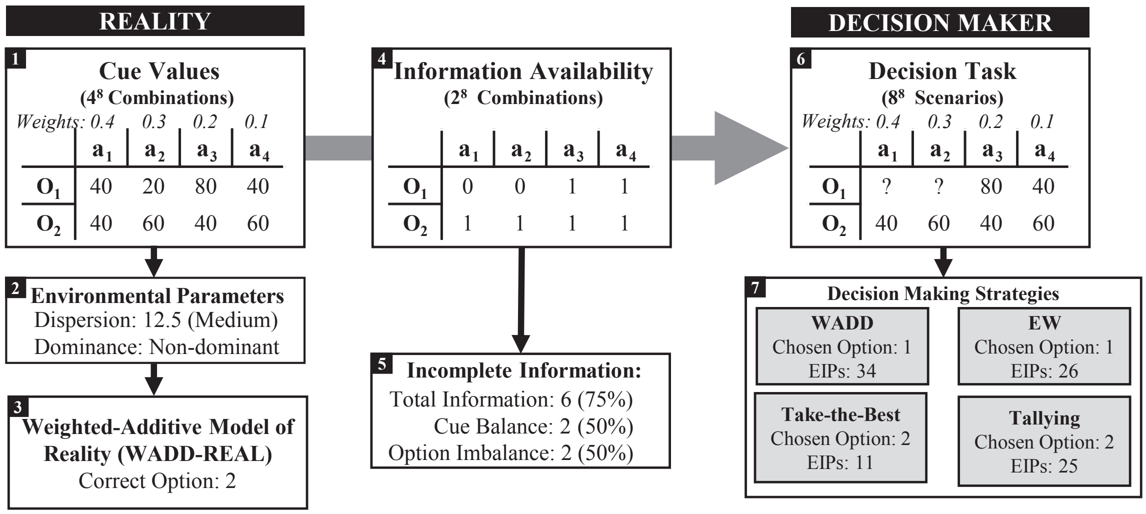

Figure 5.1: Model of the simulation engine for decision making with incomplete information used in this study: (Box 1) Combinations of cue values are generated representing a unique instance of the environment. From this set of environmental cue values, the two environmental parameters of dominance and dispersion are calculated (Box 2) and the weighted-additive environmental model identifies the correct option (Box 3). (Box 6) The decision maker receives the decision task in which some of the environmental values are unknown (Box 4). This information availability could represent sensor status or reliability. (Box 5) From the information availability (Box 4), the measures of incomplete information are calculated. (Box 7) This decision task is input into each decision making strategy which output their chosen option and the elementary information processes (EIPs) so that accuracy and effort can be measured.

5.1 Method

The focus of this study is to determine how three measures of incomplete information (total information, option imbalance, and cue balance) and two task parameters (dominance and dispersion) impact the accuracy of and effort required when using two heuristic (take-the-best, TTB; and Tallying) and two analytic strategies (weighted-additive, WADD; and equal-weighting, EW). The study also investigated how the mechanisms of the strategies themselves affect obtainable accuracy and required effort. To this end, a computational simulation was used to evaluate the decision making strategies under identical scenarios. The decision task simulation has three variables: decision strategy, cues scores, and distribution of incomplete information.

5.1.1 Scenario Generation

Each cue for each option had two possible information availability levels (cue value known, 1, or cue value unknown, 0) and four possible cue values (20, 40, 60, and 80). With these two factors and associated levels a full factorial DOE was created which resulted in $2^8 \times 4^8$ ($16,777,216$) decision tasks for each decision strategy.

5.1.2 Measuring Accuracy

For this study, correct decisions were determined by the option with the higher criterion determined by the weighted additive model of decision making without incomplete information using cue weights, $w = \{0.4, 0.3, 0.2, 0.1\}$, referred to as WADD-REAL. It is acknowledged that the research on ecological rationality which stresses that the correct strategy for defining accuracy depends on the task environment, and our future work will investigate alternative methods for selecting the correct decision. However, for this initial investigation, which is based on generic tasks with no objectively correct option, WADD-REAL is a justifiable definition of a correct decision. Reasons for selecting WADD-REAL include its use as the criterion for multi-attribute decision accuracy within the cost-benefit framework upon which this study is founded (Payne et al., 1993), and that "true" utility functions for assessing accuracy of decision strategies are still commonly assumed to be a linear function of the attributes (Katsikopoulos, 2013). Accuracy, therefore, is defined as the percentage of scenarios in which a decision making strategy chooses the correct decision.

5.1.3 Measurement of Effort

Within decision making, effort can be characterized as the total use of cognitive resources required to complete the task (Payne et al., 1990). As described in Sec. 4.3, the computational forms of the decision strategies enable them to be decomposed into elementary information processes (EIPs). Table 5.1 adapts the nine EIPs to each decision making strat-

Table 5.1: Mapping of elementary information processes (EIPs), which approximate effort, to the individual decision making strategies.

WADD/EW	TALLYING	TTB
1) `Read [#1]` all 8 cue values into the simulation (8 EIP).	**1)** `Read [#1]` all 8 cue values into the simulation (8 EIP).	**1)** `Read [#1]` two cue values into the simulation (2 EIP).
2) `Estimate [#2]` the cue scores from the cue values, using unity function or set to 50 if unknown (8 EIP).	**2)** `Estimate [#2]` the cue scores from the cue values, using unity function. Do not estimate if cue values are unknown (0 EIP if 0 total information, 8 EIP if 8 total information).	**2)** `Estimate [#2]` the cue scores from the cue values, using unity function. Do not estimate if cue values are unknown. (0 EIP to 2 EIP).
3) `Weight [#9]` each cue score by the provided cue weights (8 EIP). *EW does not perform this step.*	**3)** `Compare [#5]` all known cue scores to the cut-off score of 50 to determine positive or negative cues. (0 EIP to 8 EIP).	**3)** `Compare [#5]` the known cue scores to the cut-off score of 50 to determine positive or negative cues (0 EIP to 2 EIP).
4) `Add [#3]` the weighted cue scores together for each option to get two option scores (8 EIP).	**4)** `Add [#3]` the number of positive scores for each option. (0 EIP to 8 EIP).	**4)** `Compare [#5]` the cue scores of each option (1 EIP).
5) `Compare [#5]` the criterion for each option to determine which is larger (1 EIP).	**5)** `Compare [#5]` the criterion for each option to determine which is larger (1 EIP).	**5A)** If one option has a positive cue and the other has a negative or unknown cue, then `Choose [#8]` the option with the positive cue (1 EIP).
6) `Choose [#8]` the decision option with the higher criterion OR choose at random if criteria are equal (1 EIP).	6) `Choose [#8]` the decision option with the higher criterion OR choose at random if criteria are equal (1 EIP).	**5B)** If otherwise, repeat 1-5 process with next highest weighted cue. If no more cues, `Choose [#8]` at random (1 EIP).

Not all EIPs from Sec. 4.2.2 are used due to the nature of the strategies.

egy. EIPs have been used to measure mental effort in both empirical and simulation studies of decision making (Payne et al., 1990; Bettman et al., 1990; Mata et al., 2007).

5.2 Results

The results presented in the following subsections describe the effects of measures of incomplete information (total information, option imbalance, cues balance) and task parameters (dominance and dispersion) on the accuracy and effort for each decision strategy. Table 5.2 presents each of the measures of incomplete information, the number of generated tasks matching each level of incomplete information, accuracy, and effort required. Table 5.3 presents each of task parameter (dispersion and dominance), the number of generated tasks matching each level, accuracy, and effort required. These results do not include scenarios in which the options are equivalent (cue scores for Option 1 are identical to Option 2) or unknown (total information is 0) as these scenarios only constitute 0.8% of the total and would result in a random guess for all the decision strategies. The results include the 4.0% of the scenarios in which option scores as calculated by WADD-REAL were equal to each other such. Both options in these tasks are treated as correct. All results are reported as significant at the $\alpha = 0.001$ level to reflect the large size of the dataset. Results for accuracy were analyzed using contingency tables and logistic regression with a χ^2 distribution as the accuracy metric effectively became count data. Results for correlations were analyzed using Spearman's rank correlation coefficient. Results for effort were analyzed using a multi-phase analysis with effort requirements (EIP counts) as the response variable.

5.2.1 Strategy

Strategy had a significant effect on accuracy $(\chi^2(3, 6.66\text{E}7) = 558,294, p < 0.0001)$. This result was expected as each of the strategies have significantly different methods of choosing decision options. The following subsections will describe how the accuracy of each strategy is differentially affected by the measures of incomplete information.

Table 5.2: Average accuracy of, and effort required when using, decision strategies for each level of each measure of incomplete information.

Incomplete Information	**Tasks**	**Accuracy**				**Effort (EIP Count)**			
Total Info. Count	**Tasks**	**WADD**	**EW**	**TTB**	**Tallying**	**WADD**	**EW**	**TTB**	**Tallying**
12.5%	522,240	63.7%	63.7%	59.1%	57.9%	32.0	24.0	18.3	12.5
25.0%	1,827,840	68.8%	66.9%	63.4%	61.2%	32.0	24.0	16.3	15.0
37.5%	3,655,680	73.1%	71.4%	66.8%	64.1%	32.0	24.0	15.0	17.5
50.0%	4,569,600	77.1%	74.1%	69.8%	66.8%	32.0	24.0	14.2	20.0
62.5%	3,655,680	81.1%	78.0%	72.9%	69.6%	32.0	24.0	13.7	22.5
75.0%	1,827,836	85.4%	80.7%	76.2%	72.6%	32.0	24.0	13.5	25.0
87.5%	522,232	90.6%	84.8%	80.2%	76.1%	32.0	24.0	13.6	27.5
100.0%	65,276	100.0%	87.8%	84.8%	80.9%	32.0	24.0	14.1	30.0
Option Imbalance	**Tasks**	**WADD**	**EW**	**TTB**	**Tallying**	**WADD**	**EW**	**TTB**	**Tallying**
100%	130,560	77.0%	73.8%	57.0%	54.7%	32.0	24.0	12.3	20.0
75%	1,044,480	77.1%	74.5%	61.9%	58.5%	32.0	24.0	13.3	20.0
50%	3,655,676	77.1%	73.9%	66.8%	63.6%	32.0	24.0	14.1	20.0
25%	7,311,352	77.1%	74.7%	70.9%	68.0%	32.0	24.0	14.7	20.0
0%	4,504,316	77.4%	74.2%	72.8%	70.1%	32.0	24.0	14.8	20.1
Cue Balance	**Tasks**	**WADD**	**EW**	**TTB**	**Tallying**	**WADD**	**EW**	**TTB**	**Tallying**
100%	65,276	100.0%	87.8%	86.2%	80.9%	32.0	24.0	14.1	30.0
75%	783,352	89.2%	83.7%	80.2%	75.9%	32.0	24.0	14.2	26.7
50%	3,525,116	82.5%	78.9%	74.5%	71.5%	32.0	24.0	14.3	23.3
25%	7,050,240	76.9%	74.4%	68.9%	66.9%	32.0	24.0	14.3	20.0
0%	5,222,400	71.8%	69.7%	63.5%	62.3%	32.0	24.0	14.2	16.8

Table 5.3: Average accuracy of, and effort required when using, decision strategies for each level of each task parameter: dispersion and dominance.

Environmental Parameter	Tasks	Accuracy				Effort (EIP Count)			
Dispersion Level	**Tasks**	**WADD**	**EW**	**TTB**	**Tallying**	**WADD**	**EW**	**TTB**	**Tallying**
Low	391,680	62.1%	62.5%	55.4%	55.9%	32.0	24.0	16.9	20.0
Medium	12,586,787	76.4%	73.9%	69.2%	66.3%	32.0	24.0	14.8	20.0
High	3,667,917	81.3%	77.2%	73.7%	70.1%	32.0	24.0	13.2	20.0
Dominance	**Tasks**	**WADD**	**EW**	**TTB**	**Tallying**	**WADD**	**EW**	**TTB**	**Tallying**
Non-	11,676,946	73.9%	69.1%	68.3%	63.3%	32.0	24.0	14.3	20.0
Weak	4,308,479	83.0%	85.1%	71.8%	73.7%	32.0	24.0	15.2	20.0
Strong	660,959	97.1%	98.0%	85.0%	86.9%	32.0	24.0	13.7	20.0

Table 5.4: Analysis of effect of strategy, total information, dominance, and disperson on effort required to perform TTB.

Phase 1: Analysis 1	DF	F
Strategy	3	4,030,019*
Total Information	7	614,309*
Strategy×Total Information	21	629,299*
Total information×Dominance	14	2,289*
Strategy×Total Information×Dominance	42	2,289*
Phase 1: Analysis 2	**DF**	**F**
Strategy	3	136,935*
Total Information	7	45,839*
Strategy×Total Information	21	20,284*
Total information×Dispersion	14	7,495*
Strategy×Total Information×Dispersion	42	7,495*

DFE for all analyses: 6.66E7; *p$<$0.0001

The effect of strategy on the effort required to choose options was analyzed during the first phase of effort analysis with the two, three-way ANOVA calculations with the predictors of: (1) strategy, total information, and dominance, and (2) strategy, total information, and dispersion. Both calculations showed that all main effects, two-way interactions, and three-way interactions were significant except for dispersion or dominance as main effects and dispersion or dominance two-way interactions with strategy; see Table 5.4.

The three-way interactions, though seemingly complex, were expected because of the differential effect of total information, dispersion, and dominance, across the four strategies' effort requirements as explained in Table 5.1. The effort required to make a decision using the WADD or EW strategies is, by definition, not affected by the total information. The number of EIPs required to use either strategy for a two-option, four-cue decision remains constant at 34 and 26 EIPs, respectively, for all scenarios. This invariance to total information occurred because both strategies perform their standard processes regardless of how many attribute scores are known – by rule, unknown cue values are estimated to be 50. Conversely, for Tallying and TTB, steps 2 and 3 in their decision making processes are dependent on the number of attribute scores known so their effort requirements are affected by total information. Lastly, the effort requirements of WADD, EW, and Tallying, by definition, are not affected by dispersion or dominance because they have no mechanism in the middle of the decision making process to choose an option based on some intermediate option scores. The discrimination mechanism of TTB enables multiple decisions to be made within a single decision making process such that if the decision options' cue values have high dispersion (very different) or are strong dominant (one option has better attribute scores for all attributes) then TTB tends to select an option early, altering the amount of effort required.

5.2.2 Total Information

Total information was analyzed within each strategy and had a significant effect on the accuracy of each of the strategies: WADD ($\chi^2(7, 1.66\text{E}7) = 357,907, p < 0.0001$), EW ($\chi^2(7, 1.66\text{E}7) = 204,428, p < 0.0001$), TTB ($\chi^2(7, 1.66\text{E}7) = 192,136, p < 0.0001$), and Tallying ($\chi^2(7, 1.66\text{E}7) = 124,694, p < 0.0001$). As was expected, for each strategy there was a positive correlation between total information and mean accuracy, as total information increased so did accuracy: WADD ($\rho = 1.0000, p < 0.0001$), EW ($\rho = 1.0000, p < 0.0001$), TTB ($\rho = 1.0000, p < 0.0001$), and Tallying ($\rho = 1.0000, p < 0.0001$). The results in Table 5.2 match the conclusion made by Martignon and Hoffrage (2002) that as total information decreases, the differences in accuracy between the heuristic and the analytic strategies decrease. Specifically, as the number of known cue values decreased from all cue values known (100% total information) to one cue value known (12.5% total information), the difference in the mean accuracy between WADD and TTB decreased from 13.8% to 5.8%, and between WADD and Tallying the difference decreased from 19.1% to 5.8%.

Total information had a more complicated effect on the amount of effort required by each strategy. As shown in Fig. 5.2, the two heuristic strategies require differential information depending on the cue scores resulting in changes in effort requirements whereas the effort requirements for the two analytic strategies are invariant to total information. From the second phase ANOVA calculation, total information had a significant effect on the amount of effort required by Tallying $(F(7, 1.66\text{E}7) = 2.88\text{E}7, p < 0.0001)$ and TTB $(F(7, 1.66\text{E}7) = 5,016, p < 0.0001)$. The positive linear correlation (Fig. 5.2) between total information and the EIP count for Tallying is nearly monotonic ($\rho = 0.9596, p < 0.0001$) as a result of the strategy comparing more cue scores to the cutoff score as total information increases. The effort required to perform TTB has only a slightly negative linear correlation with total information ($\rho = -0.0250, p < 0.0001$).

The third phase of analyses of effort requirements focused on TTB by using two, two-way ANOVA calculations with TTB effort as the response variable. The first calculation included total information and dominance as the independent variables and showed that total information was a significant main effect $(F(7, 1.66\text{E}7) = 2,368, p < 0.0001)$, dispersion was not a significant main effect, but the interaction term was significant $(F(14, 1.66\text{E}7) = 2,362, p < 0.0001)$. Similarly, the second two-way ANOVA calculation included total information and dispersion as the independent variables and showed that total information was a significant main effect $(F(7, 1.66\text{E}7) = 5,672, p < 0.0001)$, dominance was not a significant main effect, but the interaction term was significant $(F(14, 1.66\text{E}7) = 7,737, p < 0.0001)$. Figure 5.3 shows the two-way effects of total information and dispersion and total information and dominance on the effort requirements.

In both Fig. 5.2 and Fig. 5.3 total information has a non-linear effect on the effort requirements of TTB. In Fig. 5.2, the highest average effort required for any total information less than 87.5% (12.8 EIPs) occurred when total information was at its lowest, 12.5%. As total information increased to 50%, the average effort decreased to its minimum of 12.1 EIPs, then increased from 50% to 100% total information to 14.1 EIPs. In Fig. 5.3, where the total information effects are separated by levels of dispersion and dominance, the minimum effort requirements for total information occur at 25% for low dispersion, 37.5% for medium dispersion, 100% for high dispersion, 50% for non-dominant, 37.5% for weak

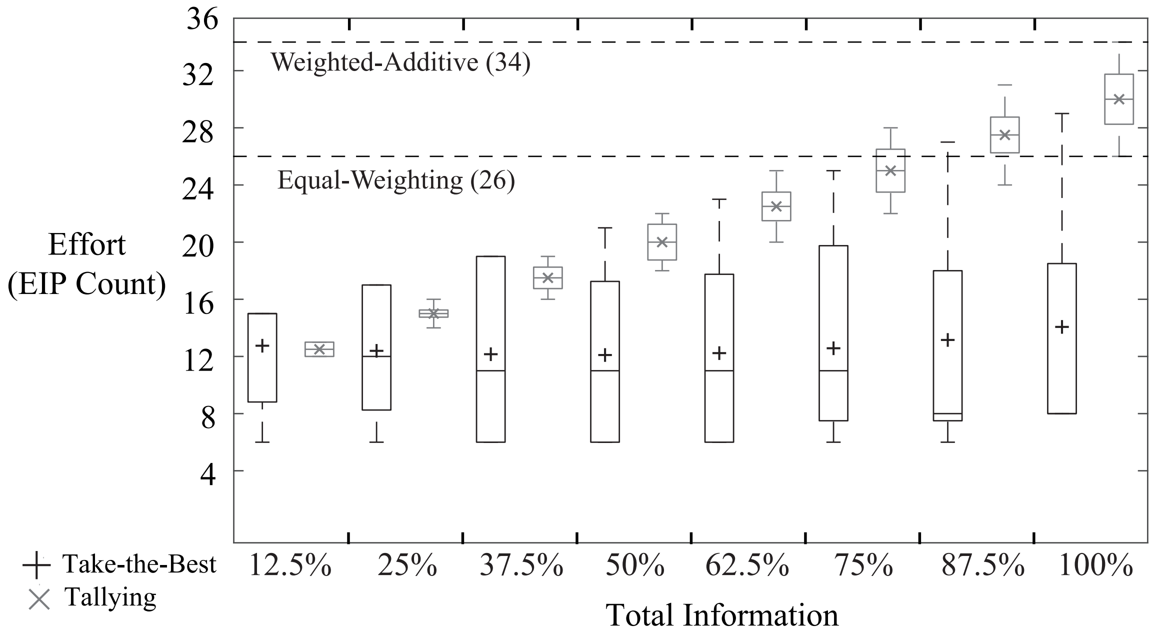

Figure 5.2: Effect of total information on the distribution of effort requirements for decision strategies. The markers indicate the mean EIP count.

dominant, and 75% for strong dominant. These non-linear results emphasize the fact that the effort for TTB is directly dependent on how many attributes it must search through until one discriminates. In strong dominant or high dispersion scenarios, the likelihood of an attribute discriminating is much higher than non-dominant or low dispersion scenarios enabling TTB to make decisions earlier. In summary, these results show that as total information increases, there is a point where TTB, on average, does not benefit from the additional attribute score information because it has already made a decision (Martignon and Hoffrage, 2002).

Examining the effort requirements by total information in Fig. 5.2 reveals that, on average, heuristics often require less effort than analytic strategies. However, the amount of effort heuristics require varies considerably. The variance in the effort required for Tallying increases from 0.25 to 2.00 as the total information increases from 12.5% to 100%. The variance is caused by Step 4 of Tallying (see Table 5.1) in which the strategy adds the number of positive cue values. Since the range of the potential number of cue values which can

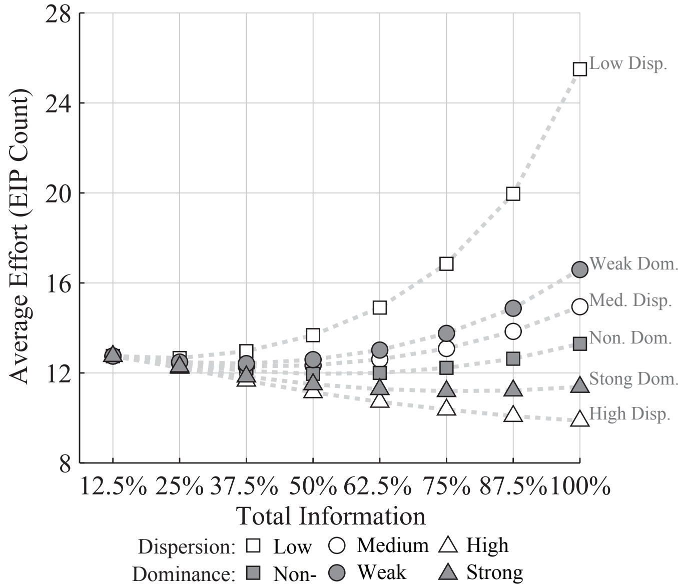

Figure 5.3: Effect of total information, dispersion, and dominance on the effort requirements of TTB.

be positive increases with increasing total information, the variability increases. Specifically, the range between the maximum and minimum number of EIPs required for Tallying at any one level of total information is equivalent to the number of cue values known. For example, with 3 cue values known (total information of 37.5%) the effort required to perform Tallying ranges from 16 to 19, whereas with 8 cue values known (total information of 100%) the effort required ranges from 26 to 34.

The range in Tallying effort requirements is relatively well-bounded, but the range for TTB is noticeably larger: from 12.5% to 100% total information, the range increases from 9 to 21. An explanation for this effect comes from the discrimination mechanism. Taking the example where only one cue score is known and it is from the highest weighted cue, if the cue score is higher than the cutoff score, then the discrimination criterion is met and TTB will have required 6 EIPs. However, if the cue score is less than the cutoff score and does not discriminate, TTB will search through the remaining cues, finding each remaining cue score unknown, resulting in a maximum possible 15 EIPs. This large range indicates that, in some decision tasks, knowing just the amount of total information and nothing else about the specifics of the strategy is insufficient to predict the effort required. For example, though TTB may, on average, require much less effort than analytic strategies or even Tallying in many scenarios, how much less will depend largely on how much information is needed to discriminate between the options. These findings suggest that in terms of effort, heuristics are only particularly effective when the number of cues required for discrimination is low. Experienced personnel tend to use heuristics, while novices tend to use analytic strategies (Garcia-Retamero and Dhami, 2009; Gigerenzer and Gaissmaier, 2011). Our results present a possible effort-based explanation of this phenomenon: experts limit their effort (and may produce decision more quickly) by relying on fewer, more important, cues whereas novices may use multiple cues. Since novices may not have preferences for cues or know which is the most important – information that experts would have gained through experience – they may spend more time evaluating cues or have a wider threshold for discrimination (several cues must discriminate, not just one), thus leading to a prolonged search and greater effort (Dieckmann and Rieskamp, 2007; Karelaia, 2006). Thus the presumed benefits of reduced effort obtained through heuristics would vanish.

5.2.3 Option Imbalance

Option imbalance was analyzed within each strategy and included in this study to augment total information by examining how the distribution of incomplete information could affect the accuracy of, and effort required to perform, decision strategies. This analysis examined the information generated by varying total information and grouped it differently to specifically focus on the distribution of information across the cues. The results show that option imbalance had a statistically significant effect on accuracy of each of the decision making strategies: WADD ($\chi^2(4, , 1.66\text{E}7) = 164, p < 0.0001$), EW ($\chi^2(4, , 1.66\text{E}7) = 894, p < 0.0001$), TTB ($\chi^2(4, , 1.66\text{E}7) = 89{,}623, p < 0.0001$) and Tallying ($\chi^2(4, , 1.66\text{E}7) = 83{,}000, p < 0.0001$). As shown in both Table 5.2 and the white markers in Fig. 5.4, average accuracy for WADD and EW decreases only 0.4% and 0.6% as option imbalance increases from 0% to 100%. However, TTB and Tallying with 0% option imbalance have an accuracy of 72.8% and 70.1%, respectively, while with 100% option imbalance (all cue values are known for one option only and no cue values are known about the other option) their respective accuracy is 57.0% and 54.7%. As expected, option imbalance had no statistically significant effect on the effort of the analytic strategies, WADD and EW. Conversely, option imbalance did have a statistically significant effect on the effort of the heuristic strategies: TTB $(F(4, 1.66\text{E}7) = 21,249, p < 0.0001)$ and Tallying $(F(4, 1.66\text{E}7) = 1,311, p < 0.0001)$. It must be noted however, that the size of the effect of option imbalance on the effort of Tallying is likely very small. The average effort required to perform Tallying is 20.0 for all levels of option imbalance except 0% when the average effort increases by 0.1.

To better understand the effect of option imbalance on the accuracy of decision strategies, a new metric was introduced, information bias: the percentage of scenarios in which a decision strategy selects the option with the most total information. In two-option decision scenarios with option imbalance (i.e. one option has more information than the other), a strategy with an average information bias of 50% indicates that in 50% of tasks, the strategy selected the option with more information and in the other 50% of scenarios the strategy selected the option with less information. Since this simulation examined all combinations of option imbalance, in half of the scenarios with option imbalance, the correct option will also be the one with more information and in the other half of those scenarios, the correct option will have less information than the incorrect option. Therefore, in this study, an information bias of 50% indicates that a strategy has no preference toward selecting options with more or less information.

Information bias was analyzed within each decision strategy and had a statistically significant effect on the accuracy: WADD $(\chi^2(1, 1.66\text{E}7) = 36, p < 0.0001)$, EW $(\chi^2(1, 1.66\text{E}7) = 17, p < 0.0001)$, TTB $(\chi^2(1, 1.66\text{E}7) = 253,293, p < 0.0001)$, and Tallying $(\chi^2(1, 1.66\text{E}7) = 239,708, p < 0.0001)$. Although information bias has a statistically significant effect on the

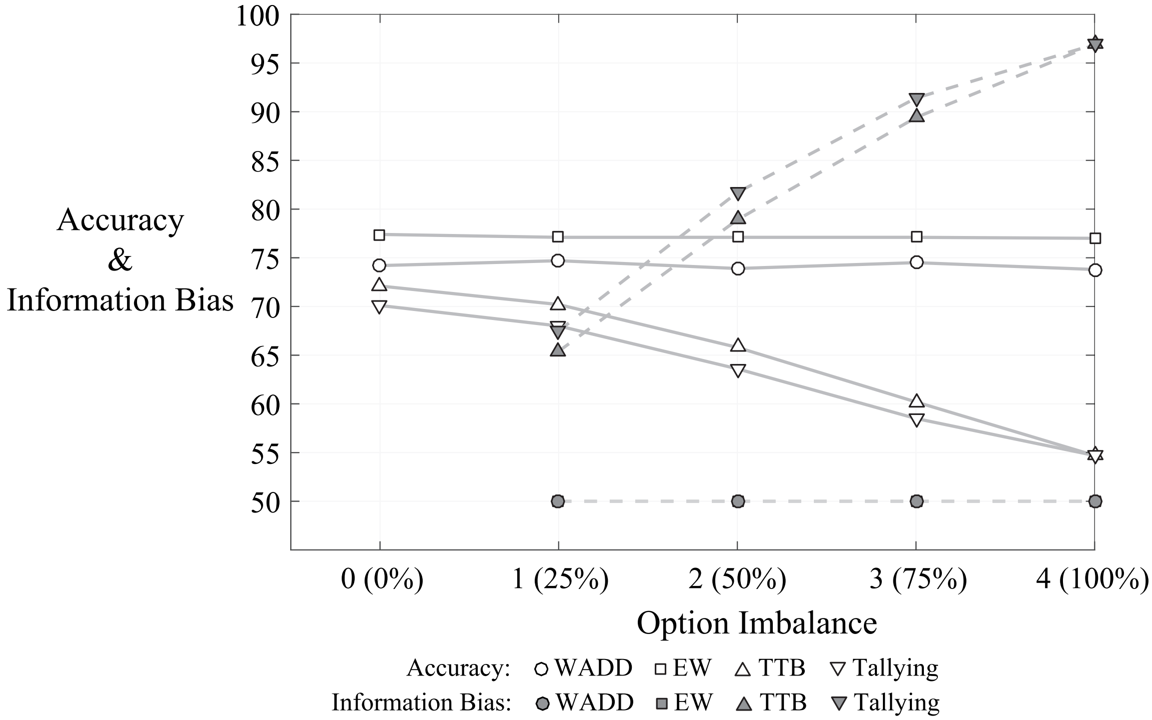

Figure 5.4: Effect of option imbalance on the accuracy and information bias of decision strategies. Accuracy data is indicated by the white markers whereas information bias data is indicated by the gray markers. Since this study examines a two-option decision task, an information bias of 50% indicates that a strategy has no preference toward selecting options with more information.

accuracy of WADD and EW, the size of the effect in these results is negligible. Figure 5.4 shows that WADD and EW had no preference for options with more or less information (50% information bias) and constant accuracy across all levels of option imbalance. In contrast, the increase in information bias to values above 50% for the heuristic strategies correlates with the decrease in their accuracy. This result indicates that when option imbalance is high, TTB and Tallying will choose the option with more total information almost independently of selecting the option with the highest overall score.

Since option imbalance measures the difference in the total information available for each decision option, it approximates the number of pairwise comparisons that TTB and Tallying have to make between known and unknown option cue values. When comparing known and unknown cue values, the discrimination mechanism is biased in favor of decision options with more information because unknown scores are treated as if they were

Table 5.5: Statistical significance of the effect of option balance on the accuracy of each decision strategy within each paired level of total information.

	WADD	**EW**	**TTB**	**Tallying**
$\boldsymbol{\chi^2(2, 2.35E6)}$ *Pair(12.5%v25%)*	4,687*	1,783*	13,234*	4,112*
$\boldsymbol{\chi^2(4, 8.23E6)}$ *Pair(37.5%v50%)*	17,288*	7,450*	61,746*	51,367*
$\boldsymbol{\chi^2(3, 5.48E6)}$ *Pair(62.5%v75%)*	16,140*	5,286*	34,775*	37,007*
$\boldsymbol{\chi^2(1, 5.88E5)}$ *Pair(87.5%v100%)*	12,161*	414*	1,379*	779*

*p<0.0001

negative scores. As a result, low option imbalance indicates that fewer of these known-unknown comparisons will be made, thus decreasing the likelihood of TTB or Tallying being biased towards the option with more information.

5.2.4 Total Information and Option Imbalance

Further analysis of the accuracy of decision strategies confirmed that the accuracy of heuristics is a function of option imbalance and total information. Although option imbalance is a function of total information, within each pair of levels of total information (12.5% and 25%, 37.5% and 50%, 62.5% and 75%, and 87.5% and 100%), the effect of option imbalance on the accuracy of each strategy was statistically significant (see Table 5.5). For example, for all scenarios in which total information was either 37.5% or 50%, option imbalance has a statistically significant effect on the accuracy of TTB (Pair(37.5%v50%),$\chi^2(4, 8.23\text{E}6), p < 0.0001$); thus supporting the results shown in Fig. 5.5-TTB that decreasing option imbalance while keeping total information relatively constant, increases accuracy.

The results support the widely-documented finding that, when using analytic strategies, accuracy increases proportionally with the amount of information about the scenario (assuming all the information is both correct and relevant). Figure 5.5 uses a heatmap format

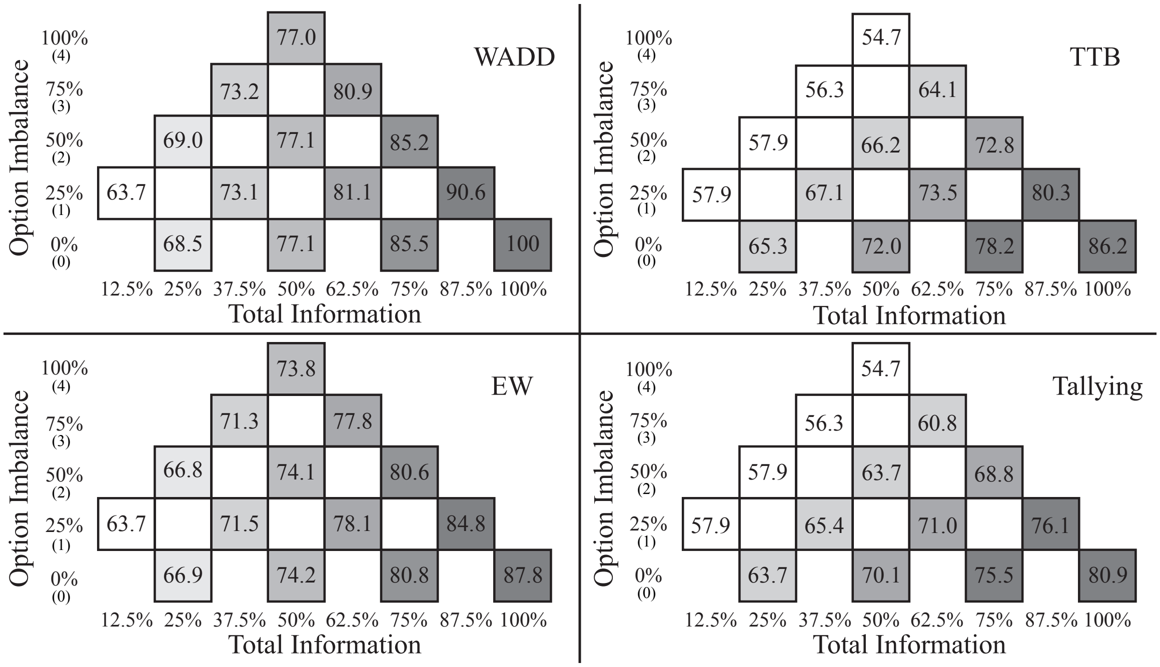

Figure 5.5: Heatmap of the effect of combinations of total information (x-axis) and option imbalance (y-axis) on the accuracy (intensity of gray shading). Empty boxes indicate that the combination of option imbalance and total information is not possible with this simulation. For example, for a total information of two, there are only two levels of option imbalance possible: zero (both cue values are known about one option) or two (one cue value is known about each option).

where total information in on the x-axis, option imbalance is on the y-axis, the values inside the boxes indicate the mean accuracy of all scenarios with that combination, and the intensity of the shading mirrors the level of accuracy (darker is better). Figure 5.5 shows that for all strategies, the combination of decreasing option imbalance and increasing total information produces higher accuracy.

More generally, for any individual level of total information, a decrease in option imbalance resulted in an increase, or equivalent, accuracy (except for 25% total information for WADD). This result is particularly applicable to the heuristic strategies in Fig. 5.5. For example, if 4 cue values are known (50% total information), when the information is distributed such that option imbalance is 100%, the accuracy of TTB is 54.7% (top row); however, if the available information is redistributed such that option imbalance is 0%, then the accuracy increases to 72.0% (bottom row).

The results in Fig. 5.5 seem to indicate a trade-off between total information and option imbalance for heuristic strategies. Movement along the diagonal in which option imbalance decreases and total information decreases may result in an increase in accuracy: 'better' option imbalance and 'worse' total information may result in 'better' accuracy. This suggests that there may be scenarios in which decision makers employing TTB or Tallying could increase their accuracy by removing correct and relevant cue values to decrease option imbalance – trading known information away for a decrease in option imbalance to increase accuracy. This is exemplified by observing the average accuracy results for TTB and Tallying at 72.8% and 68.8%, respectively, when six cue values are known (75% total information) and option imbalance is 50%. By removing one cue value from the decision task to decrease option imbalance to 25% and decrease total information to 62.5%, the accuracy of TTB and Tallying increase to 73.5% and 71.0%.

5.2.5 Cue Balance and Total Information

Cue balance is a measure of the distribution of known and unknown information defined as the number of cues which have known values for all options (see Sec. 4.3.4). Analyzing the accuracy and effort of decision making strategies with respect to cue balance in Table 5.6, shows that TTB and Tallying are affected by cue balance while WADD and EW are only affected by total information. The trends for the overall results show that decision tasks with the highest cue balance for a given level of total information indicate when TTB and Tallying will have the highest accuracy for that given level of total information. Furthermore, tasks with only balanced cues have higher accuracy than some tasks with more information but fewer balanced cues.

The exact values provided in Table 5.6 are visualized in Fig. 5.2.5. Defining the combinations of total information and cue balance in coordinates of (total information, cue balance), the tasks with the maximum number of balanced cues for each level of total information, (2,1), (3,1), (4,2), (5,2), (6,3), always have the highest accuracy for TTB and

Table 5.6: Accuracy and EIP count for each decision making strategy for all 14 combinations of total information and cue balance values.

Total Information	Cue Balance	WADD Acc.	WADD Eff.	EW Acc.	EW Eff.	TTB Acc.	TTB Eff.	Tallying Acc.	Tallying Eff.
1	0	63.7%	*32*	63.7%	*24*	57.9%	*18.1*	57.9%	*12.5*
2	0	69.0%	*32*	66.8%	*24*	61.8%	*15.6*	60.8%	*15.0*
3	0	73.2%	*32*	71.3%	*24*	64.6%	*13.5*	63.1%	*17.5*
4	0	77.0%	*32*	73.8%	*24*	66.6%	*11.8*	64.9%	*20.0*
2	1	67.1%	*32*	67.1%	*24*	63.7%	*19.2*	63.7%	*15.0*
3	1	72.9%	*32*	71.6%	*24*	67.0%	*16.4*	65.4%	*17.5*
4	1	77.1%	*32*	74.1%	*24*	69.2%	*14.2*	67.0%	*20.0*
5	1	80.9%	*32*	77.8%	*24*	70.7%	*12.3*	68.4%	*22.5*
4	2	77.0%	*32*	74.7%	*24*	72.6%	*17.3*	70.2%	*20.0*
5	2	81.3%	*32*	78.3%	*24*	74.2%	*14.9*	71.1%	*22.5*
6	2	85.2%	*32*	80.6%	*24*	75.2%	*12.9*	72.1%	*25.0*
6	3	86.6%	*32*	81.3%	*24*	79.8%	*15.5*	75.6%	*25.0*
7	3	90.6%	*32*	84.8%	*24*	80.3%	*13.5*	76.1%	*27.5*
8	4	100.0%	*32*	87.8%	*24*	86.2%	*14.1*	80.9%	*30.0*

Tallying for that level of total information. Conversely, WADD and EW are nearly invariant to cue balance as the accuracy is constantly increasing as a function of total information.

Table 5.6 and Fig. 5.2.5 also illustrate the importance of cue balance relative to total information. For TTB and Tallying, for many combinations of total information and cue balance, decreasing cue balance and increasing the total information results in a slight decrease in accuracy. This occurs for combinations (3,1) and (4,2) relative to (4,0) and (5,1), respectively.

With respect to total information, increasing total information without increasing cue balance has diminishing increases in accuracy for TTB, and Tallying. For WADD and EW, increasing total information without increasing cue balance provides consistent increases in accuracy. For TTB, increasing the total information from (2,1) to (3,1) increases the accuracy by 3.3% while increasing the total information from (6,3) to (7,3) increases the accuracy by 0.5%. Another perspective shows that as total information increases, the importance of increasing cue balance increases. Increasing the total information in such a

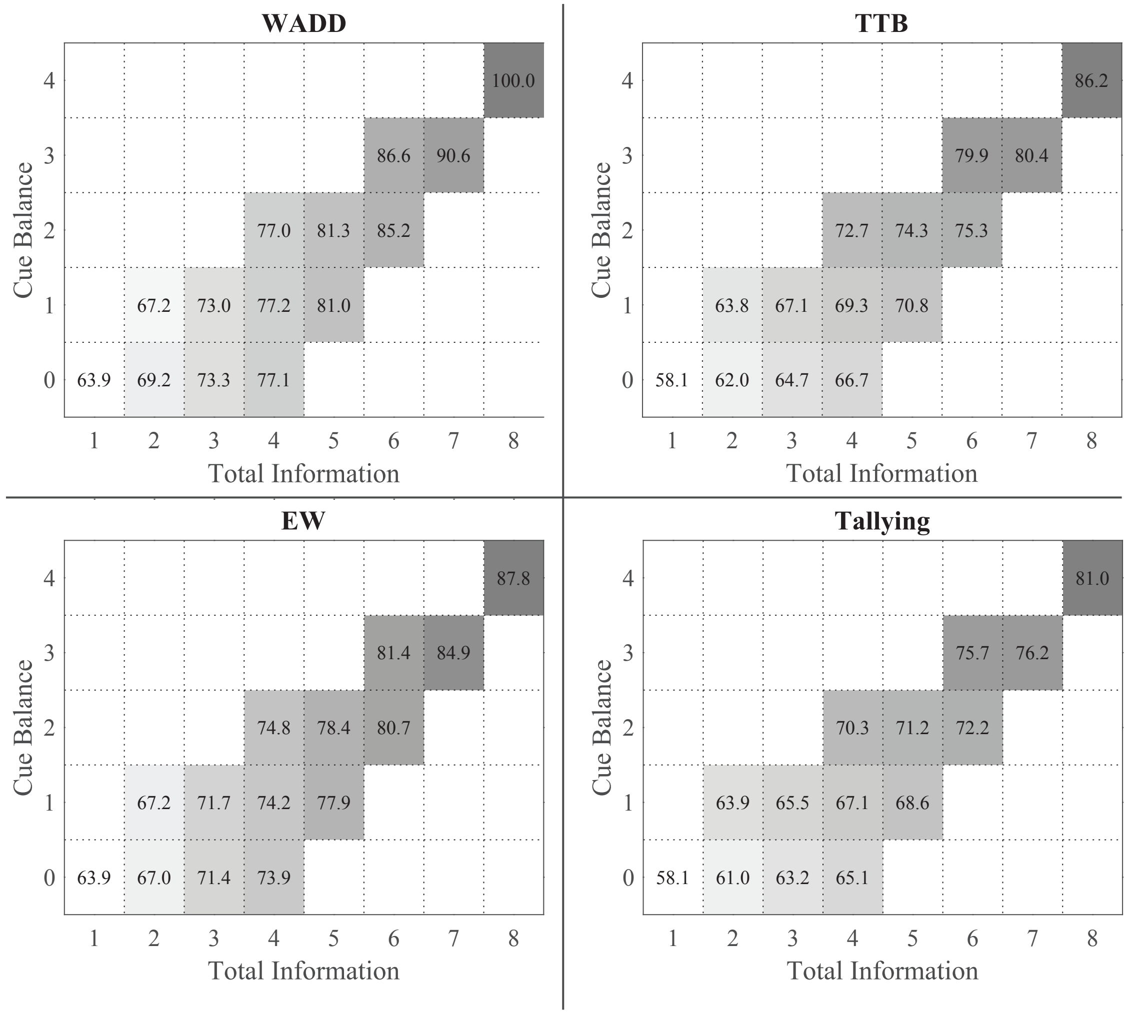

Figure 5.6: Accuracy of each decision making strategy for each combination of total information and cue balance.

way that the number of balanced cues increases, results in significantly higher increases in accuracy than increasing total information in such a way that cue balance remains the same. For example, for TTB, from (5,2) to (6,2) accuracy increases 1.0% whereas from (5,2) to (6,3) accuracy increases 5.6%. This suggests that different cue values are more valuable to search for during the information search process of decision makers. A search technique maximizing the cue balance at each step of the process would produce higher accuracy than any other technique for this 2-option decision task.

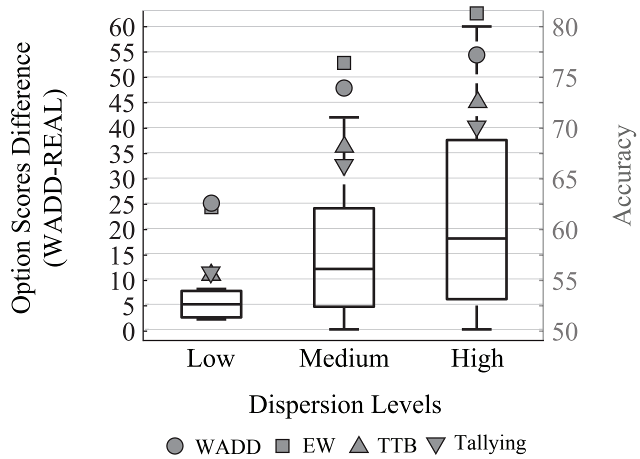

Figure 5.7: Effect of dispersion level on the difference between the WADD-REAL option scores and the accuracy of decision strategies. The box plots describe the distribution of the difference between the two option scores when the dispersion level is low, medium, or high. The gray markers indicate the accuracy of the decision strategies for each level of dispersion.

5.2.6 Dispersion

The converging nature of the accuracy of the analytic and heuristic strategies as the dispersion decreased suggests that, of the three dispersion levels, low dispersion scenarios offer the best accuracy-effort trade-off for heuristics. Comparing low and high dispersion scenarios, the difference between the most accurate analytic strategy and the most accurate heuristic increased from 6.6% to 8.6%. Additionally, Fig. 5.7 shows that the median of the differences between the overall option scores with low dispersion is 5 whereas the median for high dispersion is 18. The differences indicate that in scenarios with low dispersion, the median loss for choosing the incorrect option is 5 whereas the median loss for choosing the incorrect decision in a scenario with high dispersion is 18 yet could be as high as 60. Lastly, TTB and Tallying, on average, require less effort (14.2 and 20 EIPs, respectively) compared to WADD and EW (34 and 26 EIPs, respectively) in scenarios with low dispersion. Therefore, of the three dispersion levels, low dispersion scenarios have (1) the smallest difference between the accuracies of heuristic and analytic strategies, (2) the smallest median loss for choosing the incorrect decision, while retaining the effort requirement in favor of heuristics, and (3) heuristics maintain their decreased effort required (especially TTB). Nonetheless, in some applications a decreased accuracy of 6% in which the difference in overall score between the two options is 5 may be too much, even given the reduced effort requirement (e.g. aiming at a target). In others, particularly iterative decisions, the impact is not as severe. Therefore the question of its significance becomes more a matter of the operator's objectives and measures of significance, factors which were beyond the scope of this investigation.

5.2.7 Dominance

The effect of dominance on the accuracy of each strategy was analyzed within each strategy and shown to be statistically significant: WADD $(\chi^2(2, 1.66\text{E}7) = 376,050, p < 0.0001)$, EW $(\chi^2(2, 1.66\text{E}7) = 741,942, p < 0.0001)$, TTB $(\chi^2(2, 1.66\text{E}7) = 147,699, p < 0.0001)$, and Tallying $(\chi^2(2, 1.66\text{E}7) = 301,933, p < 0.0001)$. In contrast, the only strategy whose effort could be affected by dominance was TTB and in that case, dominance was shown to have a statistical significant effect $(F(2, 1.66\text{E}7) = 31,615, p < 0.0001)$. TTB is the only strategy whose effort requirements were affected by dominance because an increase in dominance by definition increases the likelihood that an attribute will discriminate, and thus reduces the amount of attributes TTB needs to search through. The increase in likelihood of discrimination is due to strong dominant scenarios having no attributes in which both options have the same score whereas in non-dominant scenarios, attributes with equivalent scores for both options is possible. On the other hand, the other strategies' effort requirements are not affected because dominance is a description of the attribute scores and thus aggregates all the levels of total information.

Across all levels of dominance, the analytic strategies had higher average accuracy than the heuristic strategies. However, which strategy within each type had the higher accuracy depended on the presence of dominance and the use of attribute weights. In non-dominant scenarios, the two strategies that used attribute weights to choose options, WADD and TTB, had approximately 5% higher accuracy than their respective analytic and heuristic counterparts, EW and Tallying. Conversely, in scenarios where a decision option was weakly or strongly dominant over the other option, EW had 2.1% and 0.9% higher accuracy than WADD while Tallying had 1.2% and 1.4% higher accuracy than TTB.

The use of attribute weights increased the accuracy of WADD and TTB relative to EW and Tallying in non-dominant scenarios because of the choice of WADD-REAL to determine correct decisions. In non-dominant scenarios, different attributes favor different options, so weighing attributes in a manner similar to reality is necessary to make accurate decisions. Since EW and Tallying do not use attribute weights in their decision making processes, they are less likely to combine the attribute scores in a way that matches the relative attribute importance inherent in WADD-REAL. If the correct decision was determined by an EW or Tallying strategy with full information, WADD and TTB would likely have lower accuracy than EW or Tallying. By definition, with full information in weak or strong dominant scenarios, the attribute weighting is unnecessary as any combination of the attributes will indicate the same correct option. Therefore, the slight deviation between WADD and EW, and TTB and Tallying, in weak or strong dominant scenarios is due to the changes in incomplete information – not whether the strategies used attribute weights.

This effect of varying levels of total information within different levels of dominance especially affects the effort requirements of TTB as shown in the third phase of analyses of effort requirements (see Sec. 5.2.2 and Fig. 5.3). For each level of total information except for 12.5% total information (where the three levels of dominance require the same average effort, 12.75), the order of effort requirements from least to most are strong, non-, then weak dominant scenarios. More specifically, as the total information increased, the difference in average effort requirements increased from less than 2 EIPs at 50% total information, to more than 5 EIPs at 100% total information. There are two separate reasons for the differences between the effort requirements of strong dominance and non-dominant scenarios, and the high effort requirements of weak dominance. Strong dominant scenarios require less effort than non-dominant scenarios because of strong dominant scenarios have a much higher likelihood of discriminating than non-dominant scenarios, increasing the chance that TTB will discriminate between options at earlier searched attributes. The high effort requirements for weak dominant scenarios is due to the large numbers of low dispersion scenarios being weak-dominant. Low dispersion scenarios have the highest effort requirements of the three dispersion levels because, as stated in Sec. 5.2.6, the two attributes are so similar that even with increased information TTB cannot discriminate easily. Therefore, even though the scenarios are weak dominant, the low levels of dispersion still make it difficult for TTB to discriminate, thus increasing the effort requirements.

5.3 Discussion

The results indicate how the distributions of incomplete information within decision tasks affect the accuracy and effort of decision making strategies. Two principle results were found: 1) context features matching naturalistic decision settings result in heuristic strategies being closest in accuracy to analytic strategies; 2) the variability in the distribution of the effort requirements of the heuristic strategies for each level of total information indicates that the effort requirements of heuristics may not always be as favorable as prior studies have shown; and 3) the tradeoffs between option imbalance and total information, and between cue balance and total information, suggests new insight for DSS information acquisition and restriction. These results occur even within an environment with an accuracy measurement biased toward analytic strategies. The next computer studies, therefore, leverage real-world datasets to examine the potential for these tradeoffs to apply to more general environments.

5.3.1 Heuristics in Naturalistic Decision Contexts

In our simulation, context features that resemble naturalistic decision settings commonly faced by decision makers (Orasanu and Connolly, 1993) resulted in the minimum difference between the accuracies of analytic and heuristic strategies. The naturalistic context features occurred with low total information, low dispersion, or non-dominant options. As the three context features of total information, dispersion, and dominance approached the low end of their possible values, the accuracy of analytic and heuristic strategies became more and more similar. These results extend the studies of Payne et al. (1990) who simulated decision situations to examine the impact of dispersion of the weights, while confirming their results that the presence of non-dominant options in the decision task resulted in a decrease in the difference in accuracy between analytic and heuristic strategies. Similarly, the results confirm the observation by Martignon and Hoffrage (2002) that in scenarios with low total information, the accuracy of TTB and Tallying will become close to the accuracy of WADD. While they were generated in naturalistic settings, our results provide initial, but not robust, support for the argument that heuristics can be ecologically rational in situations commonly encountered by decision makers (Todd et al., 2012) because they can produce greater accuracy than analytic strategies. With respect to performance on accuracy, our simulation was biased in favor of WADD. However, when the amount of information available, the dispersion between cue values, or the level of dominance among options decreased, the tested heuristic strategies performed almost as well as WADD in terms of accuracy. We believe this indicates that there is value in conducting additional studies of heuristic decision making strategies like TTB and Tallying under conditions of incomplete information when the standard of accuracy is not WADD-REAL, but a more objective reality.

5.3.2 Variability in Effort for Heuristics

The variability in the distribution of the effort requirements of the heuristic strategies for each level of total information highlights a more nuanced effort-accuracy trade-off than previously shown by the averaged results in Payne et al. (1990). The standard interpretation of the effort-accuracy trade-off is that people use heuristics to save effort while retaining sufficient accuracy as compared to analytic strategies. Contrary to the argument that heuristics save effort, analysis of the effect of total information on the distribution of effort requirements has shown that in some scenarios TTB and Tallying may not save much effort. The key is to use them judiciously in the correct decision context.

The level and variability of effort would be affected by changes to the assumption that all EIPs require equal time and effort. This assumption was based on prior similar work by Payne et al. (1990, 1996) in which assuming equal time and effort for each EIP produced almost identical conclusions as using time estimates for the EIPs. Though the variability in TTB and Tallying is due to the nature of the strategies, replacing counts of EIPs with time estimates for EIPs as a measure (Bettman et al., 1990) would likely result in a change in the variability. Two prominent EIPs in the decision strategies, `Compare`, used by Tallying and TTB, and `Add`, used by Tallying, take different amount of time to perform – 0.20 s for `Compare` and 0.84 s for `Add`.

The variability in effort highlights the importance of expertise in selection of decision making strategies by decision makers. Many studies have shown that the decisions of experienced professionals are more likely to be predicted by simple heuristics than WADD, whereas novice participants are better predicted by WADD (for a review see Gigerenzer and Gaissmaier, 2011). This study shows that the simple building blocks of TTB and Tallying, particularly the discrimination rule and cue-substitution rule (TTB only) cause variability in effort. These building blocks, which enable heuristics to generally achieve high accuracy with low effort, are not always beneficial and have the potential to become a detriment to accuracy and effort. This result suggests that novices may use analytic strategies due to their clear process and constant effort requirements which makes the strategies easier to implement and effort easier to predict. Heuristics, however, require expertise because knowing the important cues is not only required for high or sufficient accuracy but for minimizing effort as well – knowing when to abandon or adapt the strategy if the context features indicate that effort will not be saved (Simon, 1956; Klein, 1998).

Some caution must be taken with these results which violate the assumption that heuristics will always require significantly less information than analytic methods and that the variance in effort requirements of heuristics are significantly higher than analytic strategies. The assignment of EIPs to steps was based on the interpretation of the standard models of these decision making strategies. All EIPs were weighted equally, though, in reality, some EIPs may be more complex than others, requiring more time and effort (Payne et al., 1990). Furthermore, assignment of the EIPs requires decomposition of these decision strategies which may be open to interpretation. Consequently, different EIP weights or decompositions may yield different results.

While they did not focus on the variance of heuristic strategies, previous studies of decision making effort did not report similar variance in participants using lexicographic methods such as TTB (Mata et al., 2007; Rieskamp and Hoffrage, 2008). It is possible that the lack of variance in those studies may have been caused by the fact that real decision makers are unlikely to follow the formal processes of heuristic strategies (Hilbig, 2010). More studies of decision makers in situations of incomplete information will be necessary to determine the validity of these results.

5.3.3 Implications for Heuristic Information Acquisition and Restriction in Decision Support

The most important result is the analysis of the tradeoffs between option imbalance and total information, and cue balance and total information. The tradeoffs suggest methods for decision support systems (DSSs) to acquire and restrict information. Acquisition and restriction can mitigate information bias to increase the accuracy of heuristic strategies or decrease effort while keeping accuracy approximately the same, on average. Design for restriction of information in DSSs, which limits decision makers from performing certain actions, could use knowledge of the tradeoff between total information and option imbalance or cue balance to make specific pieces of information less salient for the decision maker. As discussed in Sec. 5.2.4, the accuracy of TTB and Tallying are almost constant when total information is decreased by one and option imbalance is decreased by one. As discussed in Sec. 5.2.5, the accuracy of TTB and Tallying increase when total information is decreased by one and cue balance is increased.

Although these analyses do not discriminate between the removal of information from the highest or lowest ranked cue, a cue value would be removed or hidden from the option with more information (option imbalance) or the cue without all values known (cue balance). Therefore, a DSS with knowledge of the relative importance of the cues and the distribution of known cue values could attempt to hide the lowest-ranked cue value of the option with more total information to increase the possibility of an accurate decision. Another possibility is for the DSS to not hide this information, but to make the information less salient or require the user to take an extra step to access this data. For example, the main page might display only the balanced cues and options, but could include a notation indicating that more information is available if needed.

Hiding available information is counter-intuitive, but it reformulates the decision task in a way to encourage the user to employ more efficient TTB and Tallying strategies. This reformulation is similar to restructuring the input to an algorithm. An example of the positive potential for restricting information is exemplified by Fig. 5.8 where cue values are not known for Option 1 - Cue 3, nor Option 2 - Cues 1 and 4. In this task, if there is no restriction on the part of the DSS then TTB and Tallying will have accuracies of 72.0% and 71.8%. However, if a DSS were to restrict information by removing the lowest-ranked cue's score (Option 1 - Cue 4) in order to decrease option imbalance and decrease the total information, the accuracy for TTB and Tallying will increase to 72.7% and 74.1% on average. This example is one of many in the simulation for which a decrease in total information resulting in a decrease in option imbalance will cause no change or a slight increase in accuracy for TTB and Tallying.

Questions regarding how the DSS would know which attribute is lowest ranked are answered by prior research showing that determining a sufficiently accurate rank of the

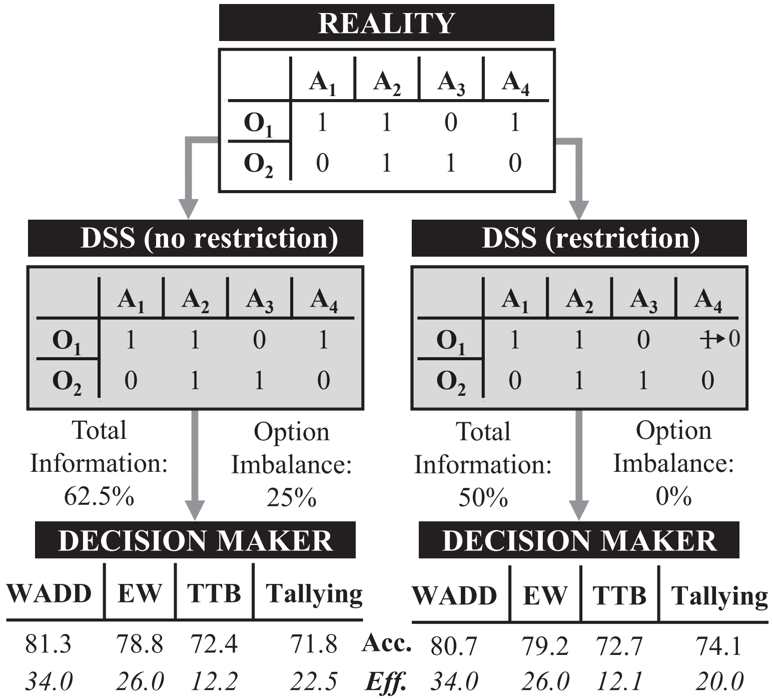

Figure 5.8: Effect of restriction of cue values by a decision support system (DSS) to reduce total information and option imbalance on the accuracy and effort of decision making strategies. The accuracy, and effort required, for each decision strategy are the averages for all combinations of cue values.

attributes would not be difficult. Katsikopoulos et al. (2010) showed that even when determining imperfect rank-orders of attributes from samples of information rather than perfect knowledge of all information, TTB was still able to outperform multiple-regression in some environments. Additionally, decision makers are known to learn rank-orders of attributes from others through teaching and imitation (Gigerenzer and Gaissmaier, 2011).

5.4 Conclusion

To better understand how to design decision support systems to accommodate a range of decision strategies, a simulation engine was designed to model the accuracy and required effort of two analytic (weighted-additive and equal-weighting) and two heuristic (take-the-best and tallying) decision strategies in over 16 million decision tasks. These scenarios had varied levels of total information available, option imbalance, cue balance, dominance, and dispersion.

Three major conclusions were drawn from the results. First, heuristic strategies are nearly as accurate as analytic strategies when decision task context features resemble naturalistic environments. The set of scenarios used in this simulation extend previous work studying these context features in isolation to show that in scenarios combining multiple naturalistic context feature levels, heuristic strategies can produce levels of accuracy that compare well with analytic strategies. Second, the effort requirements of heuristic strategies may not be as predictable as prior studies have shown. Increased information may sometimes result in more, rather than less, effort, and, surprisingly, even more effort than analytic strategies. This reversed effect is due to the fundamental mechanisms, the building blocks, of TTB and Tallying, but future studies should account for the costs associated with specific types of effort. Third, the tradeoffs between total information, option imbalance, and cue balance suggest that in certain instances, a DSS that supports heuristic strategies may be able to help the strategies achieve greater accuracy by reducing the amount of information they use. DSS designers can augment the standard analytic methods of information acquisition by acquiring or restricting specific cue values to minimize option imbalance or maximize cue balance.

CHAPTER 6

COMPUTER SIMULATION STUDY 2: HEURISTIC INFORMATION ACQUISITION AND RESTRICTION RULES FOR DECISION SUPPORT

The results in this section are a version of the following article: Canellas and Feigh, 2017.

How to handle incomplete information in decision support system (DSS) design is an open question. Particularly in terms of modifying the content presented to the decision maker (Feigh et al., 2012), are there situations when the DSS should acquire information for or restrict information from the operator? If so, how? The DSS could leverage analytic information acquisition methods requiring reliable assessments of probabilities, cue weights, and cue scores (Nelson, 2005, 2008; Meder and Nelson, 2012). However, there are multiple environmental issues with using these analytic methods in decision support (Katsikopoulos and Fasolo, 2006; Katsikopoulos et al., 2008). First, real-world problems can exhibit statistical dependencies which create computational issues in calculating probabilities and cue weights. Second, psychologically, an operator may not be able to provide accurate information to the DSS regarding the probabilities and cue weights. Third, the analytic process may seem too complex and too opaque to the operator causing the operator to be reluctant to use the DSS or accept its suggestions. Lastly, as these methods have focused on acquisition, they do not provide suggestions as to what information could be restricted from the operator.

Given these limitations, how can DSSs acquire or restrict information using methods that do not rely on information about the probabilities, cue weights, and cue scores? The promising avenue of research in the previous study (Chap. 5) examined how distributions of incomplete information affect decision making performance. Decreasing the difference in total known information between options (option imbalance) increases the accuracy of heuristic strategies and increasing the number of cues which have information known for all options (cue balance) increases the accuracy of both analytic and heuristic strategies. Based on this research, this study proposes four heuristic information acquisition and restriction rules that can guide DSSs in determining what information should be added or removed from a decision task: acquire or restrict information to decrease option imbalance, or to increase or maintain cue balance. Importantly the measures of option imbalance and cue balance only require knowledge of what information is known or unknown – no probabilities, no cue weights, no cue scores.

To determine if or when these four rules increase performance, they were simulated in 2-option decision tasks across 15 real-world datasets (Czerlinski et al., 1999) while measuring their effects on the accuracy of five decision making strategies spanning the range from analytic to heuristic. Three questions are addressed in the results: First, how often and how much do the four rules increase accuracy? To this end, the simulation measured the presence and magnitude of the changes in decision making accuracy using positive accuracy count (PAC, the percent of times in which the rules resulted in a positive change in accuracy), and average accuracy change (AAC, the average change in accuracy when using the rules), respectively. Second, how do the components of decision making strategies affect the effectiveness of the rules? Strategies are made of components that govern the type of cue scores, cue weights, and estimates of missing information (Chap. 3). Understanding the impact of components on the effectiveness of the rules enabled generalizations beyond the five strategies studied. Third, how do the characteristics of the datasets (referred to as environmental parameters) affect the effectiveness of the heuristic information acquisition and restriction rules? Environmental parameters are the general descriptions of environments in which decision makers act. By linking the performance of heuristic information acquisition and restriction rules to environmental parameters, DSSs can be more precise in their acquisition and restriction of information. This study concluded with a mathematical example showing how components and environmental parameters combine to mediate the effectiveness of the rules.

6.1 Background

To study the effectiveness of the heuristic information acquisition and restriction rules, first requires a decision task, decision strategies, and measures of the decision environment. This section begins with a description of the fitting decision task with binary cues used in this study and concludes with the definitions of the proposed heuristic information acquisition and restriction rules. More information on the five strategies can be found in Sec. 4.3.1: weighted-additive (WADD), equal-weighting (EW), Tallying, Take Two, and take-the-best (TTB). Section 4.3.2 reviewed the measures of the environment that have the potential to mediate the effectiveness of the rules: cues, predictability, redundancy, and variability. Full information accuracy from Sec. 4.2.1 was also included as a potential mediator of the rules' effectiveness.

6.1.1 Fitting Decision Task with Binary Cues

The goal of decision making is to choose one of multiple options in order to maximize some measure of utility, also known as the criterion. The options are characterized by cue scores which describe the utility of the option with respect to important aspects of the environment (cues). The situation in which decision making occurs is referred to as the decision task. The decision tasks for this study have 2 options, 3 to 5 cues, accurate cue weights and cue directions, and binary cue scores. The decision tasks are generated from 15 datasets introduced by Czerlinski et al. (1999). These datasets have become a common method for benchmarking simulations as shown by their use in numerous simulations analyzing judgment and decision making strategies (Gigerenzer and Goldstein, 1996; Czerlinski et al., 1999; Martignon and Hoffrage, 2002; Hogarth and Karelaia, 2006; Katsikopoulos et al., 2010; Katsikopoulos, 2013).

The cue weights were calculated as the ecological validity of the cues based on the complete datasets, such that this study is categorized as a fitting task (Martignon and Hoffrage, 2002). This is differentiated from prediction tasks where cue weights are derived from some subset of the datasets. Therefore, in this study there is neither deviation from the 'true' cue weights (Hogarth and Karelaia, 2007) nor variability in the prior knowledge (Katsikopoulos et al., 2010). The fitting task also results in accurate cue directions such that the correlation between the cue and the criterion is positive: as a cue score increases (decreases), the criterion increases (decreases) (Katsikopoulos et al., 2010). Accurate cue directions are particularly necessary for the accuracy of decision making strategies (Katsikopoulos et al., 2010).

The decision task consisted of binary cue scores meaning that each option was described by cue scores equal to either 0 or 1. Katsikopoulos (2013) explained that this simplification can be justified in many ways: many cues are naturally binary (e.g. color versus black and white), decision makers may only be able to differentiate between or approximate good and bad, and choices can be made more manageable by treating a continuous cue as binary. While potentially more manageable, using binary cue scores instead of their original continuous form has been shown in some cases to decrease the performance of decision strategies (Luan et al., 2014).

The use of a fitting task with accurate cue weights, accurate cue directions, and binary cue scores does simplify this study away from the reality decision makers often face. However, these simplifications were necessary for this study to isolate strategy components, environmental parameters, and incomplete information as the only factors that could affect performance. Future examinations of heuristic information acquisition and restriction rules will study prediction tasks with continuous cue scores to better generalize the results.

6.1.2 Heuristic Information Acquisition and Restriction Rules

Given a specific distribution of initial incomplete information, which cue value should be provided to or removed from the decision maker to increase accuracy? This question of information search relates to the modification of content, one of the main capabilities of decision support systems (DSSs): altering what information is presented to the decision maker, including what categories of information are presented and at what level of detail or abstraction (Feigh et al., 2012). As described in the introduction to this study, there is a need for rules that can guide DSSs as to how to add and remove information to increase the accuracy of decision makers in a computationally efficient, robust, and transparent way (Katsikopoulos and Fasolo, 2006; Katsikopoulos et al., 2008).

Within the information acquisition literature there are no heuristic information acquisition and restriction rules for decision support. There are three reasons for the lack of these rules. First, the information acquisition literature is generally focused on judgment tasks (categorization or classification) not decision making tasks (selection between independent options). The typical motivating question within information acquisition is: "How can one anticipate the usefulness of possible information queries (questions, test, or experiments), before the answer (query result, experiment result, or test outcome) is known?" (Meder and Nelson, 2012, p. 120). For categorization tasks, both analytic methods (e.g. optimal experimental design which requires reliable assessments of probabilities, cue weights, and cue scores (Nelson, 2005, 2008; Meder and Nelson, 2012)) and heuristic methods (e.g. fast-and-frugal trees which require only a rank-order of the cues and are capable of categorizing after any single cue, Martignon et al., 2008; Luan et al., 2011) have been developed.

Second, the focus of the research has been on what information should be searched for next (acquisition), ignoring the possibility that performance could potentially be increased by removing information (restriction). Actively removing information is an extension of the less-is-more effect: "when less information or computation leads to more accurate [decisions] than more information or computation" (Gigerenzer and Gaissmaier, 2011, p. 453). The analytic optimal experimental design and heuristic fast-and-frugal trees both only provide a process for acquiring more information.

Third, at first glance, heuristic decision making strategies could be considered heuristic information acquisition methods. They search through a decision task with rules describing what information they should acquire next. However, they do not account for how distributions of incomplete information affect decision making performance such as those examined in Chap. 5.

For example, imagine a driver deciding between two routes (A and B) with the goal to select the fastest route, as shown in Fig. 6.1. The routes have the following cues in order of importance: distance, amount of traffic, and road type. The driver could have prior knowledge of some of the cue scores (traffic for route A and the road type for routes A and B) but not others. If the DSS provided additional information based on TTB, it would first acquire and present information about the distance of route A because distance is the highest ranked cue. However, that would increase option imbalance and maintain the number of balanced cues, likely decreasing the accuracy of the decision made.

Four heuristic information acquisition and restriction rules are proposed which provide computationally efficient, robust, and transparent methods for adding and removing information from a decision task:

- Option imbalance acquisition (OI-A): Increase total information while decreasing option imbalance.
- Option imbalance restriction (OI-R): Decrease total information while decreasing option imbalance.
- Cue balance acquisition (CB-A): Increase total information while increasing cue balance.
- Cue balance restriction (CB-R): Decrease total information while keeping cue balance constant.

Option imbalance acquisition and restriction are based on the results in Study 1 (Chap. 5) showing that decreasing total information while decreasing option imbalance increased the

Initial Decision Task with Incomplete Information

	Distance	Traffic	Road Type
Opt. A	?	K	K
Opt. B	?	?	K

TI	OI	CB
3	1	1

Heuristic Information Acquisition and Restriction Rules

Option Imbalance Acquisition (OI-A)

	Distance	Traffic	Road Type
Opt. A	?	K	K
Opt. B	? → K	?	K

TI	OI	CB
4	0	1

Cue Balance Acquisition (CB-A)
& Option Imbalance Acquisition (OI-A)

	Distance	Traffic	Road Type
Opt. A	?	K	K
Opt. B	?	? → K	K

TI	OI	CB
4	0	2

Cue Balance Restriction (CB-R)
& Option Imbalance Restriction (OI-R)

	Distance	Traffic	Road Type
Opt. A	?	K→ ?	K
Opt. B	?	?	K

TI	OI	CB
2	0	1

Figure 6.1: Examples of the four heuristic information acquisition and restriction rules at the individual decision task level for the exemplar driving scenario. The specific distribution of known and unknown information is denoted by white K's and black ?'s, respectively.

accuracy of heuristic strategies (specifically, TTB and Tallying). In Fig. 6.1 the information is imbalanced toward route A, so adding information to route B (OI-A) or restricting information about route A (OI-R) is likely to increase accuracy of heuristic strategies.

Cue balance acquisition and restriction are also based on the results in Study 1 (Chap. 5) showing that increasing total information without increasing cue balance had negligible effects on accuracy for TTB and Tallying, but increased the accuracy of WADD and EW. Additionally, increasing total information while increasing cue balance significantly increased accuracy for the heuristic strategies (TTB and Tallying). In Fig. 6.1 there is one balanced cue for road type, so adding information to route B for traffic in order to add another balanced cue (CB-A) or restricting information about route A for traffic (CB-R) is likely to increase the accuracy of heuristic strategies.

The two proposed heuristic information restriction rules are counter to a related computational study of WADD and TTB in which the strategies' accuracies decreased when they ignored cues for which one cue value was known and the other unknown (Garcia-Retamero and Rieskamp, 2008). This method for ignoring cues was identical to CB-R in Fig. 6.1 but was not formalized as restriction rules. Nevertheless, the prior results still suggest that if decision makers do ignore cues in specific ways, then accuracy may be increased. This follows the suggestion by Gigerenzer et al. (1991) that ignoring information could lead to high accuracy if decision makers had incorrect estimates for the missing cue values.

6.2 Method

The effectiveness of the four heuristic information acquisition and restriction rules were tested in a simulation of the five decision making strategies across 15 datasets with 2 options, 3 to 5 cues, and all combinations of incomplete information. The presence and magnitude of increases in strategy accuracy at the environment level (as shown in Fig. 6.2) were used as measures of effectiveness: 1) positive accuracy count (PAC): the percent of times in which the rules resulted in a positive change in accuracy, and 2) average accuracy change (AAC): the average change in accuracy when using the rules. A positive AAC value for heuristic information acquisition indicates that adding a piece of information following the rule results in a higher accuracy than adding a piece of information that does not follow the rule. A positive AAC value for heuristic information restriction indicates that removing a piece of information using the rule increases the accuracy of the strategy compared to not removing any information.

An example of the AAC for OI-R is shown in Fig. 6.2-A: $\{0.6\%, 1.1\%, 1.3\%, 2.0\%, 2.9\%\}$. Since all five AAC values are positive, the PAC is 100%. These two metrics indicate that, in aggregate, if the decision support system restricts incomplete information using OI-R for a decision maker using TTB in the House dataset with 2-options and 3-cues, then the chance of increasing the decision maker's accuracy is 100% and the average change in accuracy is

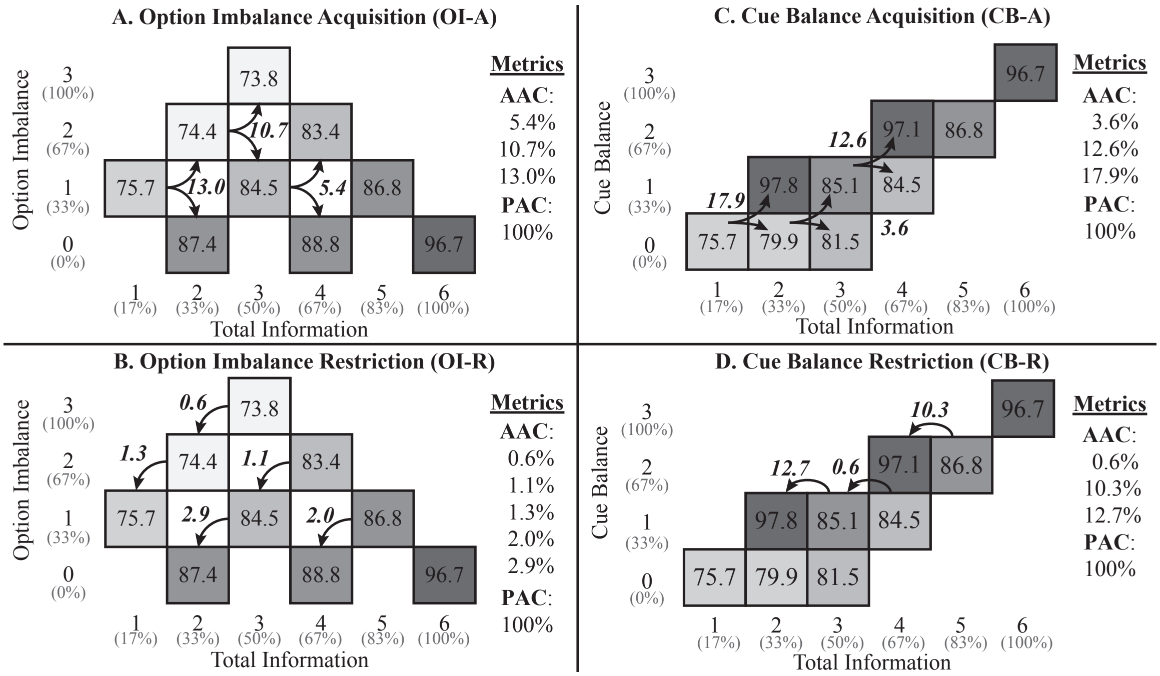

Figure 6.2: Example of option imbalance acquisition (OI-A) and restriction (OI-R) rules, and cue balance acquisition (CB-A) and restriction (CB-R) rules for TTB strategy for the House dataset with 3 cues. The comparisons of accuracy for acquisition are between the two endpoints of the arrows and the comparisons of accuracy for restriction are shown by the direction of the arrows. The change in accuracy is shown in bold-italic font.

1.58%.

Each set of AAC and PAC values are calculated for each rule as shown in Fig. 6.2. For each combination of decision environment and decision strategy there is one PAC value and between 3 and 14 AAC values – the number of AAC values changes based on the rule and the number of cues.

The datasets often included more than 5 cues and more than 1 criterion such that there were 25,069 decision environments that could have been generated from the 15 datasets with unique combinations of 3, 4, and 5 cues and 1 criterion. Therefore only one set of 3, 4, 5 cues and 1 criterion was chosen for each of the 15 datasets, resulting in 45 unique decision environments. The final 45 decision environments were selected in order to remove correlations between the environmental parameters.

Table 6.1 shows the values of the environmental parameters for 5-cue decision tasks for

Table 6.1: Environmental parameter values for the 5-cue decision environments.

Dataset	**Full Information Accuracy**					**Pred.**	**Redun.**	**Var.**
	WADD	**EW**	**Tallying**	**Take Two**	**TTB**			
Biodiversity	0.84	0.83	0.83	0.84	0.86	0.53	0.26	0.47
Boys	0.78	0.78	0.78	0.78	0.82	0.53	0.58	0.15
Car	0.75	0.78	0.78	0.75	0.72	0.57	0.21	0.16
Cities	0.81	0.83	0.83	0.81	0.81	0.86	0.22	0.23
Dropout	0.68	0.69	0.69	0.68	0.69	0.20	0.75	0.03
Fat	0.58	0.58	0.58	0.58	0.59	0.10	0.34	0.09
Fuel	0.79	0.83	0.83	0.79	0.79	0.51	0.25	0.16
Girls	0.74	0.76	0.76	0.74	0.76	0.44	0.51	0.18
Homeless	0.69	0.69	0.69	0.69	0.70	0.32	0.17	0.37
House	0.91	0.92	0.92	0.91	0.91	0.49	0.57	0.17
Oxygen	0.92	0.92	0.92	0.92	0.92	0.22	0.89	0.12
Pollution	0.77	0.79	0.79	0.78	0.77	0.48	0.34	0.17
Professor	0.81	0.81	0.81	0.81	0.82	0.70	0.30	0.43
Rainfall	0.69	0.73	0.73	0.69	0.68	0.33	0.12	0.13
Sleep	0.81	0.81	0.81	0.81	0.84	0.42	0.52	0.22

each dataset.[1] For example, the parameters for the House dataset with 5 cues are 0.91 FIA for WADD, Take Two, and TTB, and 0.92 FIA for EW and Tallying, 0.49 predictability, 0.57 redundancy, and 0.17 variability. For each level of cues, the distribution of the values for FIA, predictability, variability, and redundancy spanned the range from low to high.

Within each of the 45 decision environments, every combination of two options was combined with a full factorial combination of incomplete information. Each decision strategy made a decision on each decision task combination of incomplete information and options. Every option has a known criterion such that a decision strategy is considered to make the correct choice when it selects the option with the higher dataset criterion. The accuracy percentages in the boxes of Fig. 6.2 are the percent of tasks matching that combination of incomplete information in which the strategy selected correctly.

[1]Environmental parameters for 4-cue and 3-cue decision tasks are in Table A.5 and Table A.6, respectively.

6.3 Results

Three results are described in the following subsections. In Sec. 6.3.1, the results of AAC and PAC for each heuristic information acquisition and restriction rule for each decision making strategy shows that the rules generally increase accuracy (Fig. 6.3 and Table 6.2). In Sec. 6.3.2, the estimate of missing information is shown to be the main strategy component mediating the effectiveness of the rules (Table 6.3). In Sec. 6.3.3, full information accuracy (FIA) is shown to have the largest significant effect on AAC for every combination of rule and decision making strategy. Of the rest of the environmental parameters with significant effects, cues and predictability were the most common.

6.3.1 General Effectiveness of Heuristic Information Acquisition and Restriction Rules

The results presented in Fig. 6.3 and Table 6.2 show that the heuristic information acquisition rules tend to increase accuracy more than 86% of the time for each of the five strategies. Said another way, it is almost always beneficial to gain a new piece of information using the heuristic information acquisition rules regardless of the decision strategy. For heuristic information restriction the results are more nuanced. Both of the restriction rules tend to increase accuracy more than 71% of the time for Tallying, Take Two, and TTB, suggesting that restricting information generally works for the heuristic strategies. However, for WADD and EW, it is only beneficial to restrict a piece of information when using CB-R. Using OI-R for the two analytic strategies tends to be slightly more likely to decrease accuracy but with only a small magnitude.

Figure 6.3 shows the distributions of PAC (how often a given a rule increased the accuracy for each strategy) and AAC values (how much the accuracy of each strategy increased on average by using a rule once). Table 6.2 summarizes the results of Fig. 6.3 by showing the average PAC and AAC values. For example, for heuristics, the acquisition rules (OI-A and CB-A) had nearly a 100% chance of increasing accuracy in aggregate (PAC) with an

Table 6.2: Positive accuracy count (PAC) and average accuracy change (AAC) for each decision strategy when using heuristic information acquisition and restrictions rules averaged across all decision environments.

Strategy	**Cue Weights**	**Estimates**	**Option Balance**				**Cue Balance**			
			Acquisition		**Restriction**		**Acquisition**		**Restriction**	
			PAC	**AAC**	**PAC**	**AAC**	**PAC**	**AAC**	**PAC**	**AAC**
WADD	**Compensatory**	**Average**	88%	2.2	40%	-0.1	86%	4.4	62%	3.5
EW	**Equal**	**Average**	89%	2.7	44%	0.3	88%	4.5	60%	3.6
Tallying	**Equal**	**Negative**	100%	10.1	93%	2.9	100%	4.5	71%	2.1
Take Two	**Compensatory**	**Negative**	100%	7.3	91%	1.5	100%	5.3	88%	2.9
Take-the-Best	**Non-Compensatory**	**Negative**	100%	6.8	86%	1.2	100%	5.5	90%	2.9

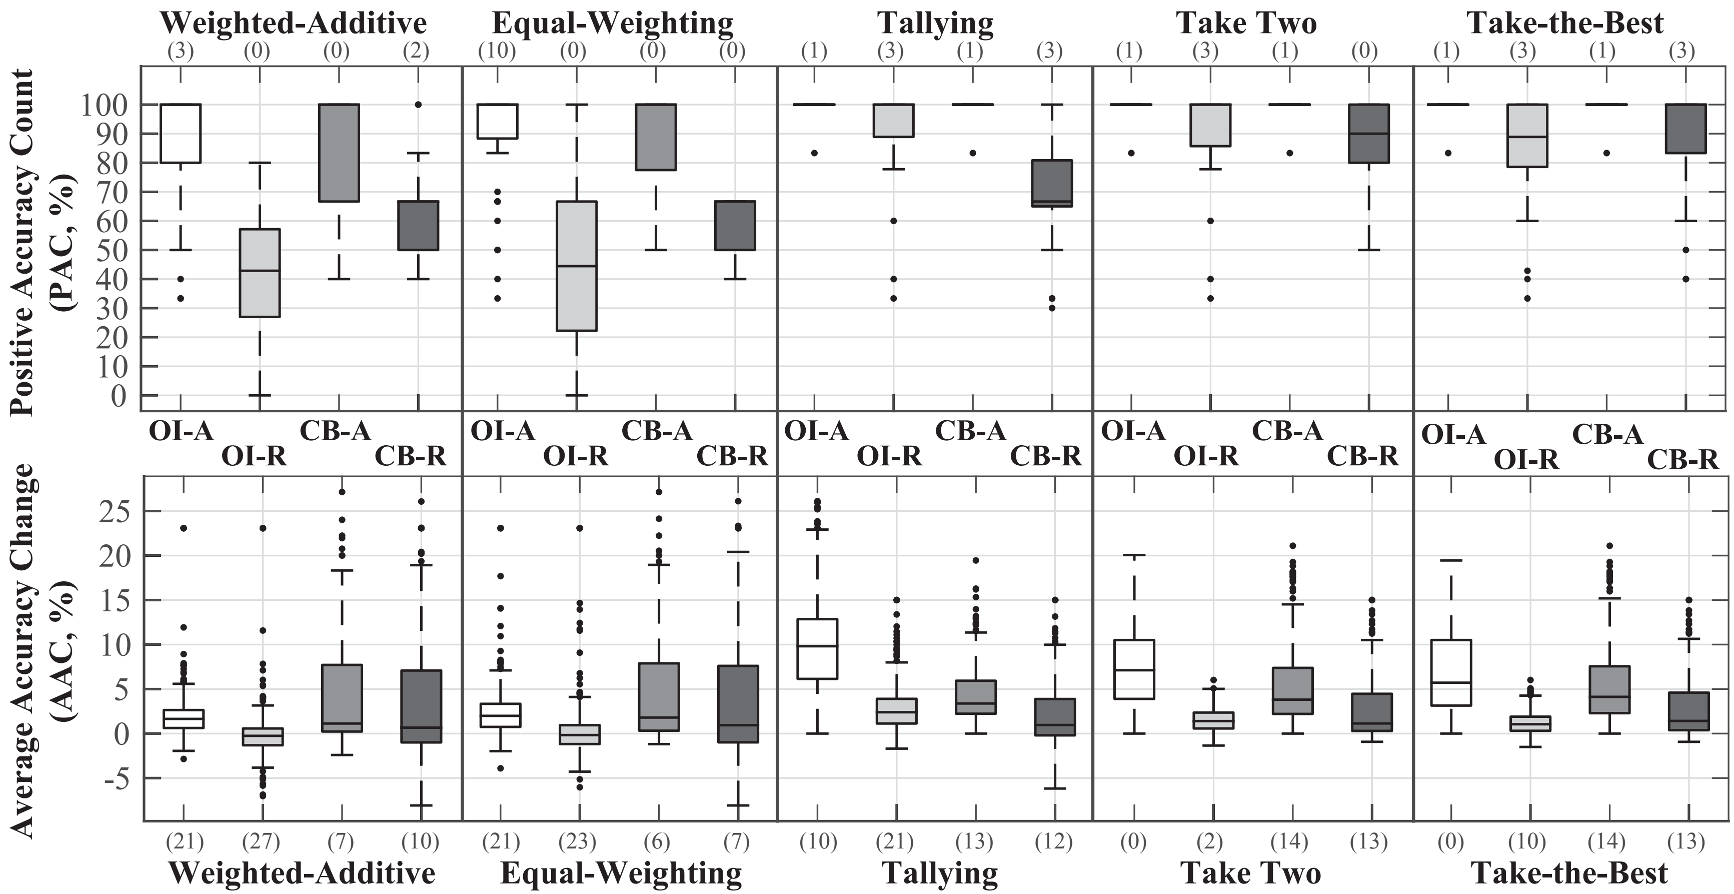

Figure 6.3: Average accuracy change (AAC) and positive accuracy count (PAC) when using each of the heuristic information acquisition and restriction rules for each decision making strategy: option imbalance acquisition (OI-A) and restriction (OI-R), and cue balance acquisition (CB-A) and restriction (CB-R). AAC measures the average change in accuracy of a strategy when using the rule and PAC measures how often the change in accuracy is positive. Each individual box plot has 45 data points for PAC, 420 data points for OB-R for AAC, and 285 data points for OB-A, CB-A, and CB-R for AAC. The box plots indicate the 25^{th} percentile, median, and 75^{th} percentile. Data points (denoted by parentheses) are outliers if they lie outside the whiskers corresponding to approximately 2.7σ if the data was normally distributed.

average increase in accuracy (AAC) greater than 4.5%.

6.3.2 Effect of Strategy Components

The results in Table 6.2 and Fig. 6.3 suggest that the heuristic information acquisition and restriction rules affect WADD and EW differently than Tallying, TTB, and Take Two. As shown in Table 4.2 the strategies differ in only two components: how they measure the relative importance of cues (cue weights) and how they estimate missing information. Furthermore, only how the strategies estimate missing information matches the differences found in Table 6.2 and Fig. 6.3: WADD and EW estimate missing cue weights as the average (0.5) while Tallying, Take Two, and TTB estimate missing information as negative (0). This suggested that the main cause of the differences in effectiveness of the heuristic information acquisition and restriction rules is due to how strategies estimate missing information.

Linear models accounting for the main effects and interactions of cue weights and estimates of missing information were fit to the PAC and AAC results for each rule. The results in Table 6.3 begin with the baseline (intercept) values describing the estimated PAC and AAC values for strategies such as TTB which have non-compensatory cue weights and estimate missing information as negative. Each additional column describes the relative change in PAC and AAC from the baseline (non-compensatory cue weights and negative estimates) to a different combination of cue weights and negative estimates. Interactions were considered significant at the $\alpha = 0.001$ level.

For example, the third column describes strategies like Tallying which use equal cue weights and negative estimates of missing information. Compared to the baseline of non-compensatory cue weights and negative estimates, this combination of components does not have a statistically different PAC for OI-A, OI-R, or CB-A, but did have a statistically significant difference in CB-R of -12%, indicating that the absolute PAC value is 66%.

The results in Table 6.3 show that how strategies estimate missing information has a

Table 6.3: Significance of strategy components on the PAC and AAC values for each heuristic information acquisition and restriction rule.

Cue Weights		**Non-Compensatory**	**Compensatory**	**Equal**	**[Non-Compensatory]**	**Compensatory**
Estimates of missing information		**Negative**	**[Negative]**	**[Negative]**	**Average**	**Average**
Positive	**Option imbalance acquisition**	95%**	0%	0%	-5%*	0%
Accuracy	**Option imbalance restriction**	66%**	0%	3%	-24%**	-1%
Count	**Cue balance acquisition**	94%**	-1%	0%	-6%**	-1%
(PAC)	**Cue balance restriction**	78%**	-3%	-12%**	-6%*	-7%*
Average	**Option imbalance acquisition**	4.4**	0.4	2.0**	-3.7**	1.1**
Accuracy	**Option imbalance restriction**	0.6**	0.1	1.0**	-1.3**	0.5**
Change	**Cue balance acquisition**	5.1**	-0.2	-0.7**	0.0	-0.5
(AAC)	**Cue balance restriction**	3.4**	-0.2	-0.5	0.8**	-0.5

*$<$ 0.001, **$<$ 0.0001. The baseline (intercept) values are non-compensatory weights with negative estimates of missing information. Each following column is the relative change in PAC and AAC from this baseline when the components are changed.

significant effect on the effectiveness for every rule except CB-A – likely because CB-A reduces the use of estimates of missing information. Furthermore, for each significant effect except CB-R, estimating missing information as 0.5 rather than negative resulted in a decrease in the likelihood of a rule increasing accuracy (PAC) or the average increase in accuracy (AAC).

The results in this subsection support previous research showing that how strategies estimate missing information can have a more significant effect on performance than how strategies integrate information (cue weights) (Garcia-Retamero and Rieskamp, 2008). Conceptually, it makes sense that the effectiveness of the heuristic information acquisition and restriction rules are mediated by how strategies estimate missing information. Estimates of missing information are only necessary when information is incomplete. The estimates make assumptions about cue scores which can be inaccurate. The assumption of the three heuristics – that unknown cue values are negative – is a strong assumption that was based on a predefined characteristic of the strategies, not information about the decision environment. Therefore, the rules increase accuracy for these strategies because the rules attempt to restructure the decision task so that decisions are made with a focus on the cues which have known information for all options. This attempts to decrease reliance on estimates of missing information which may be incorrect (Gigerenzer et al., 1991).

6.3.3 Effect of Environmental Parameters

The heuristic information acquisition and restriction rules were shown in Table 6.2 to generally increase the accuracy of decision making strategies. The next objective is to determine if any environmental parameters mediate the magnitude of the average accuracy change (AAC) when using the rules. Table 6.4 shows the significance of the estimates of the coefficients for each environmental parameter regressed against the AAC values for all 20 combinations of the four rules and the five strategies. Results shown in Table 6.4 were deemed significant at the $\alpha = 0.001$ level.

Table 6.4: Significance of environmental parameters on the AAC values for each combination of decision making strategy and heuristic information acquisition and restriction rule using linear models.

		Cues	FIA	Predictability	Redundancy	Variability
WADD	**OI-R**	-0.4	8.6**	-3.2**	5.9**	-7.1**
	OI-A	-1.1**	16.1**	-0.5	6.8**	-7.9**
	CB-R	-2.3**	22.0**	2.1	8.0**	-8.7
	CB-A	-2.4**	18.6**	6.0*	4.6	-5.8
EW	**OI-R**	-0.6	8.8**	-3.7**	6.7**	-8.4**
	OI-A	-1.6**	17.0**	-1.3	8.2**	-10.3**
	CB-R	-2.2**	20.4**	2.6	6.9**	-7.5
	CB-A	-2.1**	16.5**	6.7**	3.0	-3.1
Tallying	**OI-R**	-0.6*	11.1**	0.9	2.8**	-0.2
	OI-A	-1.9**	30.5**	7.5**	3.0	5.7
	CB-R	-1.1*	12.7**	1.1	4.7**	-3.8
	CB-A	-1.3**	14.7**	6.5**	0.1	2.7
Take Two	**OI-R**	-0.2	3.0**	1.4**	-0.6	2.0**
	OI-A	-1.1*	16.5**	7.5**	-1.5	9.9**
	CB-R	-1.4**	16.2**	1.2	5.7**	-4.7
	CB-A	-1.8**	18.4**	6.4**	2.4	1.2
TTB	**OI-R**	-0.3**	2.2*	1.0*	-0.4	0.9
	OI-A	-1.1	15.0**	6.9**	-1.2	9.1**
	CB-R	-1.5**	15.6**	1.3	5.7**	-5.4
	CB-A	-1.8**	18.2**	6.6**	2.7	0.5

285 data points per statistical test for OI-A, CB-A, CB-R,
and 420 data points per statistical test for OI-R, *$<$ 0.001, **$<$ 0,0001

The results show that full information accuracy (FIA) had the largest, positive, significant effect on AAC results for all 20 combinations. The rest of the environmental parameters in order of number of significant cases were cues (13), predictability (12), redundancy (10), and variability (7). The values in Table 6.4 for cues are the changes in AAC if the number of cues is increased by one. The values in Table 6.4 for FIA, prediction, redundancy, and variability are to the changes in AAC if the parameters increased by 100%. Therefore, if the full information accuracy (FIA) of a decision maker using TTB increases by 10% for a given dataset then the average accuracy change when using CB-A should increase by 1.82% (10% of 18.2%).

Similar to the explanation of the estimates of missing information results in Sec. 6.3.2, the positive correlation between FIA and AAC is likely because the rules are restructuring the distribution of information so that decisions are made with a focus on the cues which have full information – approximating decisions with FIA. For example, if the CB-R was used as much as possible within a two-option decision task, it would reduce the distribution of incomplete information such that there were only balanced cues. Therefore, the accuracy of the decisions would depend on the accuracy of the cues with full information, which are correlated with strategies' full information accuracy.

The reason that FIA had larger and more significant effects on AAC than the other environmental parameters (cues, predictability, variability, and redundancy) is that changes in FIA is the culmination of other environmental parameters affecting accuracy. As described in Sec. 4.3.2, most of the previous research on environmental parameters examined how they affected accuracy in decision tasks with full information. In other words, those previous studies were measuring the effect of other environmental parameters on FIA. Therefore, changes in predictability, redundancy, and variability, alone without changes in FIA, should not generally affect the effectiveness of the rules.

Changes in the number of cues require a footnote to this generalization because they can affect the specific magnitude of individual AAC values. As shown in Table 6.4, the coefficient estimates for cues are always negative, meaning that as the number of cues increase, the AAC decreases for all strategies. This negative correlation is a result of the increased size of the decision task. For example, as the number of cues increases, there is an increase in the amount of information needing to be acquired or restricted to decrease option imbalance by the same percentage. Thus, with the same FIA, an increase in number of cues will result in each individual acquisition or restriction affecting AAC by a smaller amount – not necessarily decreasing accuracy.

6.4 A Mathematical Explanation of Estimates and FIA as Mediators of Rule Effectiveness

The two main mediators of the effectiveness of the rules were the estimates of missing information and the full information accuracy of the strategy. A brief mathematical result can show how these two mediators can be integrated into a more full understanding of the rules: FIA sets the upper limit of the accuracy the rules can help the strategies' achieve and estimates determine how much accuracy changes when altering the distribution of incomplete information. Further mathematical analyses will need to be conducted but this example should provide an intuition.

Let's imagine a one cue decision environment with two options, A and B. The correct decision is to select option A which has a higher dataset criterion value than option B. The decision maker will therefore make the correct decision when the decision maker's utility for option A is greater than option B as shown in Eq. 6.1 – cue weights are not included because there is only one cue being considered. Using the component-based general linear model of decision making from Chap. 3, the utility functions can be rewritten based on cue scores (a), estimates of missing information (e) and incomplete information (z) as shown in Eq. 6.2.

$$U(A) > U(B) \tag{6.1}$$

$$e + [a(A) - e]z(A) > e + [a(B) - e]z(B) \tag{6.2}$$

Table 6.5 shows the resulting accuracy for the four possible combinations of incomplete information. The full information accuracy of the strategy is set when information is known for both options (OI, 0%; CAP, 100%). Let's assume that the cue scores are binary as was the case in this study such that $a(A) = 1$ and $a(B) = 0$. Therefore, the FIA for this decision task is the ecological validity of the cue (v, Martignon and Hoffrage, 2002): the conditional probability of A being the correct option when $a(A) = 1$ and $a(B) = 0$.

Table 6.5: Resulting decision equations and accuracy as a function of estimates of missing information for a two-option, one-cue decision task with incomplete information.

$\mathbf{z(A)}$	$\mathbf{z(B)}$	**Option Imbalance**	**Cue Balance**	**Decision Equation**	**Accuracy** $(\mathbf{e=1})$	$(\mathbf{e=0.5})$	$(\mathbf{e=0})$
1	1	0%	100%	$1 > 0$	v	v	v
1	0	100%	0%	$1 > e$	50%	v	v
0	1	100%	0%	$e > 0$	v	v	50%
0	0	0%	0%	$e > e$	50%	50%	50%

The results of the decision equations in Table 6.5 show that the accuracy of the strategy with only one cue known (OI, 100%; CAP, 0%) depends on the estimate of missing information. If the strategy estimates missing information as positive ($e = 1$), the strategy would be as accurate as the cue's validity when only Option B is known (v) but would select randomly when only Option A is known (50%). The converse is true for estimating missing information as negative ($e = 0$), when only Option A is known the accuracy is the cue's validity but when only Option B is known the accuracy is random. Furthermore, if the strategy estimates missing information as an average cue score between 0 and 1, the decision will be as accurate as the cue's validity (v) both when only Option A is known and when only Option B is known. Therefore, the average accuracy for the strategy for scenarios with 100% IB and 0% CB is the average of 50% and v ($avg(50\%, v)$) for the positive and negative estimates, and v for the average estimate.

These small scale mathematical results support the overall simulation results. The goal of the heuristic information acquisition and restriction rules is to reduce option imbalance and increase cue balance. The FIA and the estimate of missing information determines the effectiveness of the rules because they set the upper and lower limit of the strategy's achievable accuracy. From the math example, if a strategy was estimating missing information as positive or negative (e.g. Tallying, Take Two, or TTB), then acquiring information would increase the average accuracy from the average of 50% and v, to v (PAC, 100%, assuming v is greater than $avg(50\%, v)$; AAC, $v/2 - 25\%$). However, if the strategy was estimating missing information as an average (e.g. WADD and EW), using the heuristic information acquisition rule would not increase the average accuracy as both were equivalent to the cue's validity (PAC, 0%; AAC, 0%).

6.5 Discussion

The motivating question for this research was to identify how and when decision support systems should acquire information for or restrict information from the operator. The results of the simulation and brief math analysis show that heuristic information acquisition and restriction rules increase the accuracy of decision strategies in a majority of cases (measured by PAC) and often substantially (measured by AAC), and are mediated by the strategies' full information accuracy and estimations of missing information. The following subsections describe a synthesized version of the heuristic information acquisition and restriction rules, expand upon the usefulness of the component-level abstraction of decision making strategies, and identify the implications of the rules for developing decision support systems.

6.5.1 Rules for Heuristic Information Acquisition and Restriction

The general rules for heuristic information acquisition and restriction are:

> **Acquire** information to decrease option imbalance or increase cue balance, or both.
>
> **Restrict** information to decrease option imbalance, and maintain or decrease cue balance.

The pair of heuristic information acquisition rules or the pair of restriction rules cannot always be performed together. However, if both the option imbalance and cue balance rules were to be combined, the accuracy should likely be increased more than either in isolation. The rules cannot always be performed together because option imbalance and cue balance are coupled, as shown in Fig. 6.1. For example, when acquiring information to decrease option imbalance (OI-A), cue balance can either be increased (CB-A) or maintained – never decreased. Given that OI-A and CB-A had high probabilities of increasing accuracy (PAC) and high values of accuracy increases (AAC), it is inferred that decreasing option imbalance while increasing cue balance would be the best option; but the current results are not conclusive. Now that the general rules of heuristic information acquisition and restriction have shown to be effective at the environment level, future work will need to measure PAC and AAC values at the individual decision task level.

6.5.2 Components of Strategies: A New Level of Abstraction

The results extend the perspective that viewing strategies as a unique combination of components will be a useful level of abstraction for studies investigating the determinants of decision making performance. To understand the importance of this perspective, compare the strategy-level insights versus the component-level insights.

At the strategy-level, WADD and EW are viewed as analytic strategies and Tallying, Take Two, and Take-the-Best as heuristic strategies. Reviewing the results from Table 6.2 (excluding the component descriptions) and Fig. 6.3 would therefore suggest that there is something specifically separating the two types of strategies. The heuristic information acquisition rules increase accuracy less often for analytic strategies (86%-89%) than for heuristic strategies (100%). Furthermore, for the heuristic information restriction rules, OB-R increases accuracy less than 50% of the time for analytic strategies but 86% of the time or more for heuristic strategies. However, what is it about the two groups of strategies that makes the rules differentially effective? There is nothing obvious in the definitions of heuristics, that they use simple rules and ignore part of the information (Gigerenzer and Gaissmaier, 2011), that clearly explains these results.

The component-level abstraction of strategies in Sec. 6.3.2 showed that the differences in effectiveness of the heuristic information acquisition and restriction rules were mostly caused by how strategies estimated missing information and were not particularly affected by cue weights – a result first found by Garcia-Retamero and Rieskamp (2008). Furthermore, the mathematical analysis based on the component-level abstraction in Sec. 6.4 provided preliminary justification for the interaction of estimates of missing information and the full information accuracy.

The importance of accurately estimating missing information can be seen by the main objective of the rules: to increase the cue-wise and option-wise balance of the distribution of information. The rules acquire or restrict information to decrease option imbalance and increase cue balance so that decision making strategies do not have to make decisions in which one options' cue value is known and another options' cue value is unknown. Essentially, the rules attempt to restructure the decision task so that decisions are made with a focus on the cues which have all the information known thus reducing the reliance on estimates of missing information.

6.5.3 Heuristic Decision Support and Decision Support for Heuristics

The effectiveness of the heuristic information acquisition and restriction rules showed both the potential for heuristic decision support methods and a need for further research into designing decision support specifically for heuristic decision makers.

The results show that decision support systems themselves can utilize heuristics that are effective and efficient in order to help decision makers be effective and efficient. First, the rules are generally effective. The acquisition rules are highly likely to increase accuracy across a variety of strategies. The restriction rules are likely to increase accuracy for strategies that estimate missing information as negative or positive and do not negatively impact accuracy for strategies with average estimates. Second, the rules are transparent and easy to communicate. The rules alter the distribution of information in order to create a balance of information between options and within cues. Third, the rules require little information and process that information quickly. The rules only require a count of how many cues are known for each cue and each option.

The differences in the effectiveness of the rules between strategies and components suggest that effective support for heuristic decision making strategies can sometimes be very different from effective support for analytic strategies. Using OI-R as the example, the rule was more likely to increase accuracy and a larger change in accuracy for the heuristic strategies (Tallying, Take Two, and TTB) than the analytic strategies (EW and WADD). While it was determined that this difference is largely due to how the strategies estimate missing information, the results still suggest that decision support system design should account for the entire range of strategies and constituent components.

Furthermore, removing available information from a decision maker (OI-R and CB-R) was shown to be an effective, if counterintuitive, method of decision support. In application, heuristic information restriction rules could be used if there is not enough time (Rieskamp and Hoffrage, 2008), money (Bröder, 2000), or capacity (Fasolo et al., 2007) for the decision maker to process all of the available information; or if the decision maker has a tendency to exhibit information bias. Further research is needed to determine how to implement information restriction appropriately in a decision support tool and to relate it to previous research on decision making strategies which ignore cues with missing information (Garcia-Retamero and Rieskamp, 2008).

6.6 Conclusion

This study presented heuristic information acquisition and restriction rules for decision support which can address some of the psychological and environmental issues that make analytic information acquisition challenging. The simulation results showed that the rules are generally effective for increasing the accuracy of decision making strategies (especially for acquisition), transparent and easy to communicate (create a balance of information between options and within cues), and require little information to perform (only requiring knowledge of the distribution of known and unknown information).

Simulation and mathematical analyses based on a component-level abstraction of strategies were used to identify the main mediators of the rules' effectiveness as the strategies' estimate missing information and full information accuracy. Further research is needed to fully describe how strategy components and environmental parameters interact to affect the performance of the rules at the task level. The results should ultimately provide an impetus for the development of heuristic methods of decision support and decision support methods specifically for heuristic decision making strategies.

CHAPTER 7

COMPUTER SIMULATION STUDY 3: DETERMINANTS OF DECISION MAKING WITH INCOMPLETE INFORMATION

This study aims to comprehensively address the questions of incomplete information and heuristic information acquisition and restriction only initially addressed in the previous two studies. The first study (Chap. 5) showed that measuring total information is insufficient for fully characterizing the effects of incomplete information on decision making performance. Using an artificial dataset with a weighted-additive model for defining the criterion, the accuracy of heuristic strategies like take-the-best and tallying were shown to be affected not just by total information but by the balance of information between options and within cues. Option imbalance, the difference in the amount of information known for the most well-known and the least-well known, was shown to negatively affect the accuracy of heuristic strategies as it increased – essentially forcing the heuristic strategies to select at random when sufficiently high. Cue balance, the count of how many cues have known values for all options, was shown to positively affect that accuracy of heuristic strategies as it increased. In fact, for heuristic strategies, the results suggested that reducing information while reducing option imbalance or increasing cue balance could actually increase decision making accuracy.

These tradeoffs between total information, option imbalance, and cue balance were formalized into heuristic information acquisition and restriction rules in the second study (Chap. 6): option imbalance acquisition (increase total information while decreasing option imbalance), option imbalance restriction (decrease total information while decreasing option imbalance), cue balance acquisition (increase total information while increasing cue balance), and cue balance restriction (decrease total information while keeping cue balance constant). Using 15 real-world datasets and measuring the changes in accuracy using aggregate-level averages, the rules were shown to be quite effective. The acquisition rules were shown to increase accuracy by 2%-4% for for weighted-additive and equal-weighting, and by 5%-10% for take-the-best and tallying. Even the restriction rules showed promise, particularly cue balance restriction which increased accuracy on average by 2%-4%.

To goal of this third study is to address the methodological limitations of prior studies through a comprehensive simulation of decision making with incomplete information. Instead of artificial datasets, seven real-world datasets resulting in 306 environments, from the UCI Machine Learning Database will be used (Lichman, 2013). Instead of studying the well-known, named strategies of take-the-best, tallying, and weighted-additive, three fundamental components of decision making strategies will be varied independently: cue weights, estimates of missing information, and cutoff type (see Chap. 3). Instead of measuring the effectiveness of the information acquisition and restriction rules at the aggregate-level, the change in accuracy for individual acquisition and restriction of each cue value will be measured.

The paper is structured around the three motivating questions of this dissertation. First, how does incomplete information affect decision making accuracy? The single variable results show that generally, increasing total information, reducing option imbalance, and increasing complete attribute pairs, increase accuracy of all strategies.

Second, how are strategies differentially affected by these distributions of incomplete information? Strategies' estimates of missing information are shown to have the most significant effect on decision making accuracy with incomplete information, followed by cutoff type, and lastly, cue weights. This confirms and expands previous work by Garcia-Retamero and Rieskamp (2008) showing that to achieve highest accuracy with incomplete information, a strategy must match the estimate of missing information to the environment, nearly independent of the cue weighting method. The two-way interactions between the measures of incomplete information show that when the estimates matched the environment (in this study, median estimates matched the uniform distribution of missing information), the accuracy of the decision making strategies become almost entirely dependent on total information, irrespective of option imbalance or cue balance. However, when strategies' estimates did not match the environment (positive and negative estimates in this study), option imbalance became the major determinant of decision making accuracy, more so than total information or cue balance.

Third, can information acquisition and restriction methods, which only rely on knowledge of the distribution of incomplete information, increase decision making accuracy? In short, yes, they can, depending on the estimates of missing information. On average, heuristic information acquisition methods were able to increase accuracy by 1% for strategies whose estimates matched the environment and up to 4% for strategies whose estimates did not match the environment. In restriction, the best outcomes were to achieve an increase in accuracy of less than 1%, especially for strategies whose estimates did not match the environment.

This study concludes with discussions of how cue weights and total information are insufficient ways to characterize decision making with incomplete information, and the implications of the heuristic information acquisition and restriction rules.

7.1 Background

7.1.1 Incomplete information

Decision making with incomplete information refers to unknown cue values within the decision task. Incomplete information is often caused by time pressure, high information acquisition costs, too much information, or incomplete information in the environment (see Sec. 2.2). Incomplete information within a decision task can be characterized along three dimensions: total information, option imbalance, and cue balance. Total information (TI) measures how many cue values are known within the decision task. Previous work is in agreement that more information generally increases accuracy of decision making strategies though the amount of increase is dependent on the fit between the strategy and the environment (see Chap. 5; Martignon and Hoffrage, 2002; Garcia-Retamero and Rieskamp, 2008). Option imbalance (OI) measures the difference in amount of known cues values between the most-known option and the least-known option. Using an artificial dataset and a linear model as the criterion, Chap. 5 showed that decreasing option imbalance drastically increased the accuracy of heuristics like take-the-best and tallying but had no effect on weighted-additive or equal-weighting. Cue balance (CB) measures the number of cues in which cue values are known for all options. Using the same artificial dataset and a linear model for accuracy, Chap. 5 showed that increasing cue balance has a positive effect on all the strategies studied: take-the-best, tallying, weighted-additive, and equal-weighting.

The most interesting results arose when examining the interaction of the measures of incomplete information, showing that reducing total information while reducing option imbalance or maintaining cue balance could actually increase the accuracy of take-the-best and tallying. The accuracy of weighted-additive and equal-weighting were only affected by total information. These notable results for the two heuristic strategies suggested that accuracy could be increased by reformulating the distribution of incomplete information through acquisition or restriction to minimize option imbalance and maximize cue balance.

7.1.2 Heuristic Information Acquisition and Restriction Rules at the Aggregate-Level

In Chap. 6, the tradeoffs were formalized as heuristic information acquisition and restriction rules: methods of adding or removing cue values from a decision task that do not require knowledge of the cue values, cue weights, or probabilities, only the distribution of incomplete information. Four specific rules were defined: option imbalance acquisition (increase total information while decreasing option imbalance), option imbalance restriction (decrease total information while decreasing option imbalance), cue balance acquisition (increase total information while increasing cue balance), and cue balance restriction (decrease total information while keeping cue balance constant). The four rules were tested for their effectiveness in increasing accuracy at the aggregate-level across 15 datasets with two options, and three to five cues. The aggregate-level of analysis refers to measures of changes in decision making accuracy were calculated based on the averages of an entire dataset, rather than each individual task.

The results of the simulation showed that the heuristic information acquisition and restriction rules had promise. Option imbalance acquisition resulted in an average accuracy increase of 2% for weighted-additive and equal-weighting, and 6% to 10% for take-the-best and tallying. In comparison, cue balance acquisition was more effective for weighted-additive and equal-weighted, with an average accuracy increase of 4%, but less effective for take-the-best and tallying, with an average accuracy increase of 5%. Even the restriction rules showed promise. Of two the restriction rules, cue balance restriction had more positive results, having an average accuracy increase for each strategy between 2% and 4%. Option imbalance restriction was effective in increasing the accuracy of take-the-best and tallying by 1.2% and 2.9%, respectively, but having essentially no positive or negative effect on the accuracy of weighted-additive or equal-weighting.

Using statistical analysis and a mathematical result, the effectiveness of the rules were shown to be dependent on the estimates of missing information and not on the cue weights. Extending the prior results by Garcia-Retamero and Rieskamp (2008, 2009), the strategies that correctly modeled the relationship between the missing information and cue scores (i.e. weighted-additive and equal-weighting estimated the uniformly missing information as a median value) were more robust to the effects of incomplete information, and thus did not benefit as much from the rules altering the distribution of information as those strategies with incorrect models of the relationship (i.e. take-the-best and tallying estimated the missing information as negative).

The analysis at the aggregate-level had limitations that the task-level analysis in this study will overcome. The analysis did not account for the fact that cue balance and option imbalance are coupled despite its mention in the discussion. Therefore, it is unknown how the various acquisition and restriction rules should or should not be combined. Furthermore, the aggregate-level analysis did not measure what specific cues were being removed relative to the application of each rule, eliminating the possibility of examining how knowledge of cue weights could impact the performance of the rules.

7.2 Method

The goal of this study was to characterize the effect of cue weights, cutoff types, and estimates of missing information as interacting, independent variables on decision making accuracy within various distributions of incomplete information and on the effectiveness of acquiring or restricting information at the decision task level.

The accuracy of the decision making strategies was tested across the 7 datasets in Table 7.1 from the UCI Machine Learning Database (Lichman, 2013). For each dataset, all 2-option decision tasks with 5 cues and 1 criterion were investigated. To study incomplete information, each 2-option decision task was combined with all combinations of missing information and provided to each decision making strategy. Eighteen strategies were constructed from all combinations of cue weights (compensatory, equal, and non-compensatory), estimates of missing information (positive, negative, and median), and cutoff types (prior and relative). The cue weights, estimates, and directions were calculated based on all options in the dataset. To study the effectiveness of acquiring or restricting cue values, within each 2-option decision task, all acquisitions and restrictions of each cue value were performed and categorized as one of eight possible information acquisition or restriction types in Table 7.2.

7.2.1 Strategies

The decision tasks in this study are paired comparison tasks in which the goal is to predict which of two options has the higher criterion value. For example, it can be asked which of two wind tunnel conditions causes higher sound pressure for a specific airfoil. The decision can be based cues including the angle of attack or the free-stream velocity.

Table 7.1: Characteristics of the seven datasets used in Study 3.

Dataset[a]	Description	Cues, Criteria	Options	Environments	Total Tasks
Concrete Slump	Choose the concrete with higher slum, flow, and compressive strength based on components including cement, slag, fly ash, water, coarse aggregate, and fine aggregate.	7, 3	103	63	3.31 E 5
Computer Hardware	Choose the computer hardware with higher performance (published or estimated) based on components including memory, channels, and cycle time.	6, 2	209	12	2.61 E 5
Yacht Hydrodynamics	Choose the sailing yacht with higher resistance based on coefficients including prismatic, beam-draught, and Froude numbers.	6, 1	308	6	2.84 E 5
Forest Fires	Choose the forest location with higher area burned based on weather factors including wind speed, temperature, relative humidity, and rain.	8, 1	517	56	7.47 E 6
Energy Efficiency	Choose the building with higher heating and cooling loads based on features including surface area, wall area, overall height, and glazing area.	8, 2	768	112	3.30 E 7
Concrete Compress	Choose the concrete with higher compressive strength based on components including cement, slag, fly ash, water, coarse aggregate, and fine aggregate.	8, 1	1030	56	2.97 E 7
Airfoil Self-Noise	Choose the wind tunnel condition with higher sound pressure for NACA 0012 airfoil based on conditions including angle of attack, frequency, and free-stream velocity.	5, 1	1503	1	1.13 E 6

The datasets are available at the UCI Machine Learning Database: http://archive.ics.uci.edu/ml/datasets/.

Table 7.2: Information acquisition and restriction types for the 2-option decision tasks.

Name	Rule	Cue Balance	Option Imbalance
-4	Restriction	-1	-1
-3	Restriction	-1	1
-2	Restriction*	0	-1
-1	Restriction*	0	1
1	Acquisition	0	1
2	Acquisition*	0	-1
3	Acquisition*	1	1
4	Acquisition*	1	-1

*Examples of the heuristic information acquisition and restriction rules studied in Chap. 6.

The decision making strategies in this study are combinations of three components described by the general linear model of decision making in Chap. 3 and Eq. 7.1. Each decision making strategy in this study is modeled using the binary utility function which ultimately converts the cue values to binary cue scores before integrating them with the cue weights. This choice is justified because the non-compensatory weights cannot be used without binary scores, and the equal weights and compensatory weights from cue validities require that the cue values are normalized in a consistent manner. However, future studies will include exact utility functions so that models like multiple linear regression and logistic regression can be compared.

$$C_i \quad = \quad \sum_{j=1}^{n} w_j \cdot H\Big[d_j\Big(e_j + (a^v_{i,j} - e_j)z_{i,j}\Big) - (d_j \cdot c_j) + \Delta_j\Big] \tag{7.1}$$

Cutoff type denotes whether the strategy requires the cue values to be dichotomized using a cutoff value defined prior to the decision task (prior, P) or using a cutoff value defined relative to the cue values within the task (relative, R). All strategies using prior cutoffs in this study had cutoff values equivalent to the median of all cue values ($c_j = \tilde{a}_j$). Relative cutoffs compare each original cue value to either the maximum or minimum cue value within the decision task, depending on the cue direction ($c_j = \text{max/min } a_j$). Relative cutoffs can approximate the use of continuous cues scores for non-compensatory weights like take-the-best and other strategies which require binary cue scores Katsikopoulos et al. (2010).

Cue weights (w_j) define the relative importance of cues and generally describe how the cue scores are integrated. Equal weights (E) sets all the cue weights equal ($w_j = 1$). Non-compensatory cue weights (N) represent strategies in which once a cue discriminates in favor of an option, the decision will not change regardless of the information processed for lower-ranked cues (Martignon et al., 2008) ($w_j = 4^{1-j}$). Compensatory cue weights (C) technically describes all other distributions of cue weights which are not equal nor non-compensatory (Hogarth and Karelaia, 2006). However, in this study, compensatory cue weights describes a type of strategy which does not have any preference for compensatory or non-compensatory ($w_j = v_j$).

The cues weights for compensatory strategies and cue ordering for non-compensatory strategies were calculated based on the frequentist cue validity (Gigerenzer and Goldstein, 1996),

$$v = \frac{R}{R + W} \tag{7.2}$$

where R and W are the number of correct and incorrect decisions, respectively, based on an individual cue.

Estimates (e) denote how the strategy estimates and replaces missing cue values: positive (P, maximum cue value, $\max a_j^v$), median (M, $\hat{a}_j^v$), and negative (N, minimum cue value, $\min a_j^v$). As described in Sec. 3.1.3, the maximum and minimum cue values are defined independently of the determination of cue directions (d_j) – whether the cues are positively or negatively correlated with the criterion. For each strategy the threshold was set to zero ($\Delta = 0$) meaning that any difference between two cue values is large enough to tell them apart. This is the standard threshold value used in simulation studies as it has often been ignored and thus equivalent to zero.[1]

Combinations of the two cutoff types, three cue weights, and three estimates resulted in 18 unique strategies which are presented in Table 7.3. The components represent large sets of the three families of well-studies strategies based on their cue weights: tallying and equal-weighting (EW), take-the-best (TTB), and weighted-additive (WADD). The traditional form of tallying is represented by the PNE model with prior cutoffs (P), negative estimates (N) and equal weights (E), so that it chooses the option with the higher number of positive cue scores. The traditional form of TTB is represented by the PNN model with prior cutoffs (P), negative estimates (N), and non-compensatory weights (N), so that it chooses the option that discriminates on the highest-ranked cue i.e. a cue that has a positive cue score for one option and an unknown or negative cue score for the other. The traditional form of WADD is represented by the RMC model with relative cutoffs (R), median estimates (M), and compensatory weights (C), so that it sums the product of each cue score and cue validity, selecting the option with the highest criterion.

$$C_{PNE} = \sum_{j=1}^{n} H\Big[d_j\Big(\min a_j^v + (a_{i,j}^v - \min a_j^v)z_{i,j}\Big) - (d_j \cdot \tilde{a}_j^v)\Big] \tag{7.3}$$

$$C_{PNN} = \sum_{j=1}^{n} 4^{1-j} \cdot H\Big[d_j\Big(\min a_j^v + (a_{i,j}^v - \min a_j^v)z_{i,j}\Big) - (d_j \cdot \tilde{a}_j^v)\Big] \tag{7.4}$$

$$C_{RMC} = \sum_{j=1}^{n} v \cdot H\Big[d_j\Big(\tilde{a}_j^v + (a_{i,j}^v - \tilde{a}_j^v)z_{i,j}\Big) - (d_j \cdot \max / \min a_j)\Big] \tag{7.5}$$

7.3 Results

7.3.1 One-Variable Effects

Figure 7.1 plots the average accuracy of each strategy as a function of each individual measure of incomplete information. The plots for total information (top plot, Fig. 7.1)

[1] See Payne et al., 1990, for an exception and Luan et al., 2014, for a comprehensive study of thresholds.

Table 7.3: Strategy codes for the 18 variations of the components of the strategies based on the general linear model of decision making.

Strategy Family	Strategy Code	Cutoff Type	Estimates	Cue Weights
Tallying and equal-weighting (EW)	PME	(P) Prior	(M) Median	(E) Equal
	RME	(R) Relative	(M) Median	(E) Equal
	PNE	(P) Prior	(N) Negative	(E) Equal
	RNE	(R) Relative	(N) Negative	(E) Equal
	PPE	(P) Prior	(P) Positive	(E) Equal
	RPE	(R) Relative	(P) Positive	(E) Equal
Take-the-Best (TTB)	PMN	(P) Prior	(M) Median	(N) Non-compensatory
	RMN	(R) Relative	(M) Median	(N) Non-compensatory
	PNN	(P) Prior	(N) Negative	(N) Non-compensatory
	RNN	(R) Relative	(N) Negative	(N) Non-compensatory
	PPN	(P) Prior	(P) Positive	(N) Non-compensatory
	RPN	(R) Relative	(P) Positive	(N) Non-compensatory
Weighted-additive (WADD)	PMC	(P) Prior	(M) Median	(C) Compensatory
	RMC	(R) Relative	(M) Median	(C) Compensatory
	PNC	(P) Prior	(N) Negative	(C) Compensatory
	RNC	(R) Relative	(N) Negative	(C) Compensatory
	PPC	(P) Prior	(P) Positive	(C) Compensatory
	RPC	(R) Relative	(P) Positive	(C) Compensatory

and cue balance (bottom plot, Fig. 7.1) showed the biggest differences occurred between strategies using prior or relative cutoffs. Generally, strategies with prior cutoffs showed a higher accuracy than strategies with relative cutoffs. The strategies, in order of decreasing full information accuracy, were: prior non-compensatory (71%), prior compensatory (70%), prior equal (68%), and then relative non-compensatory, compensatory, and equal all at 63%. The trends of total information and cue balance showed that the accuracy of strategies with prior cutoffs had larger increases than relative cutoffs.

As expected, median estimates, which correctly match the uniform distribution of missing information, do make the accuracy robust to option imbalance, regardless of the cue weights or the cutoffs (middle plot, Fig. 7.1). The 6 strategies with median estimates maintained or even slightly increased their accuracy as option imbalance increased. Positive or negative estimates, which incorrectly assume a positive or negative relationship between the missing cue values and the criterion, were drastically affected by option imbalance. When option imbalance increased from 0 to 5, accuracy reduced by 8% to 14%. When the information between options was balanced, the accuracy of the strategies generally converged the same accuracy, especially amongst strategies whose only differences were how they estimated of missing information.

7.3.2 Two-Variable Interaction Effects

Each decision task is defined by not just one of these measures of incomplete information, but by all three. So although, increasing total information, decreasing option imbalance, and increasing cue balance, individually tend to result in greater accuracy for strategies, one has to examine the two-way interactions as shown in Fig. 7.2. In sum, there are two types of interactions based on whether the strategy estimated the missing information incorrectly as positive or negative (left column of plots in Fig. 7.2), or whether the strategy estimated the missing information correctly as the median (right column of plots in Fig. 7.2).

Strategies that estimated missing information incorrectly showed the tradeoffs discov-

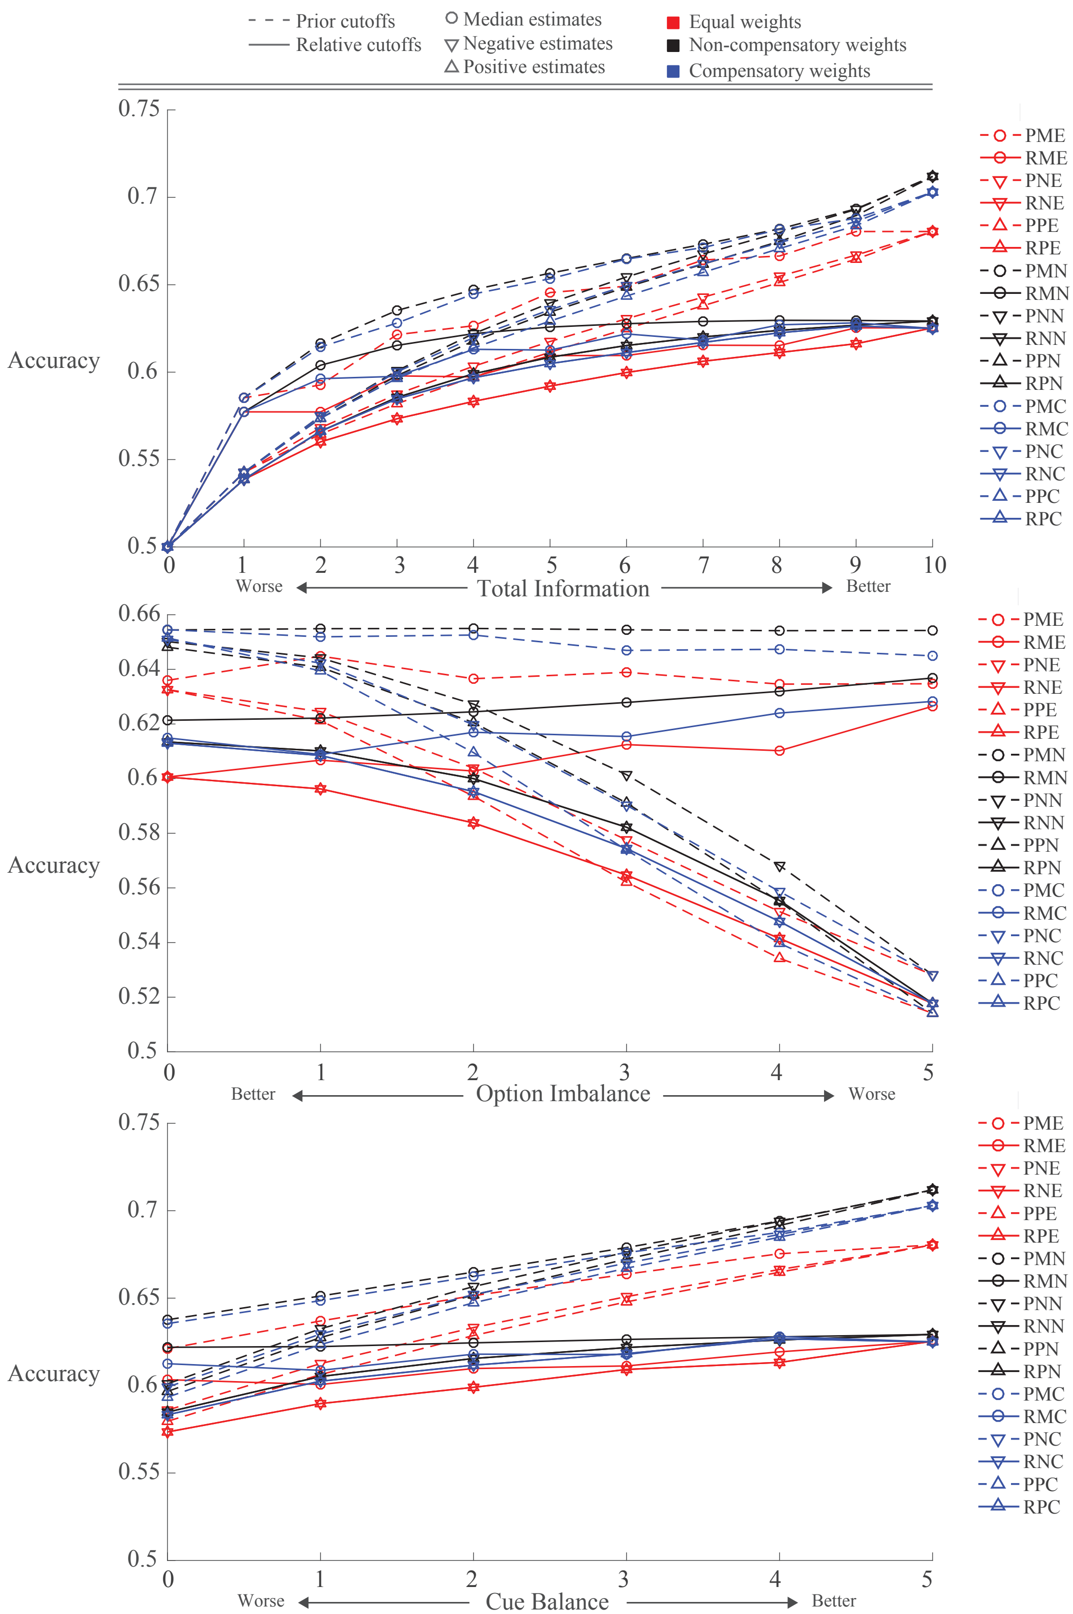

Figure 7.1: Accuracy of decision making strategies as a function of total information, option imbalance, and cue balance.

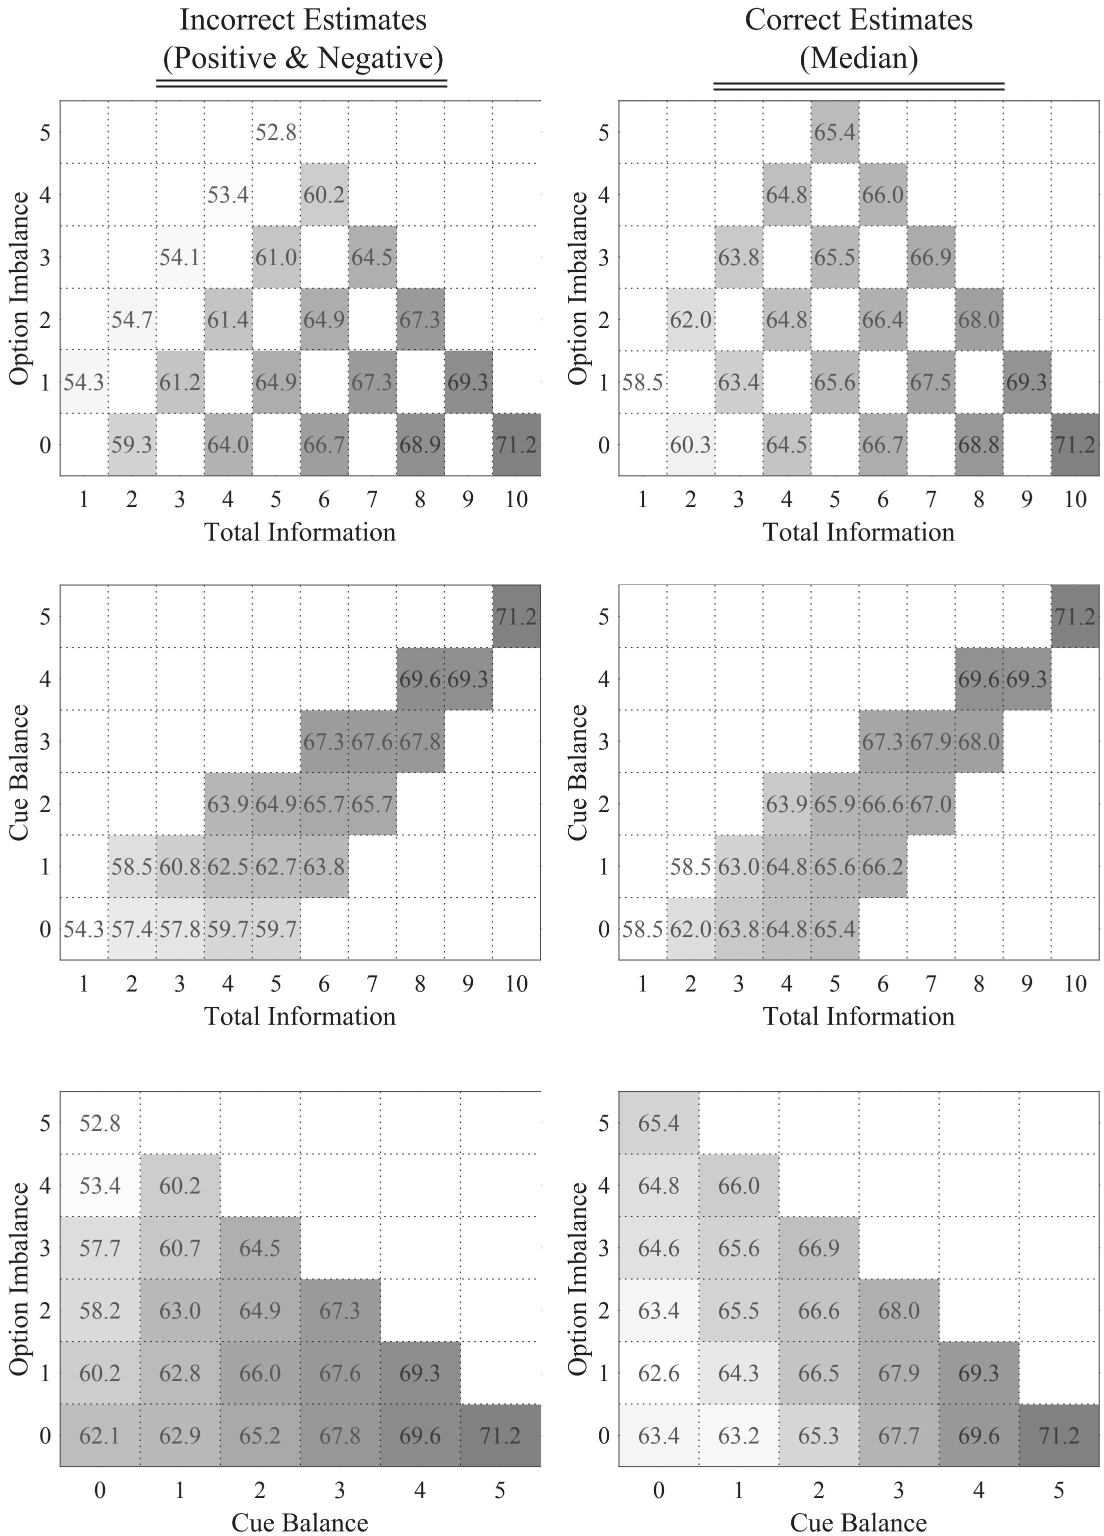

Figure 7.2: Exemplar two-way interactions between the three measures of incomplete information. The left column showing the results of the PNN strategy represents strategies whose estimates of missing information did not match the environment. The right column showing the results of the PMN strategy represents strategies whose estimates of missing information did match the environment.

Table 7.4: Accuracy of decision making strategies with incomplete information. Cell colors are conditionally formatted with a blue-white-red transition from low to high accuracy.

TI	RNN	PNN	RNE	PNE	RMC	PMC	RPN	PPN	RPE	PPE	RMN	PMN	RME	PME	RNC	PNC	RPC	PPC
0	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50
1	0.54	0.54	0.54	0.54	0.58	0.59	0.54	0.54	0.54	0.54	0.58	0.59	0.58	0.59	0.54	0.54	0.54	0.54
2	0.57	0.57	0.56	0.57	0.60	0.61	0.57	0.57	0.56	0.56	0.60	0.62	0.58	0.59	0.57	0.57	0.57	0.57
3	0.59	0.60	0.57	0.59	0.60	0.63	0.59	0.60	0.57	0.58	0.62	0.64	0.60	0.62	0.58	0.60	0.58	0.60
4	0.60	0.62	0.58	0.60	0.61	0.64	0.60	0.62	0.58	0.60	0.62	0.65	0.60	0.63	0.60	0.62	0.60	0.61
5	0.61	0.64	0.59	0.62	0.61	0.65	0.61	0.63	0.59	0.61	0.63	0.66	0.61	0.65	0.60	0.64	0.60	0.63
6	0.62	0.65	0.60	0.63	0.62	0.66	0.62	0.65	0.60	0.62	0.63	0.67	0.61	0.65	0.61	0.65	0.61	0.64
7	0.62	0.67	0.61	0.64	0.62	0.67	0.62	0.66	0.61	0.64	0.63	0.67	0.62	0.66	0.62	0.66	0.62	0.66
8	0.62	0.68	0.61	0.65	0.63	0.68	0.62	0.67	0.61	0.65	0.63	0.68	0.62	0.67	0.62	0.67	0.62	0.67
9	0.63	0.69	0.62	0.67	0.63	0.69	0.63	0.69	0.62	0.66	0.63	0.69	0.63	0.68	0.63	0.69	0.63	0.68
10	0.63	0.71	0.63	0.68	0.62	0.70	0.63	0.71	0.63	0.68	0.63	0.71	0.63	0.68	0.62	0.70	0.62	0.70

CB	RNN	PNN	RNE	PNE	RMC	PMC	RPN	PPN	RPE	PPE	RMN	PMN	RME	PME	RNC	PNC	RPC	PPC
0	0.58	0.60	0.57	0.59	0.61	0.64	0.58	0.60	0.57	0.58	0.62	0.64	0.60	0.62	0.58	0.60	0.58	0.59
1	0.61	0.63	0.59	0.61	0.61	0.65	0.61	0.63	0.59	0.61	0.62	0.65	0.60	0.64	0.60	0.63	0.60	0.62
2	0.62	0.66	0.60	0.63	0.62	0.66	0.62	0.65	0.60	0.63	0.62	0.66	0.61	0.65	0.61	0.65	0.61	0.65
3	0.62	0.68	0.61	0.65	0.62	0.68	0.62	0.67	0.61	0.65	0.63	0.68	0.61	0.66	0.62	0.67	0.62	0.67
4	0.63	0.69	0.61	0.67	0.63	0.69	0.63	0.69	0.61	0.66	0.63	0.69	0.62	0.68	0.63	0.69	0.63	0.68
5	0.63	0.71	0.63	0.68	0.62	0.70	0.63	0.71	0.63	0.68	0.63	0.71	0.63	0.68	0.62	0.70	0.62	0.70

OI	RNN	PNN	RNE	PNE	RMC	PMC	RPN	PPN	RPE	PPE	RMN	PMN	RME	PME	RNC	PNC	RPC	PPC
0	0.61	0.65	0.60	0.63	0.61	0.65	0.61	0.65	0.60	0.63	0.62	0.65	0.60	0.64	0.61	0.65	0.61	0.65
1	0.61	0.64	0.60	0.62	0.61	0.65	0.61	0.64	0.60	0.62	0.62	0.65	0.61	0.64	0.61	0.64	0.61	0.64
2	0.60	0.63	0.58	0.60	0.62	0.65	0.60	0.62	0.58	0.59	0.62	0.65	0.60	0.64	0.60	0.62	0.60	0.61
3	0.58	0.60	0.56	0.58	0.62	0.65	0.58	0.59	0.56	0.56	0.63	0.65	0.61	0.64	0.57	0.59	0.57	0.57
4	0.56	0.57	0.54	0.55	0.62	0.65	0.56	0.55	0.54	0.53	0.63	0.65	0.61	0.63	0.55	0.56	0.55	0.54
5	0.52	0.53	0.52	0.53	0.63	0.64	0.52	0.51	0.52	0.51	0.64	0.65	0.63	0.63	0.52	0.53	0.52	0.51

ered in Chap. 5. Decreasing total information while decreasing option imbalance, tends to increase or maintain accuracy. Decreasing total information while maintaining cue balance tends to only slightly decrease accuracy. Overall, within the same amount of total information, the best distribution of information is that in which cue balance is maximized and option imbalance is minimized.

Strategies that estimated missing information accurately were robust to changes in option imbalance or cue balance: their accuracy simply increased when total information increased and decreased when total information decreased. For these strategies, within the same amount of total information, decreasing option imbalance or increasing cue balance had essentially no effect.

These results expand the inferences made by Garcia-Retamero and Rieskamp (2008) who showed that for strategies to achieve highest accuracy, estimates of missing information must match distribution of missing information in the environment. In this simulation, missing information was uniformly distributed. Therefore, missing information should be estimated as the median, not as positive or negative. The results in Fig. 7.2 show that when strategies match the environment they do not only have better accuracy overall, but their accuracy becomes a function of only total information, largely unaffected by other aspects of the distributions of incomplete information.

7.3.3 Heuristic Information Acquisition and Restriction Rules

As expected, the effectiveness heuristic information acquisition and restriction rules shown in Fig. 7.3 follow the results of the two-way interactions: for the strategies that estimated missing information correctly as the median, the effectiveness of the rules were driven by total information, but for strategies that estimated missing information incorrectly as positive or negative, the effectiveness was driven by option imbalance. The performance of the various heuristic information acquisition and restriction rules was measured by positive accuracy count (PAC, the percent of times in which the rules resulted in a positive change in accuracy) and average accuracy change (AAC, the average change in accuracy when using the rules).[2] The relative impact of the strategy components in order, again, ranged from estimates of missing information which had a significant effect, to the cutoffs which had a moderate effect, and finally to the cue weights which had almost no effect.[3]

For the strategies that estimated missing information correctly as the median, the results are actually best described through the differential effects of cutoffs. Relative cutoffs with median estimates were essentially unaffected by the nuances of option imbalance or cue balance: acquiring information resulted in an average increase of around 1% while restricting information led to an average decrease of 1%. However, prior cutoffs preferred that information be acquired or restricted in such a way that cue balance was maintained or reduced, respectively.

For the strategies that estimated missing information incorrectly as positive or negative, the best rules can be summarized as those that acquire or restrict information to reduce option imbalance, regardless of cue balance. Acquiring information while decreasing option imbalance resulted in a 2% to 4% increase, increasing accuracy between 57% and 64% of the time. Restricting information to decrease option imbalance resulted in 0.5% decrease to 0.5% increase, increasing accuracy between 48% and 54% of the time.

The lack of influence of cue balance on these results is a little surprising – the negative relationship for some strategies even more so. Prior work in Chap. 6 showed that restricting total information while maintaining cue balance was an effective rule at the aggregate-level, increasing accuracy by 2% to 4% regardless of estimates of missing information. That work showed that acquiring total information to increase cue balance was even more effective, increasing accuracy by 4% to 6% regardless of estimates of missing information. It was thought that particularly for the strategies that estimated missing information incorrectly as positive or negative cue values, increasing cue balance would be effective as it reduced the

[2]Note that the measures of acquisition in Study 2 compared acquiring information following the rules to acquiring information not following the rules. Whereas in this study, the measures of acquisition compare acquiring information following the rules to the current distribution of incomplete information.

[3]Complete numerical results are presented in Tables A.7 and A.8.

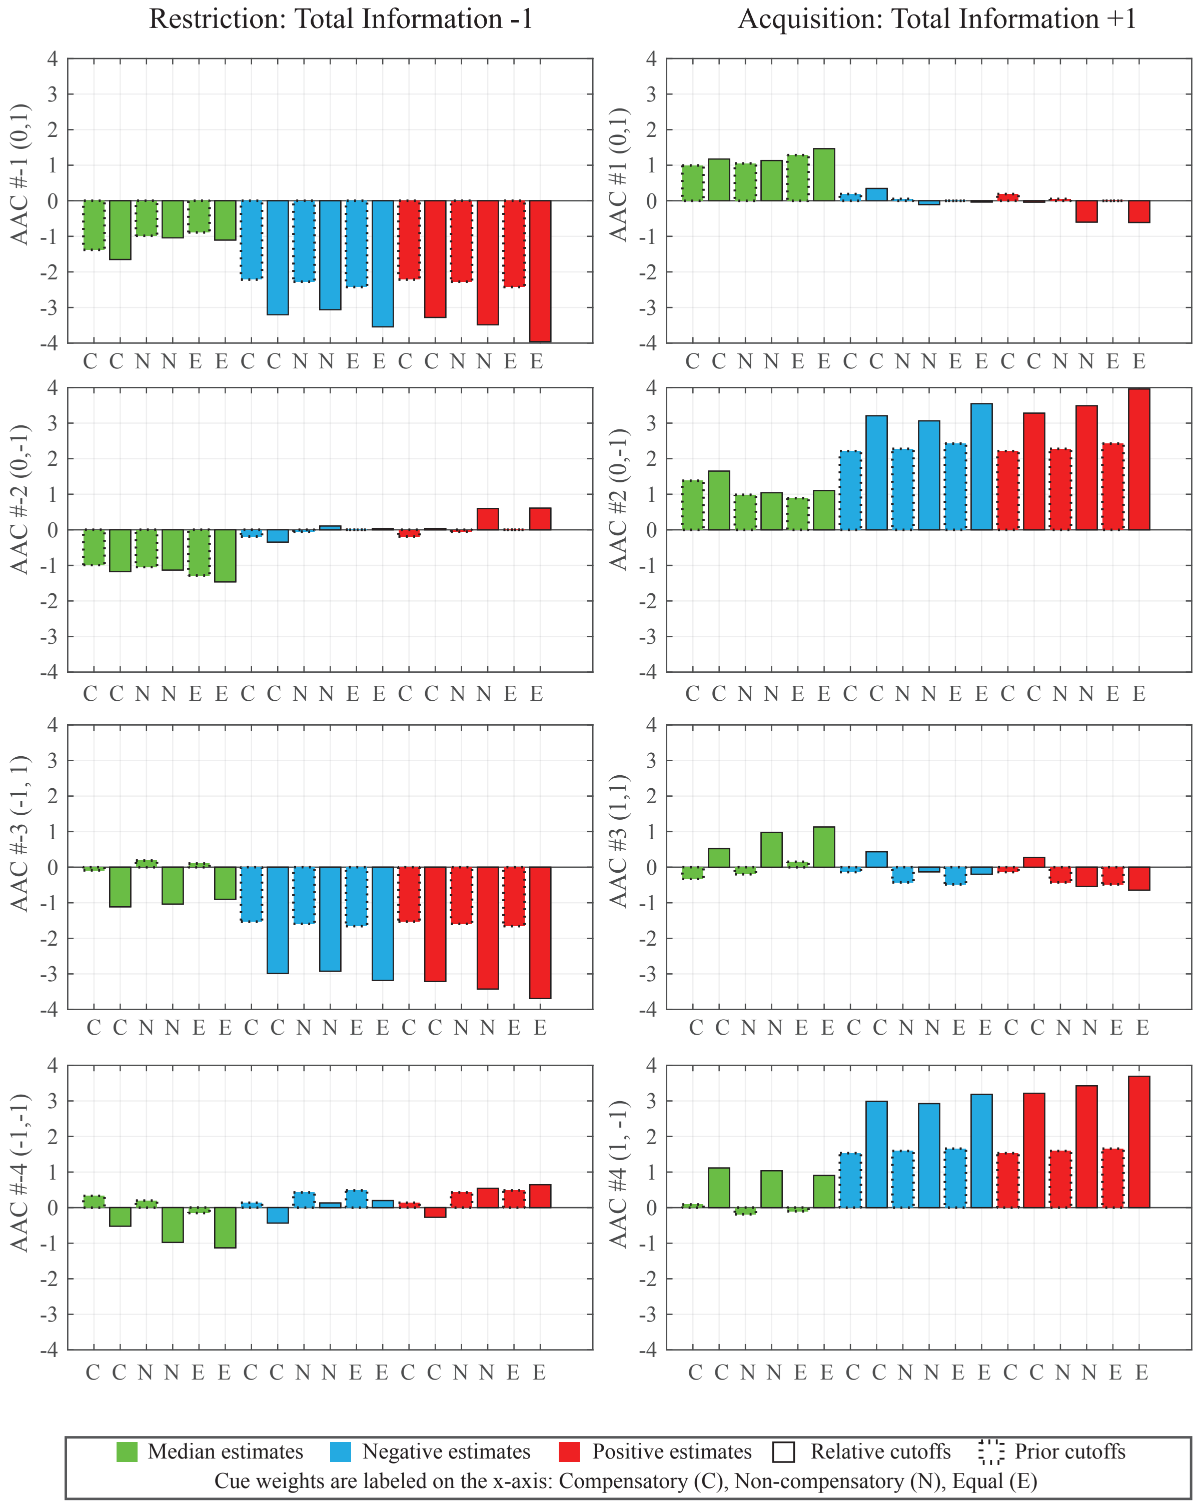

Figure 7.3: Average accuracy change (AAC) for each information acquisition and restriction type. Parentheses denote the change in cue balance and option imbalance, respectively.

likelihood that decisions were made using the incorrect estimates of missing information. It could be that since this prior analysis did not measure option imbalance and cue balance simultaneously as was done in this study, the the option imbalance results were embedded in the cue balance acquisition and restriction results.

The alternating heights of the bars in Fig. 7.3 is a result of the alternating use of prior or relative cutoffs. The prior cutoffs almost always have the same direction of effectiveness for each rule (increasing or decreasing accuracy) as their complimentary relative cutoffs, but with a decreased magnitude. Therefore, it seems that in addition to estimating missing information correctly, converting cues to binary using the median cue values rather than relative cue values may make strategies more robust to changes in incomplete information.

In sum, the task-level results are more conservative than the aggregate-level results, but still show the promise of heuristic information acquisition and restriction rules.[4] When the estimates do not match the environment, heuristic information acquisition and restriction rules should try to reduce option imbalance regardless of their effects on cue balance. However, when estimates do match the environment, heuristic information acquisition and restriction rules should acquire information to not affect cue balance. The most effective heuristic information restriction rule for median estimates was to reduce cue balance, but that only achieved mixed results.

7.3.4 How Valuable is Cue Rank Information?

One of the key aspects defining the heuristic information acquisition and restriction rules is that they do not use cue rank to determine which cue to acquire or restrict. However, in many cases, there are potentially multiple cues that can be acquired or restricted that satisfy the rules for reducing option imbalance or maintaining cue balance. Figure 7.4 answers the question: what if the heuristic information acquisition and restriction rules had

[4] Future work will measure the change in achievement, rather than accuracy, to make studies more comparable. It could be that the environments used in this study are just more difficult for the strategies than the ones used in Chap. 6.

cue rank information? By plotting the effectiveness of the rules against cue rank, it is clear that the rules are at their most effective when restricting lower ranked cues and acquiring higher ranked cues. This is not a surprising result but shows just how much more effective the rules can be if combined with knowledge of cue rank. When acquiring the highest ranked cue to reduce option imbalance, accuracy can increase between 1% and 5%. When restricting the lowest ranked cue to reduce option imbalance, accuracy can decrease 1% or even increase up to 2%. Alternatively, it is clear how variable the results of heuristic information acquisition and restriction by cue rank, ranging by up to 2% for the same rule and strategy, just based on the rank of the cue.

7.4 Discussion

7.4.1 Beyond Total Information and Cue Weights

Decision making with incomplete information is shown to be a fundamentally different regime than the full information decision tasks commonly studied. Well-known decision making strategies have been largely characterized and named by how they integrate the cue values via cue weights. Heuristics like take-the-best use non-compensatory weights that can make decisions using only one cue. Others, like tallying or equal-weighting, don't use cue weights, instead integrating all cues equally. These heuristics are typically contrasted with strategies or models that use compensatory weights like multiple linear regression or weighted-additive. These differences are important in full information decision tasks. However, these differences break down as total information is reduced.

When there is incomplete information it becomes clear that two aspects largely determine the performance of decision making strategies: the estimates of missing information and option imbalance. When the estimates of missing cue values match the environment (i.e., median estimates if the cue values are uniformly missing as in this study), total information becomes the main determinant of decision making accuracy. However, when the estimates of cue values do not match the environment (i.e., the positive and negative esti-

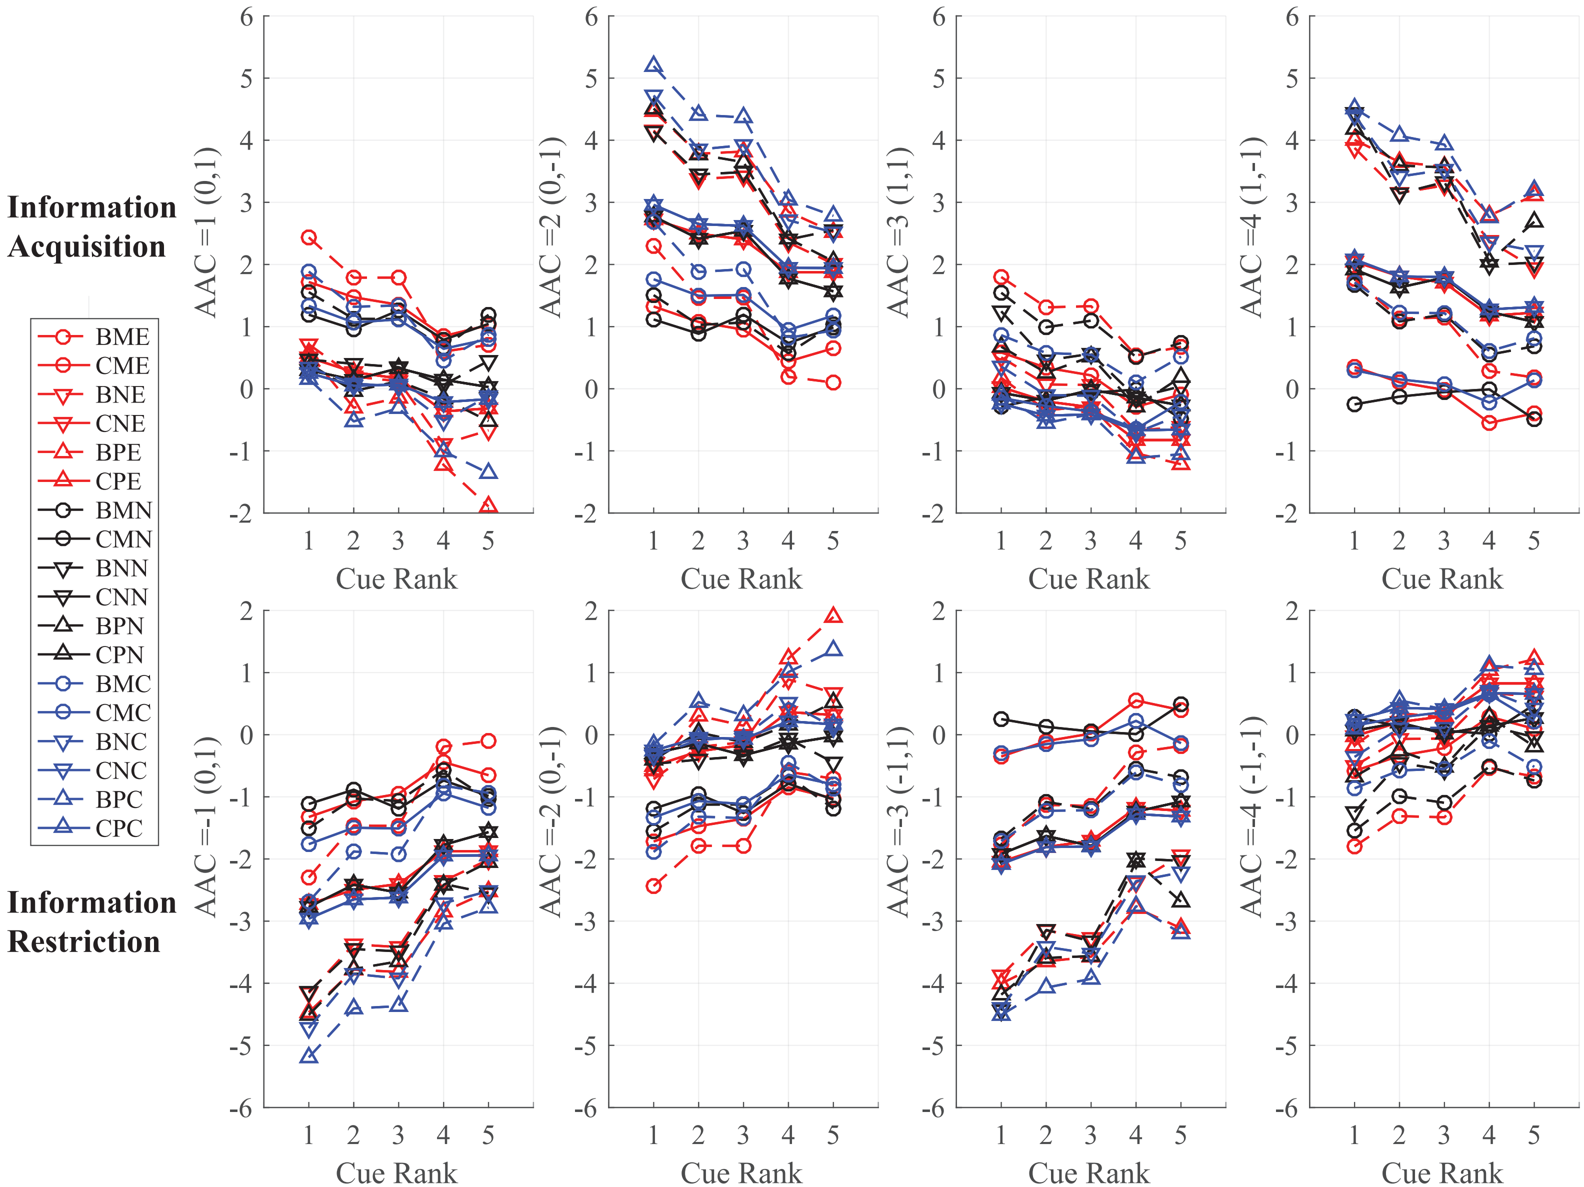

Figure 7.4: Average accuracy change (AAC) for each information acquisition and restriction type by cue rank. Parentheses denote the change in cue balance and option imbalance, respectively.

mates in this study), option imbalance becomes the main determinant of decision making accuracy, not total information.

In the context of training and decision support for situations with incomplete information, these results put the focus on how distributions of incomplete information are presented to decision makers and how decision makers can be supported to estimate missing cue values accurately. How exactly to modify and present information in accordance with minimizing option imbalance and maximizing cue balance is not yet known – particularly, with respect to hiding and restricting information. In terms of information acquisition, direct studies will be required to compare these heuristic information acquisition strategies to the analytic, Bayesian information acquisition methods (Nelson, 2005; Meder and Nelson, 2012).

Additionally it is still an open question as to how good people are at estimating missing information (Rieskamp et al., 2003; Rieskamp, 2006b). Garcia-Retamero and Rieskamp (2009) showed that people could quickly adapt their estimates of missing information from average estimates when information was uniformly missing to negative estimates when missing information was correlated with negative cue values. However, studies like Rieskamp (2006a) showed that people were slow to adapt to new environments, described as an "inertia effect." Another way to address the question of estimating the value of a missing cue is to frame it as a judgment problem as in Chap. 3 and von Helversen and Rieskamp (2009).

This issue of estimating missing information would be another useful point of synergy between the fast-and-frugal heuristics community (where most of the direct studies of estimates of missing information have been done) and the naturalistic decision making community (who study expert decision makers in familiar contexts) (Keller et al., 2010). One of the eight characteristics of naturalistic decision settings is uncertain, dynamic environments with incomplete and imperfect information (Orasanu and Connolly, 1993). Lipshitz and Strauss (1997, p. 156) showed that participants dealing with a lack of information typically used assumption-based reasoning to reduce uncertainty: "Construct a mental model of the situation based on beliefs that are (1) constrained by (though going beyond) what is more firmly known, and (2) subject to retraction when and if they conflict with new evidence or with lines of reasoning supported by other assumptions." Modeling or supporting this assumption-based reasoning might be a path toward better estimates of missing information and thus, better decision making with incomplete information.

Alternatively, instead of relying on training or a person's experience to estimate missing cue values, decision support tools might be able to use analytic methods for estimating the missing values (Arbuckle, 1996).

7.4.2 Information Acquisition Without Cue Weights, Cue Values, or Probabilities

Standard analytic information acquisition methods require reliable assessments of probabilities, cue weights, and cue scores (Nelson, 2005, 2008; Meder and Nelson, 2012). However, these requirements face multiple environmental issues when used in the real-world (Katsikopoulos and Fasolo, 2006; Katsikopoulos et al., 2008): First, real-world problems can exhibit statistical dependencies which create computational issues in calculating probabilities and cue weights. Second, psychologically, an operator may not be able to provide accurate information to the DSS regarding the probabilities and cue weights. Third, the analytic process may seem too complex and too opaque to the operator causing the operator to be reluctant to use the DSS or accept its suggestions. Lastly, as these methods have focused on acquisition, they do not provide suggestions as to what information could be restricted from the operator.

The rules examined in this study overcome these environmental issues to develop new methods of information acquisition and restriction. When the estimates do not match the environment, heuristic information acquisition and restriction rules should try to reduce option imbalance regardless of their effects on cue balance. However, when estimates do match the environment, heuristic information acquisition and restriction rules should acquire information to not affect cue balance. In general, the acquisition rules were shown to increase the accuracy of strategies despite not using the rank order of the cues, the cue values, or the probabilities. Accuracy increased more than 1% per acquisition for strategies with correct estimates of missing information and up to 4% per acquisition for strategies with incorrect estimates.

The restriction rules showed more mixed results. For strategies that estimated missing information correctly, those that used prior cutoffs can restrict information to reduce cue balance with negligible effects on accuracy, while those that used relative cutoffs should not restrict information. For strategies with incorrect estimates of missing information, restricting information to reduce option imbalance resulted in accuracy changes of plus-or-minus 0.5%. In sum these restriction results provide a more formal foundation to the prior results showing that specific cue values in a decision task can be removed without compromising accuracy (see Sec. 2.3).

Future work will be required to directly compare the effectiveness of these heuristic rules against analytic information acquisition methods which do use probabilities, cue weights, and cue scores.

7.5 Conclusion

The previous studies of decision making with incomplete information in Chap. 5 and Chap. 6 provided initial results toward new hypotheses of decision making performance with incomplete information: that there is more to measuring incomplete information than just total information; that there is more to strategy definitions than just their cue weights; and that there are ways to effectively acquire and restrict information without knowledge of cue weights and cue values. The results were labeled as initial because they had some methodological limitations such as the use of artificial datasets, few decision making strategies, and aggregate-level analysis. This study re-examined these hypotheses using a comprehensive study of seven real-world datasets, eighteen decision making strategies, and analysis at the task-level. The results confirmed each of these hypotheses while also providing deeper explanations. The balance of information between options and within cues, interact with total information to determine decision making accuracy. Matching estimates of missing information to the environment is essential for being robust in environments with incomplete information. Acquiring and restricting information heuristically can be an effective method for increasing or maintaining accuracy, respectively, even without the knowledge cue weights, cue values, or probabilities.

With these hypotheses confirmed, the next step is to determine how the distributions of incomplete information affect people's decision making accuracy. In particular, does having correct estimates of missing information make people's decision making accuracy robust to various distributions of incomplete information?

CHAPTER 8

VALIDATION HUMAN-SUBJECTS STUDY

Computational studies in prior chapters have presented fairly conclusive results showing the estimates of missing information and distributions of incomplete information are the main determinants of decision making accuracy with incomplete information regardless of decision strategy (Chap. 5-7). When the estimates of missing information match the environment (estimating missing information as average cue values when missing information is uncorrelated), only total information matters: more information increases accuracy, regardless of its distribution. However, when the estimates of missing information do not match the environment (estimating missing information as negative when when the missing information is uncorrelated), option imbalance and cue balance are more important that total information: decreasing option imbalance and increasing cue balance increases accuracy. This experiment attempts to address an important follow-up question: does having correct estimates of missing information make people's decision making accuracy robust to various distributions of incomplete information? [1]

8.1 Background

An experiment was designed based on the contextual model of decision making accuracy described in Chap. 4. The level of difficulty (task parameter) and incomplete information were manipulated directly in the experiment and serve as the independent variables. Participants were not controlled as to how they processed information (strategies) or estimated missing information. The decision tasks were designed to elicit how participants estimated missing information (positive, negative, average) but not how they processed information

[1] See Appendix A.3 for documentation of approval from Georgia Tech's Institutional Review Board and associated consent forms.

because prior computer simulation studies (Chap. 7 and Garcia-Retamero and Rieskamp, 2008) and human-subjects studies (Garcia-Retamero and Rieskamp, 2009) have shown that decision making performance with incomplete information is more determined by how people estimate missing information than how they weigh and integrate the information. To make the decision tasks independent of strategy, in each task, the traditional heuristic (take-the-best, TTB) and analytic (weighted-additive, WADD) strategies with average estimates of missing information would select the same option without selecting at random. The dependent variables measured by the output of each decision task were 1) the accuracy of the decision and 2) the time required to make the decision.

8.1.1 Difficulty

Difficulty is a task parameter which describes the magnitude of the difference between the criteria of the two presented options. As the difference between the two options decreases, the two options become more similar, and the difficulty of making the correct decision is presumed to increase. The difference between option scores varied from 0 to 1. Difficulty was categorized into 3 levels: low [0.60-0.85], medium [0.25-0.45], and high [0.05-0.13].[2]

8.1.2 Incomplete Information

Three measures are used to describe the distribution of incomplete information: total information, option imbalance, and cue balance (Sec. 4.3.4). Total information describes the number of known cues available to the decision maker. Option imbalance measures the difference in total known cues between two options. Cue balance counts the number of cues for which cue scores are known for both options. For a two-option, four-cue decision task, there are 22 unique distributions of incomplete information as described by these three measures. This experiment studied the 13 distributions in which 2, 4, 6, and 8 cue scores

[2]These categories do not cover the entire span of 0 to 1 because the categories only correspond to the range of differences represented in the selected strategies. Differences not categorized in these 3 levels were not used in the experiment.

were known.

The 13 distributions shown in Table 8.2 enabled the experimental results to be comparable to the simulation results of previous research of decision making with incomplete information. By studying all 13 distributions, accuracy can be studied relative to each individual measure of incomplete information and also the interactions. Previous research has shown that, in general, accuracy of decision making strategies increases when total information increases (Martignon and Hoffrage, 2002; Garcia-Retamero and Rieskamp, 2008; Chap. 5 and Chap. 7), option imbalance decreases (Chap. 5 and Chap. 7), and cue balance increases (Chap. 5 and Chap. 7).

The comparison within total information provided insights into how known cue scores should be distributed within a decision task. The computer simulation studies in Chap. 5 and Chap. 7 showed that with fixed amount of total information, cue scores should be distributed such that option imbalance is minimized and cue balance is maximized. For example, this study examines five different ways that 4 known cue scores could be distributed between two options, varying from 100% option imbalance to 100% cue balance.

8.1.3 Estimates of Missing Information

There at many ways that individuals estimate missing information of which three will used in this experiment: positive, i.e., replace the missing cue value with a high cue value (assuming the missing information positively correlated with the criterion), negative, i.e., replace the missing cue value with a low cue value (assuming the missing information negatively correlated with the criterion), or average, i.e., replace the missing cue value with the average value of the cue (assuming missing information is not correlated with the criterion). Garcia-Retamero and Rieskamp (2009) showed that decision makers have the capability to correctly adapt their estimates of missing information to situations when the missing cue scores tend to be negative and when the missing cue scores are uncorrelated with any particular value. Participants correctly matched the environments by estimating missing information as negative and average cue scores, respectively. This adaptive behavior is ecologically rational as it achieved the highest accuracy in those environments. In this experiment, decision tasks are designed such that the missing information is uncorrelated with average estimates resulting in the correct decision in every decision task for TTB and WADD.

8.1.4 Interaction of Estimates of Missing Information and Distributions of Incomplete Information

Estimates of missing information also have significant mediating effects on how distributions of incomplete information affect accuracy. Chapter 7 showed that when estimates of missing information match the environment, accuracy of a wide range of decision making strategies is only determined by total information and is unaffected by option imbalance or cue balance (see Fig. 7.2). When estimates do not match the environment, option imbalance and cue balance do mediate decision making accuracy. Since participants should be able to adapt to using the correct average estimates of missing information, the decision making accuracy of the participants should not be affected by changes in distributions of incomplete information measured by option imbalance and cue balance.

8.2 Method

8.2.1 Participants

30 people (24 men and 6 women) with an average age of 24 years old participated in the experiment, which took approximately 1.5 hours. Most students were from the School of Aerospace Engineering at the Georgia Institute of Technology, with a total of 19 graduate students and 11 undergraduate students participating. Participants received a show-up payment of \$10, and an additional payment based on their performance (up to a maximum total payment of \$25.60), with an average total payment of \$24.36.

8.2.2 Decision Environment

The decision environment was based on characteristics of the decision making environment of an anti-aircraft warfare coordinator (AAWC) on a military ship (Rothrock, 1995; Rothrock and Kirlik, 2003). Based on this background, the participants' goal as an AAWC is to determine which of two incoming targets is more dangerous (the decision criterion) and should be engaged, based on four binary cues: altitude (low, high), speed (slow, fast), distance from neutral air corridor (near, far), and size (small, large).

Constructing the relationship between the cue scores and criterion scores is an important consideration when studying decision making accuracy - both specifically related to human-subjects studies (Olsson et al., 2006; Pachur and Olsson, 2012) and generally with respect to the environmental parameters described in Chap. 4 (validity, predictability, and compensatoriness). The criterion model created for this experiment is a non-linear, non-monotonic model as shown in Eq. 8.1. The first parentheses of Eq. 8.1 generates the estimate of the danger of the incoming target based on a linear model of the target's characteristics: the target's altitude, speed, and size. When the binary cue values equal 1, the target is more dangerous than if the binary cue value is 0. Figure 8.1 visualizes the two exemplar targets. The missile, a small ($a_4 = 1$), high speed ($a_2 = 1$) target at low altitude ($a_1 = 1$) is more dangerous than the transport aircraft, a large ($a_4 = 0$), low speed ($a_2 = 0$) target at high altitude ($a_1 = 0$). To make the decision environment more difficult and more related to real-world tasks, the danger also had to be considered with respect to the target's distance from an air corridor used for non-military transports. If the target is near the corridor ($a_3 = 0$) the danger is 70% less than if it were far from the corridor ($a_3 = 1$).

$$C = (0.6 \cdot a_1 + 0.3 \cdot a_2 + 0.1 \cdot a_4) \cdot (0.3 + 0.7 \cdot a_3) \tag{8.1}$$

Table 8.1 show the 16 unique targets that were generated the binary cues. The decision environment is characterized by the following environmental parameters: cue weights (as

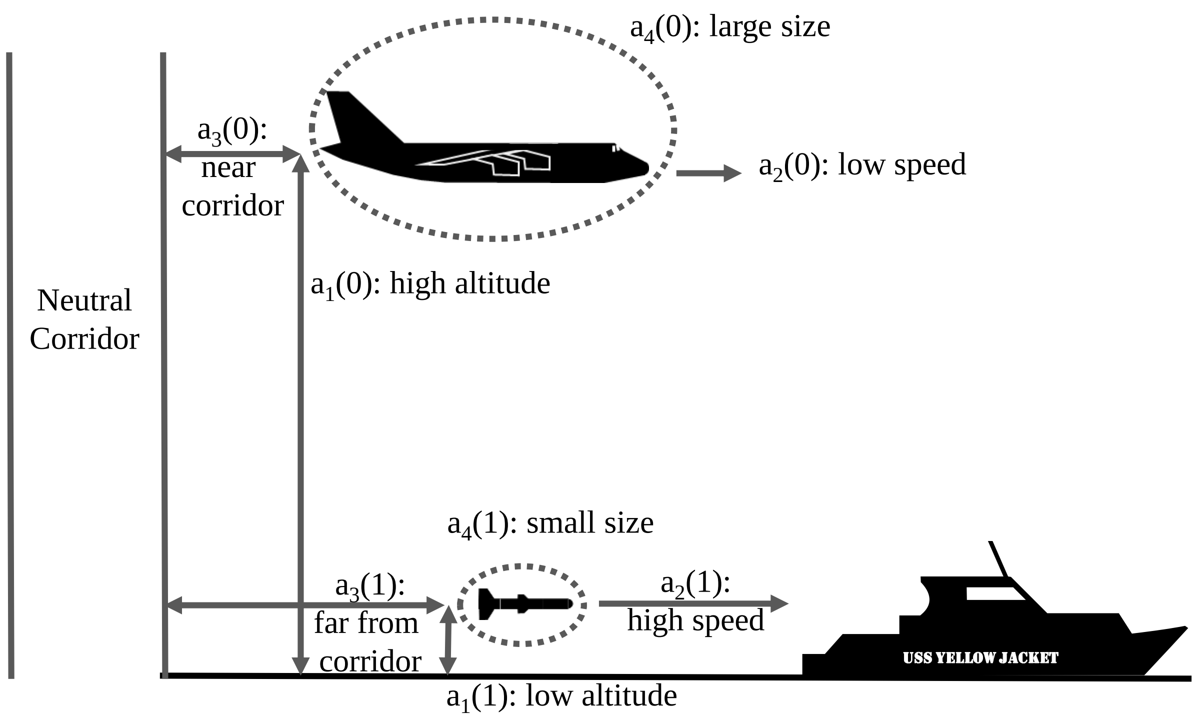

Figure 8.1: Visualization of the decision environment with reference to the two exemplar targets: a missile with a high level of danger and a transport with a low level of danger.

calculated by frequentist ecological validity) $(0.88, 0.73, 0.78, 0.59)$, compensatory class-5, variability (0.29), redundancy (0.00), and predictability (0.85).

To better mimic the context of a AAWC, instead of presenting values of 0 or 1 to the participants, the binary scores were transformed into context-relevant values based on Rothrock (1995) and Rothrock and Kirlik (2003), as shown in Table A.10. For example, size was mapped to radar cross-section with relevant values in units of m^2. The bounds of the cues were selected to be far enough apart that distinguishing between the two binary states would be clear. The specific tasks are provided in Table A.3.2.

8.2.3 Task Design

To appropriately study the interaction of estimates of missing information and distributions of incomplete information, the tasks were selected such that a participant's selected options will distinguish between their estimates of missing information and their ability to select

Table 8.1: Options and criterion scores.

Option No.	Cue 1, a_1 Altitude	Cue 2, a_2 Speed	Cue 3, a_3 Distance from Corridor	Cue 4, a_4 Size	Criterion Score (C) Danger
1	1	1	1	1	1.00
2	1	1	1	0	0.90
3	1	1	0	1	0.30
4	1	1	0	0	0.27
5	1	0	1	1	0.70
6	1	0	1	0	0.60
7	1	0	0	1	0.21
8	1	0	0	0	0.18
9	0	1	1	1	0.40
10	0	1	1	0	0.30
11	0	1	0	1	0.12
12	0	1	0	0	0.09
13	0	0	1	1	0.10
14	0	0	1	0	0.00
15	0	0	0	1	0.03
16	0	0	0	0	0.00

options correctly. Within each decision task used in this study, the correct decision can be made by estimating missing information as the average, in between the two limits of the high and low cue values. The process for generating and analyzing these decision tasks is outlined in Fig. 8.2. The goal is to generate two decision tasks for each combination of incomplete information (Table 8.2) such that an error matrix can be constructed from the results with two dimensions of the decision maker's accuracy (correct or incorrect) and two estimates of missing information (negative estimates make correct decisions or positive estimates make correct decisions). These pairs of decision tasks are constructed such that, for the same pair of options, flipping the distribution of incomplete information results a reversal of the positive or negative estimates selecting correctly. From this pair of decisions, participants' strategies can be categorized in one of four ways: 1) negative estimates of missing information, 2) positive estimates of missing information, 3) average estimates of missing information with accurate decisions, and 4) average estimates of missing infor-

1 **Two decision tasks with the same two options but with flipped distributions of incomplete information.**

Decision Task 1

	Cues			
	1	**2**	**3**	**4**
1A – Correct	1	0	0	0
1B - Incorrect	**? (0)**	0	0	**? (0)**

Decision Task 2

	Cues			
	1	**2**	**3**	**4**
2A - Correct	**? (1)**	0	0	**? (1)**
2B - Incorrect	0	0	0	0

2 **Choosing an option in both tasks elicits two characteristics of the strategy: the estimate of missing information and the accuracy. The strategy can then be classified into four categories.**

Negative Estimates

Participant Chooses

	Correct	Incorrect
Negative Estimates Correct	1A	0
Positive Estimates Correct	0	2B

Average Estimates – Correct

Participant Chooses

	Correct	Incorrect
Negative Estimates Correct	1A	0
Positive Estimates Correct	2A	0

Positive Estimates

Participant Chooses

	Correct	Incorrect
Negative Estimates Correct	0	1B
Positive Estimates Correct	2A	0

Average Estimates – Incorrect

Participant Chooses

	Correct	Incorrect
Negative Estimates Correct	0	1B
Positive Estimates Correct	0	2B

Figure 8.2: Designing decision tasks to elicit estimates of missing information, information bias, and accuracy.

mation with inaccurate decisions. For the five combinations of incomplete information in which estimates could not be elicited, tasks were selected such that all strategies (positive, negative, average) select correctly.

Table 8.2: 13 distributions of incomplete information examined in the experiment.

Total Information	Option Imbalance	Cue Balance	Decision Task Measures: Accuracy	Decision Task Measures: Estimates
2	0	0	●	○
2	0	1	●	○
2	2	0	●	●
4	0	0	●	●
4	0	1	●	●
4	0	2	●	○
4	2	0	●	●
4	2	1	●	●
4	4	0	●	●
6	0	2	●	●
6	0	3	●	○
6	2	2	●	●
8	0	4	●	○

8.2.4 Materials and Procedure

For each decision task, participants were required to determine which of two targets was more dangerous. Both targets were presented simultaneously as a set of 4 cues with their corresponding cue weights, as described in Sec. 8.2.2: altitude (0.88), speed (0.73), distance (0.78), and size (0.59). The decision tasks were presented and answered using a computer interface as shown in Fig. 8.3. The two targets and four cues were displayed in a matrix of two rows and four columns. All cue scores that were designated as known cue values were presented for the entire duration of the task. Cues scores were not searched for, unlike other similar studies (Rieskamp and Hoffrage, 2008).

In the instructions, participants were told: “You will be asked to take on the role of an anti-aircraft warfare coordinator (AAWC) with the duty to determine which of two incoming targets is more dangerous and should be engaged by your weapons system. The data collection phase of the study consists of 156 decision tasks of varying difficulty and level of missing information. Within each decision task you will have 20 seconds to determine which of the two incoming targets is more dangerous - there is always one target which is

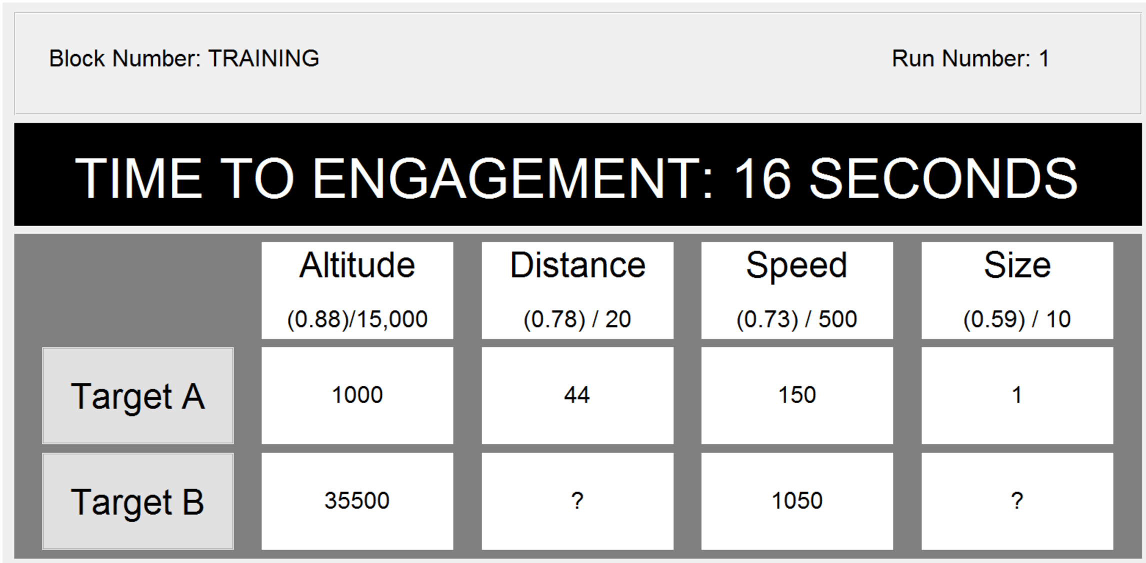

Figure 8.3: Example of the interface used in Experiment 1.

more dangerous. In addition to the show-up payment of $10, you will earn $0.10 for each correct answer. Therefore, you have a chance to earn up to $25.60." After each decision, the interface indicated whether the participants choice was correct or incorrect, or that the participant's time expired. The decision task time limit was selected at 20 seconds for both experimental and operational relevance. A previous study of time pressure by Rieskamp and Hoffrage (2008) showed that participants experiencing high time pressure made decisions in approximately 20 seconds for a 4-option, 5-cue decision task. With respect to the operational constraints of integrated ship defense this time limit is at (or even slightly above) the upper bound of reasonable "man-in-the-loop" processes due to the Earth's curvature, the sensor antenna height, and energy propagation effects (Prengaman et al., 2001, p. 529).

The cues, cue scores, cue weights, thresholds, and cue directions were described to each participant. For example, the altitude cue was described as: "The height of the target above sea level. Low altitude targets are considered to be more dangerous than high altitude. Experienced AAWCs use 15,000 ft as the threshold. The information will be presented in a table format with the most useful variables decreasing from left to right. The interface presents information about usefulness and the threshold next to the titles for each variable. Usefulness is defined in the following way: for example, when comparing targets to determine which is more dangerous, if you always choose the target with the lower altitude, you would make a correct decision in approximately 88% of the decision tasks." As told to the participants, the cue weights and thresholds were displayed within the interface and the cues were ordered from left to right with decreasing weights as shown in Fig. 8.3. This information was provided so that the variability between participant performance would be largely based on incomplete information and how participants estimated missing information.

The participants were then trained until their accuracy plateaued so as to minimize learning effects during the data collection phase. Participants were presented with blocks of 11 medium-difficulty decision tasks accounting for all 11 distributions of incomplete information in Table 8.2 in which each variant of WADD and TTB strategy and estimates would select the correct option. This excludes 2 distributions of incomplete information: 2 total information with 2 option imbalance, and 4 total information with 4 option imbalance. In these distributions the strategies with negative estimates will never select the same option as the positive estimates. The participants were considered to have completed learning if the number of correct decisions in each of the previous 3 completed blocks differ by one or fewer.

Each participant was trained within the first three blocks of training. Meaning that for the first three blocks, their maximum and minimum block accuracy were different by no more than one decision. The 30 participants' accuracy after the three blocks (33 tasks) were distributed as: 100% ($n = 20$); 97% ($n = 4$); 94% ($n = 3$); and 88% ($n = 1$).[3]

The data collection phase consisted of 6 blocks of 26 decision tasks. Within each block there were 13 combinations of incomplete information each with 2 variations of correspondence accuracy to determine the participants mechanism of estimating missing information

[3]Two participants' training data was not saved due to computer malfunction.

and accuracy. Three blocks were constructed for each level of difficulty which was measured by the magnitude of the difference between the option scores. The resulting experiment has 13 levels of incomplete information, 2 accuracy-estimate variations, 2 levels of difficulty, and 3 repetitions for a total of 156 decision tasks. To reduce order effects, the order of blocks and the order of tasks within each block were counterbalanced using Latin hypercube designs.

8.3 Results

The first analysis identified the largest set of blocks without significant learning effects on accuracy and whether participants appropriately adapted their estimates to the environment. Then, difficulty was examined as a potential mediator of decision making accuracy and time required. Lastly, accuracy and time required were analyzed with respect to the distributions of incomplete information in one-variable, two-variable, and three-variable combinations.

The analyses of the results are reported as follows. Analyses with accuracy as the dependent variable used χ^2 tests and analyses with time required as the dependent variable used ANOVA tests. Significant results are reported at the $\alpha = 0.001$ level. One standard error bands are calculated as $\sigma/\sqrt{n}$.

8.3.1 Did Participants Adapt Their Estimates of Missing Information?

Participants were shown to best match strategies with average estimates throughout each block: assuming missing information was a value between the limits of high and low categories (see Table 8.3). These estimates were identified by comparing the participants' decisions to how each mechanism for estimating missing information would have decided. Using outcome-based estimate identification only considers the participants' decisions and not the search processes (Garcia-Retamero and Rieskamp, 2009; Rieskamp and Hoffrage, 2008). Since using average estimates was the only mechanism that would result in correct decisions in every task (see Table 8.3), these results provide further evidence for the hy-

Table 8.3: Number of tasks in which strategies using different mechanisms for treating missing information make different predictions.

Task Type	**Prediction**			**Number of Discriminating Tasks (All Blocks, Blocks 3-6)**
Treating missing information as negative vs. as the average or positive	$\text{TTB}_{negative}$ $\text{WADD}_{negative}$	vs.	$\text{TTB}_{positive}$ $\text{TTB}_{average}$ $\text{WADD}_{positive}$ $\text{WADD}_{average}$	(48, 32)
Treating missing information as positive vs. as the average or negative	$\text{TTB}_{positive}$ $\text{WADD}_{positive}$	vs.	$\text{TTB}_{negative}$ $\text{TTB}_{average}$ $\text{WADD}_{negative}$ $\text{WADD}_{average}$	(48, 32)
Correct use of TTB or WADD using any estimate of missing information	Correct $\text{TTB}_{positive}$ $\text{TTB}_{negative}$ $\text{TTB}_{average}$ $\text{WADD}_{positive}$ $\text{WADD}_{negative}$ $\text{WADD}_{average}$	vs.	Incorrect	(60, 40)

pothesis that participants adapt their estimates to their environment (Garcia-Retamero and Rieskamp, 2009).

Though participants were trained to a consistent accuracy, participants' accuracy and time required still showed learning effects during the data collection phase as shown in Fig. 8.4. An ANOVA with predicted inferences as the dependent variable and block order as the dependent variable, showed that block order across all six blocks had significant effects on the participants' use of positive [$F(5, 4671) = 3.65$, $p < 0.003$], negative [$F(5, 4671) = 3.87$, $p < 0.002$], and average estimates [$F(5, 4671) = 14.83$, $p < 0.0001$].

Given that participants were learning, ANOVA tests were conducted for each consecutive series of blocks with Block 6 (2 through 6, 3 through 6, 4 through 6, and 5 and 6) to determine at what point participants started to use the average estimates consistently. Blocks 3 through 6 was the largest set for block order have no effect on participants' use of

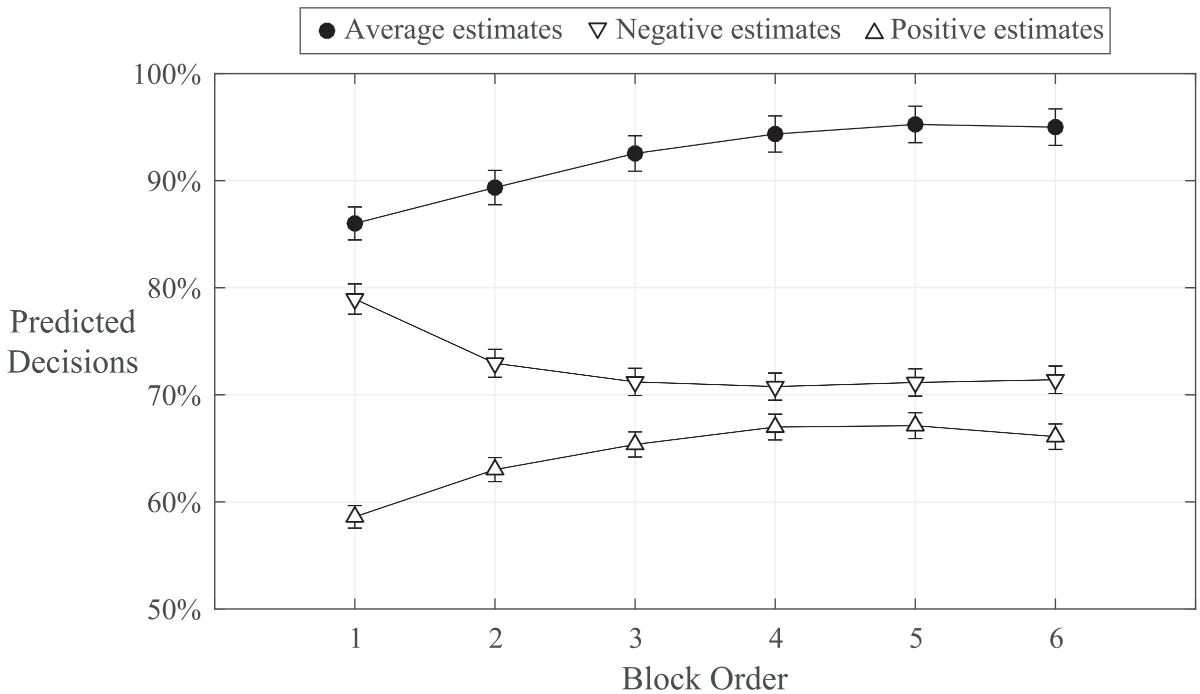

Figure 8.4: Effect of block order on estimate used. For a participant using average estimates, the average estimates would predict decisions in 100% of the tasks, and positive and negative estimates would predict decisions in 69% of the tasks. Error bars represent one standard error.

positive [$F(3, 3114) = 0.24$, $p < 0.87$], negative [$F(3, 3114) = 0.03$, $p < 0.99$], or average estimates [$F(3, 3114) = 2.17$, $p < 0.09$]. Block order for blocks 2 through 6 did have a significant effect on participants' use of average estimates [$F(4, 3893) = 7.52$, $p < 0.0001$], but not on positive [$F(4, 3893) = 0.97$, $p < 0.42$] or negative estimates [$F(4, 3893) = 0.27$, $p < 0.90$]. Therefore, only the blocks 3-6 will be used in all following analyses.

For blocks 3 through 6, there was no significant effect of block order on accuracy [$\chi^2(3,3120) = 7.69$, $p < 0.05$] but there was an effect on time required [$F(3,3116) = 19.06$, $p < 0.0001$]. As shown in Fig. 8.5, there was a clear increase in accuracy (from 85.9% to 95.0%) and decrease in time required (from 5.54s to 3.93s) as participants faced more tasks from block 1 to block 6.

The likely reason for the effect of block order on accuracy and time required was the difference in the accuracy-estimate task type between the training and the data collection phase of the experiment – in fact, this is the only major difference between the two phases.

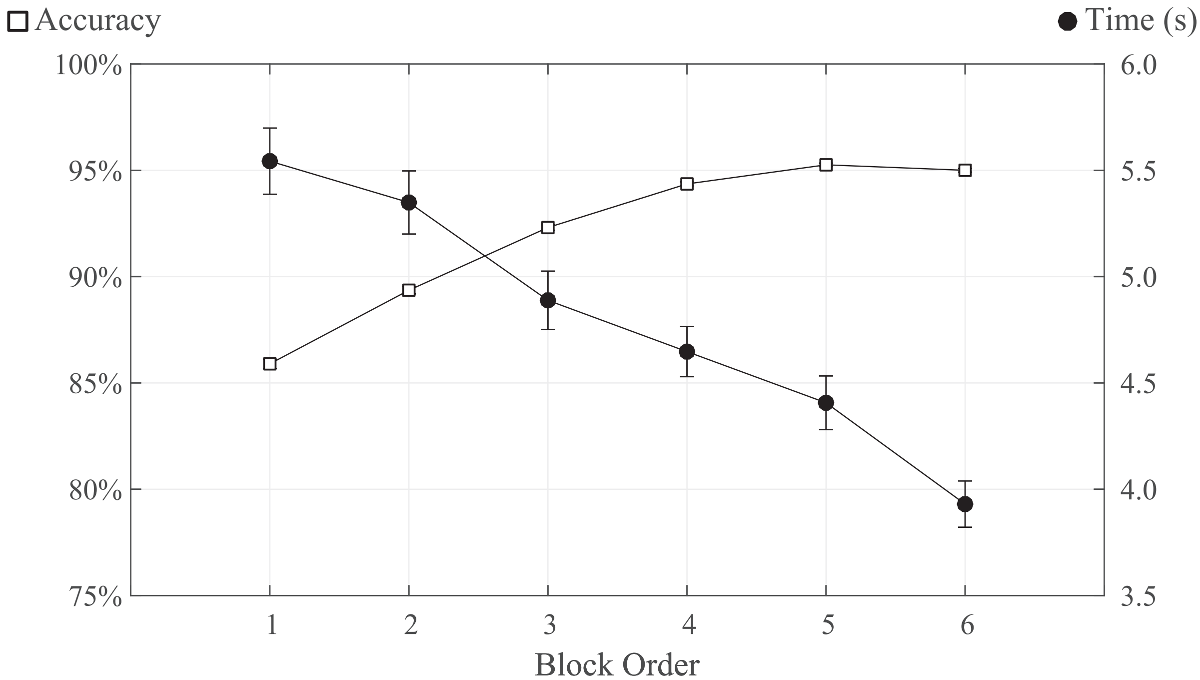

Figure 8.5: Effect of block order on accuracy and time required. Error bars represent one standard error.

In the training phase, tasks were selected such that all three estimate mechanisms would select correctly in each task. However, in the data collection phase only average estimates would select correctly in each task, with positive or negative estimates being accurate in 69% of tasks. Previous studies have shown that participants adapt their estimate mechanisms to the environment (Garcia-Retamero and Rieskamp, 2009). Therefore, participants would not have been required to adapt their estimate mechanisms in the training, only once the data collection phase began. By plotting the percent of participants' decisions predicted by each estimate in each block (Fig. 8.4), it is clear that the participants were correctly adapting from a combination of negative and average estimates to only average estimates. Therefore, in future experiments, the training tasks should be selected in the same manner as data collection tasks in order to restrict learning to the training phase.

8.3.2 Does Computational "Difficulty" Result in Human Decision Making "Difficulty"?

Difficulty was postulated as an environmental parameter that could mediate the performance of human participants. Results showed that difficulty had no significant effect on accuracy [χ^2(1,3120) = 5.03, p $<$ 0.03] or time required [F(1,3118) = 2.65, p $<$ 0.1]. The lack of effect is likely because what may be difficult according to the difference in criterion value of the options may not be an observable difference for the participants. This is particularly true if the participants used heuristic strategies like take-the-best which do not calculate or estimate an overall criterion value, unlike weighted-additive.

8.3.3 Did Distributions of Incomplete Information Affect Decision Making Performance?

Participants were shown to use accurate, average estimates in greater than 94.3% of their decision tasks in blocks 3 through 6, but the variance in accuracy and time required were not explained by block order or difficulty. Based on the computational studies of decision making with incomplete information in Chap. 7, because participants used average estimates, total information should be the only measure of incomplete information to have any significant effect on accuracy. These results show that total information did have significant effects, but so did option imbalance and cue balance (Table 8.4).

Starting with the original measure of incomplete information, total information did have a significant effect on accuracy [χ^2(3,3120) = 19.84, p $<$ 0.0002] and time required [F(3,3116) = 21.45, p $<$ 0.0001]. However, accuracy had a slight U-shape and the time required plot is an inverted-U shape as shown in Fig. 8.6. From the perspective of total information as a solitary measure of incomplete information, as total information increases, accuracy and time required should increase monotonically. Computational representations exemplify this assumption (Chap. 5 and Chap. 7; Payne et al., 1990; Garcia-Retamero and Rieskamp, 2008). However, the U-shaped curves suggest that accuracy and time required are being affected by factors in addition to total information.

The U-shaped curves can be explained as the effect of changing distributions of incom-

Table 8.4: The average accuracy and time required for each of the 13 distributions of incomplete information measured, sorted by accuracy.

Total Information	**Option Imbalance**	**Cue Balance**	**Time Required (Sec)**	**Accuracy**
2	0	1	3.05	99.2%
4	0	2	3.52	99.2%
6	0	3	3.80	99.2%
8	0	4	4.08	99.2%
2	0	0	3.78	98.3%
4	0	1	4.65	97.1%
4	2	0	4.99	96.3%
6	0	2	5.19	95.8%
4	4	0	4.72	95.0%
4	0	0	4.88	94.2%
2	2	0	4.76	86.3%
4	2	1	5.27	85.4%
6	2	2	5.41	80.0%

plete information. Accuracy and time required were both significantly affected by both option imbalance [χ^2(2,3120) = 137.12, $p < 0.0001$; F(2,3117) = 47.39, $p < 0.0001$] and cue balance [χ^2(4,3120) = 30.49, $p < 0.0001$; F(4,3117) = 8.26, $p < 0.0001$], respectively. For option imbalance, the highest accuracy (97.7%) and lowest time required (4.12 sec) occurred when option imbalance was zero. For cue balance, the average accuracy and time required with 0, 1, or 2 cue balance was 93.1% and 4.55 seconds, but participants showed increased accuracy (99.1%) and decreased time required (3.94 sec) when cue balance was higher (3 or 4).

A more thorough explanation of the single-variable effects lies in the analysis of two-variable interactions of the measures of incomplete information as shown in Fig. 8.7. Once again, even though the average estimates are accurate and should make accuracy only a function total information, the accuracy and effort results in Fig. 8.7 are clearly affected by option imbalance and cue balance, matching computational results of inaccurate estimates in Chap. 7. In both two-variable interactions with total information, the maximum accuracy and minimum time required occurred simultaneously and occurred independently of total

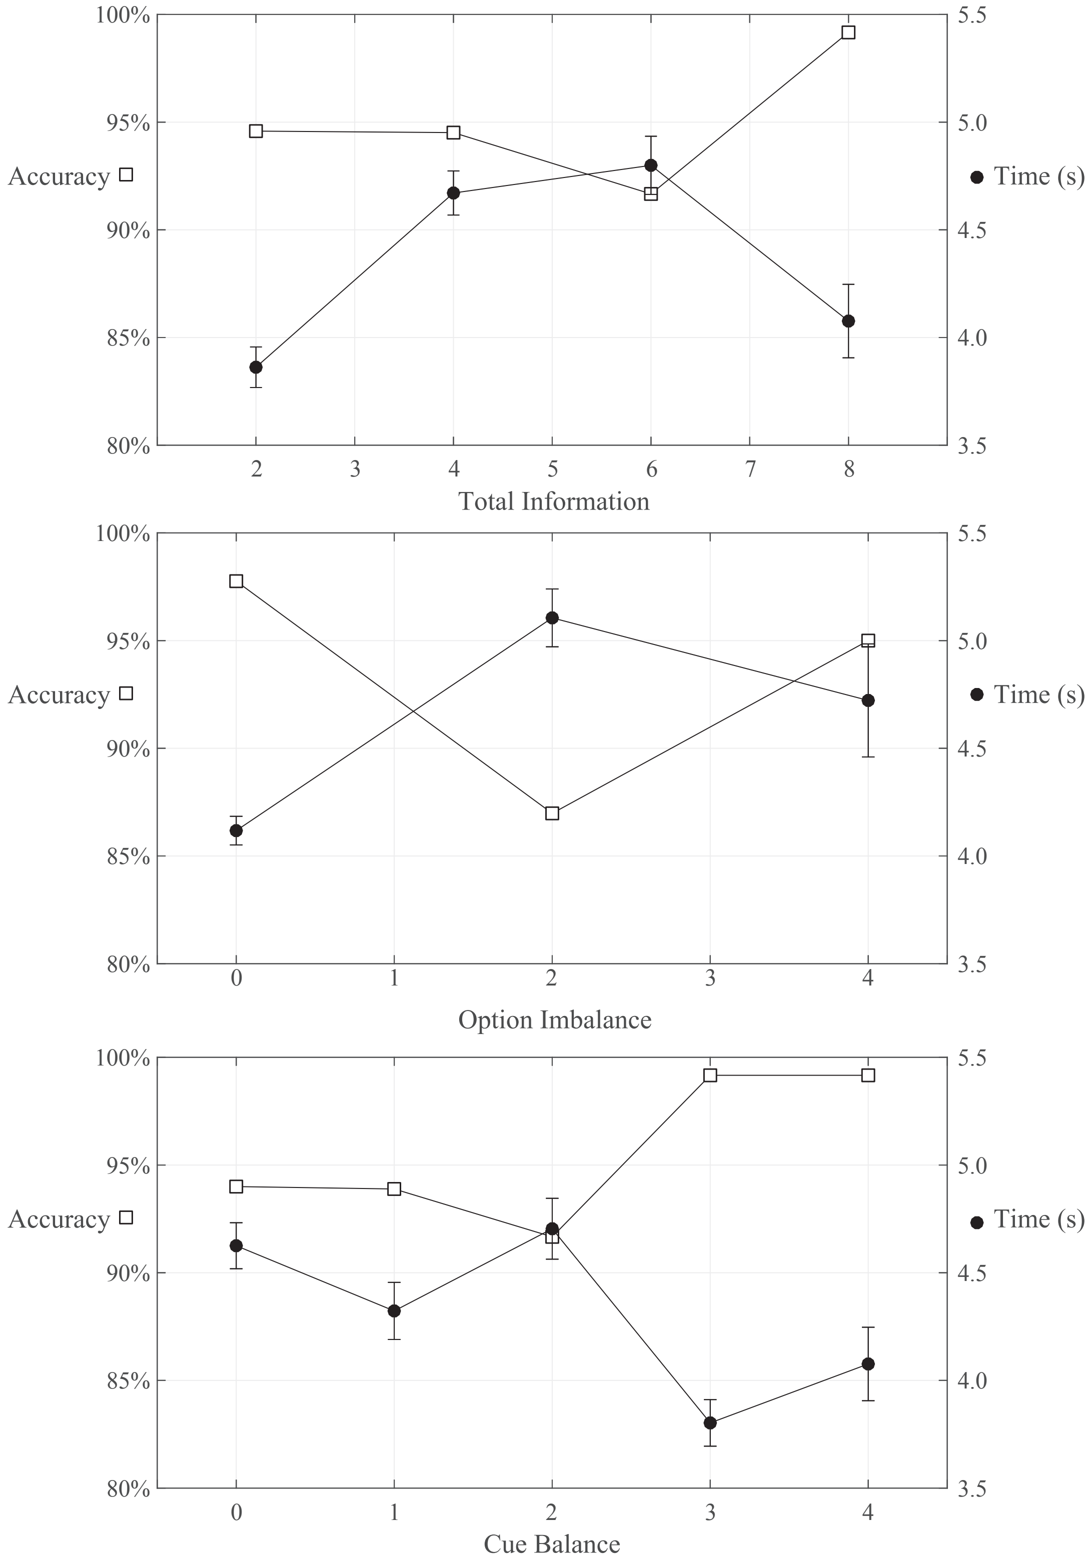

Figure 8.6: Effect of total information, option imbalance, and cue balance on participants' decision making accuracy and time. Error bars represent one standard error.

information, when option imbalance was minimized and when cue balance was maximized. Comparing cue balance and option imbalance directly showed that minimizing option imbalance has a more positive effect on accuracy and time required than maximizing cue balance.

Ultimately, the specific values of accuracy and time required are determined by the collective interaction of the three variables of the distribution of incomplete information (total information, option imbalance, and cue balance). Figure 8.8 displays a more nuanced and accurate picture of how accuracy and time required are affected by total information that cannot be seen in one-variable (Fig. 8.6) or even two-variable (Fig. 8.7) analyses. The 13 distributions of incomplete information are categorized into three groups:

– Green: These four distributions have cue balance maximized within each level of total information [(2,0,1); (4,0,2); (6,0,3); (8,0,4)]. Each distribution resulted in the same highest average accuracy, 99.2%, for any distribution. These distributions minimized time required within each level of total information.

– Yellow: These six distributions do not have the maximum cue balance, but, in all but one case, they do not have the maximum option imbalance either: [(2,0,0); (4,0,1); (4,2,0); (4,4,0); (4,0,0); (6,0,2)]. Each distribution showed reduced accuracy of 3% to 5% from the green group distributions and an average increase in time required of 1 second.

– Red: These three distributions have the maximum option imbalance for two and six cue scores [(2,2,0); (6,2,2)], and an option imbalance of two and cue balance of one for four cue scores [(4,2,1)]. These three had lowest average accuracy (below 87%) and the longest time required within each total information. Within this group, as total information increased, accuracy decreased and time required increased.

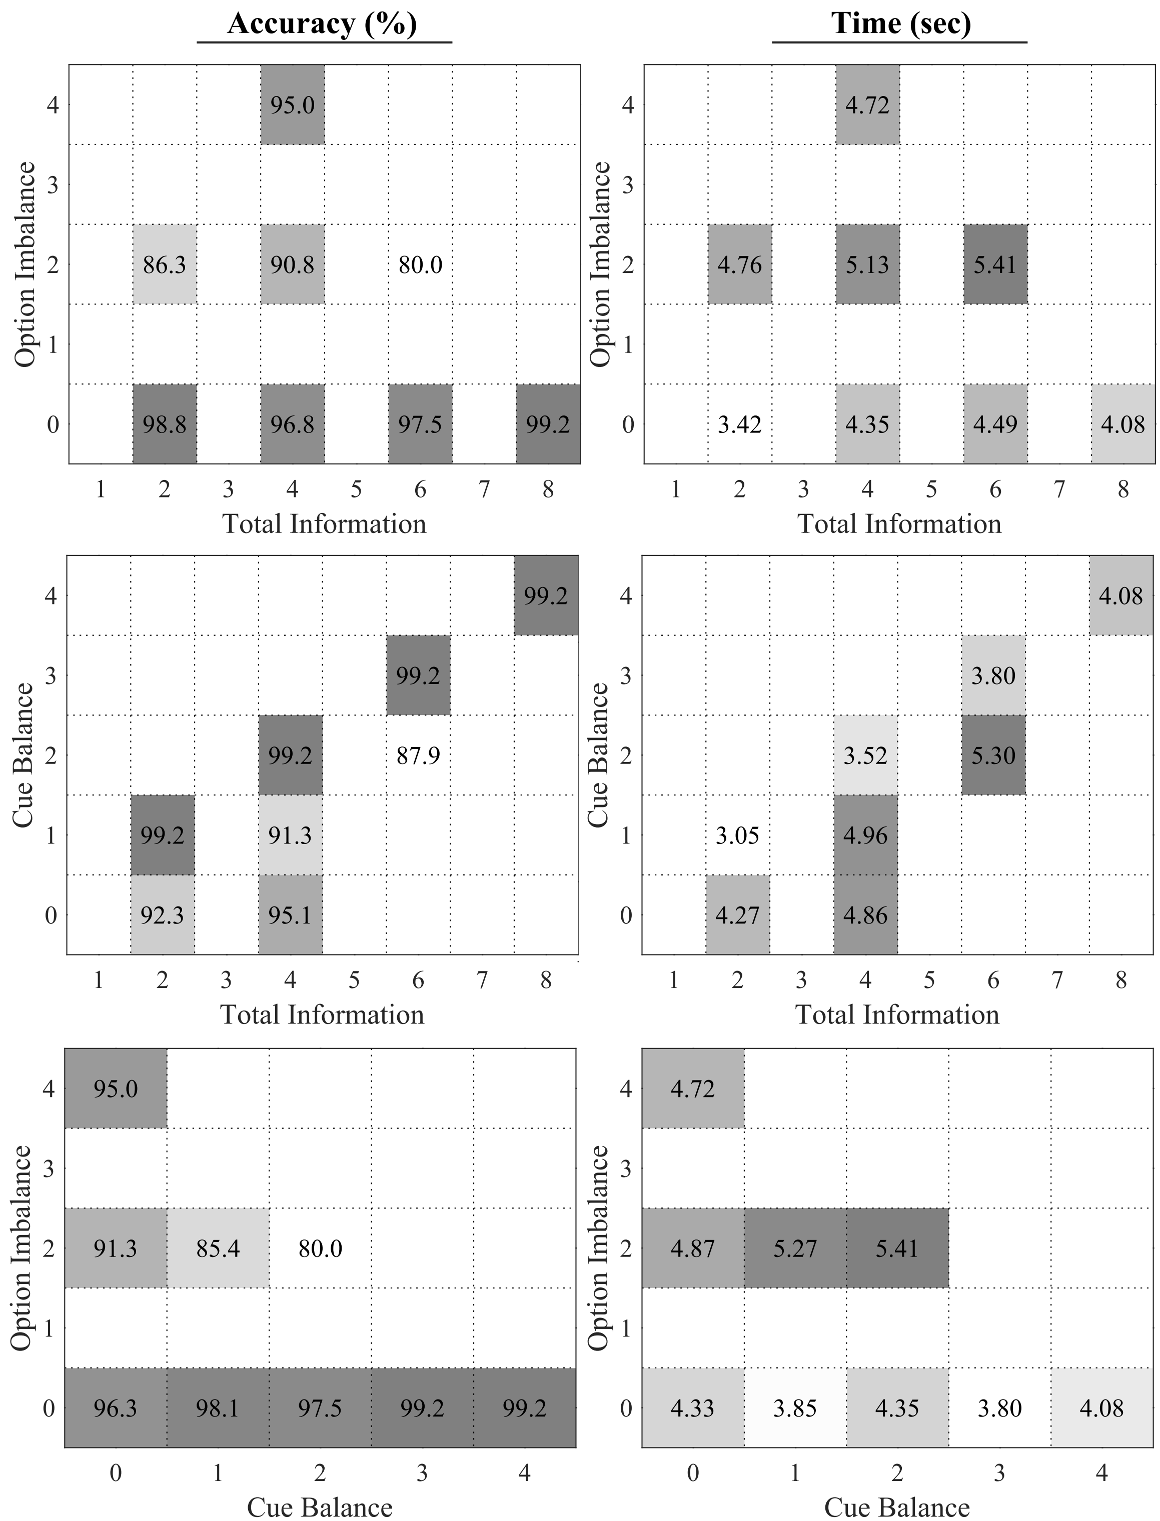

Figure 8.7: Two-variable interactions between the three measures of incomplete information with performance measured in accuracy (left column) and time required (right column).

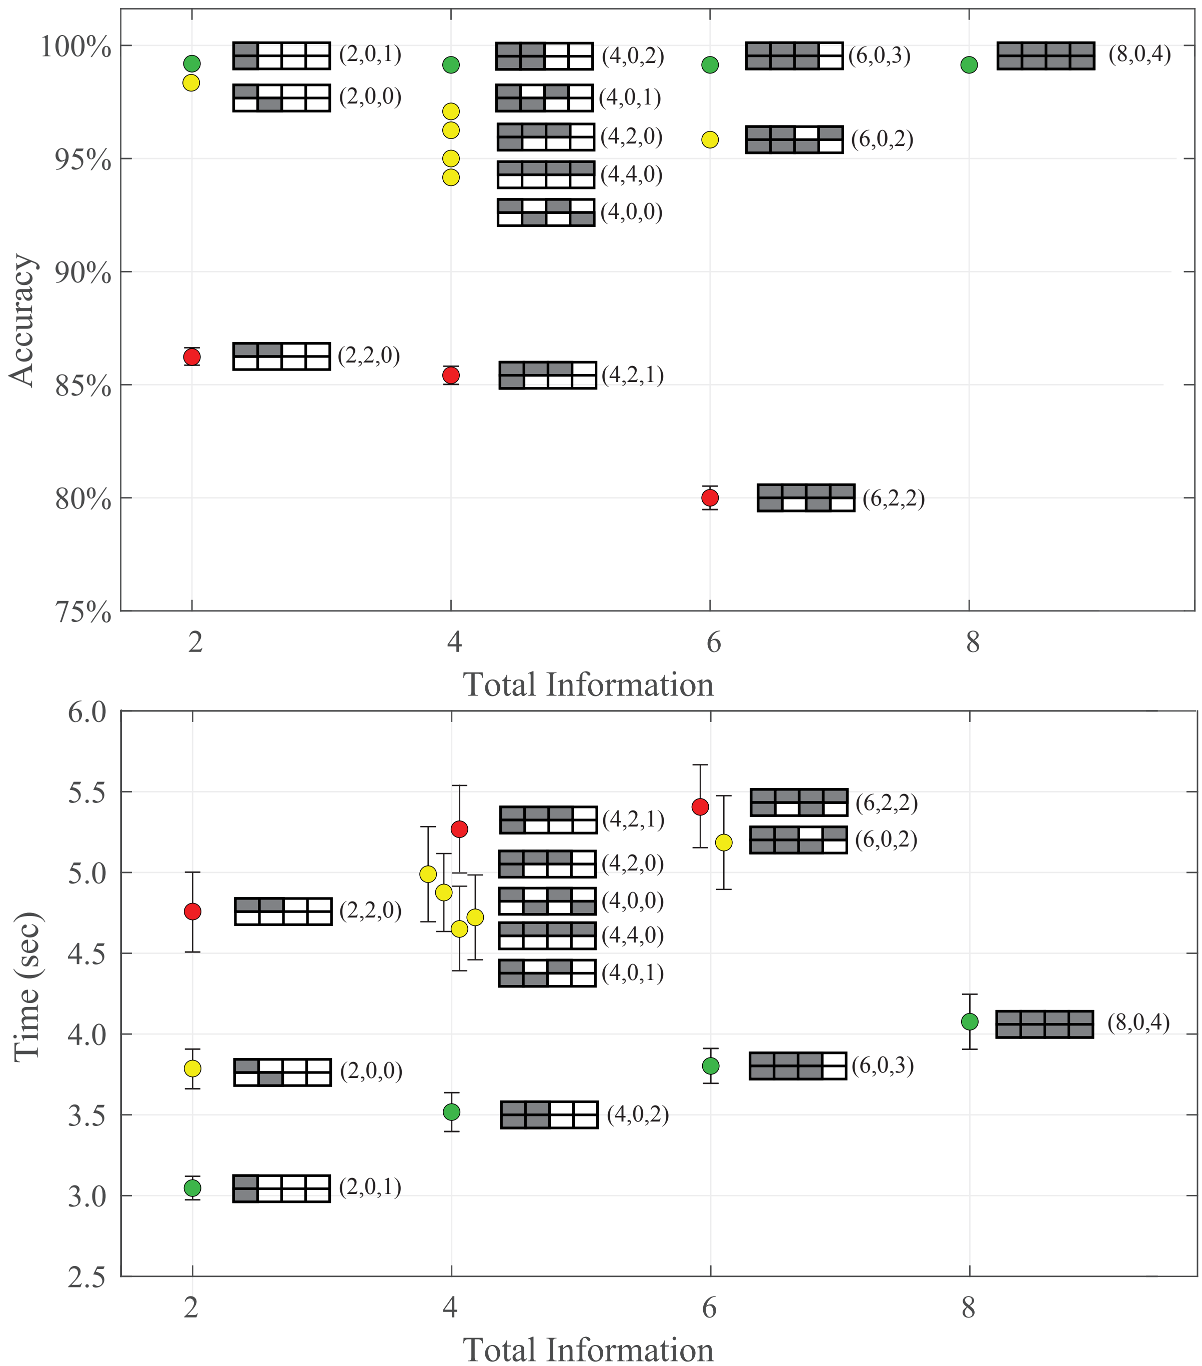

Figure 8.8: Interaction of the distribution of incomplete information (total information, option imbalance, and cue balance) on participants' decision making accuracy and time required, as a function of total information. Exemplar decision tasks with visual distributions of incomplete information are provided for each data point. The dark blocks indicate known cue scores and white blocks indicate unknown cue scores. Error bars represent one standard error.

8.4 Discussion

Did the conclusions from computer simulations in Chap. 7 about the interaction between estimates, distributions, and performance extend to human decision makers? Participants successfully adapted to using correct, average estimates of missing information, but the effects of the distributions of incomplete information matched the computational results of decision makers with incorrect estimates of missing information. This raises some important questions about a potential reality gap between computational models of human decision making with incomplete information and the humans themselves (Sec. 8.4.1), whether certain distributions are just 'unfriendly' or 'friendly' regardless of estimates of missing information (Sec. 8.4.2), and how future work can generalize the current results (8.4.3).

8.4.1 Falling into the Reality Gap

A reality gap occurs when virtual models struggle to reproduce the stochasticity of the real world. Though this term is pulled from the robotics community (e.g., Jakobi et al., 1995), the reality gaps have consistently been the motivational hurdles in the field of decision making. Bounded rationality was posited by Simon (1955) as a counterpoint to the prevailing models of humans as a completely rational 'economic man.' Tversky and Kahneman's heuristics-and-biases research characterized just how unlike economic models humans are, finding that people do not necessarily judge uncertainty according to the rules of probability and statistics, since they are influenced by framing and other cognitive biases (e.g., Tversky and Kahneman, 1974; Kahneman et al., 1982; Kahneman and Tversky, 1984). Then, two additional decision making paradigms have been developed to address what other researchers believed were limitations in the heuristics-and-biases research to explain how people used the bounds on their rationality in successful ways. The naturalistic decision making framework, showed that people faced with time pressure, uncertainty, ill-defined goals, and other complexities in familiar and meaningful environments, can perform very well by using various strategies such as matching strategies to situations, experience-based knowledge modeling, and implementable-based strategy selection (e.g., Klein et al., 1988; Orasanu and Connolly, 1993; Klein and Calderwood, 1996; Klein, 2008; Lipshitz et al., 2001). Furthermore, the fast-and-frugal heuristics framework has shown that people use the bounds on rationality of simplicity, speed, and frugality as a mechanism for simple, robust, and accurate strategies that adapt to the environment and ecology (e.g., Gigerenzer and Goldstein, 1996; Gigerenzer et al., 1999; Gigerenzer and Gaissmaier, 2011; Todd et al., 2012).

Following in the succession of reality gaps between mathematical models, computational simulations, and human-subjects studies, decision making with incomplete information is simply a new one. According to the computational simulation in Chap. 7, when estimates match the environment, only total information should really affected accuracy, and when the estimates do not match the environment, option imbalance and cue balance should affect accuracy more so than total information. However, the results presented in this study suggest that, in reality, there might only be one classification: regardless of whether a decision maker's estimates of missing information matches the environment, their accuracy will be affected by the distribution of incomplete information. In this study, decreasing option imbalance and increase cue balance generally increased accuracy and decreased time required.

These results were unexpected as the mathematical model of decision making used in the computational simulation in Chap. 7 accounts for a very broad range of decision making strategies that have been shown to generally match participant's strategies (Chap. 3 and Chap. 4). Notably, the human-subjects study by Garcia-Retamero and Rieskamp (2009), which explicitly studied estimates of missing information, did not report accuracy or time required as a function of estimates of missing information and thus did not report any difference between that study and their corresponding computer study (Garcia-Retamero and Rieskamp, 2008). At this point, it may be that even when participants are shown to match their estimates of missing information to the environment, they should still be supported and modeled as if their estimates do not match the the environment. At least until the general linear model, and the broad number of strategies that model represents, can account for the relative difficulty of certain types of distributions of incomplete information.

8.4.2 Friendly and Unfriendly Distributions of Incomplete Information

The results presented in this study suggest that some distributions of incomplete information are simply difficult, even when a decision maker's estimate of missing information matches the environment. Leveraging the language of Shanteau and Thomas (2000), there seem to be 'friendly' and 'unfriendly' distributions of incomplete information. Friendly distributions have low levels of option imbalance and high levels of cue balance, whereas unfriendly distributions have high levels of option imbalance and low levels of cue balance. These definitions leave out total information because at this point it seems to be a secondary mediator of decision making performance as compared to option imbalance and cue balance. As shown in Fig. 8.8, any friendly distribution (denoted by the color green) had a higher accuracy and lower time required than any unfriendly distribution (denoted by the color red), regardless of total information.

If these trends can be generalized through repeated studies and different task environments, there are significant implications for decision support. The goal of decision support systems would then be to modify the presentation of information from unfriendly to friendly distributions. Within most unfriendly distributions, there is a friendly distribution with less total information. Similarly, when starting from an unfriendly distribution, there is almost always a clear path toward a friendly distribution through acquiring information.

Within this perspective, the computational results in Chap. 6 and Chap. 7 showed that the effectiveness of heuristic rules for acquiring and restricting information depended on whether the estimates matched the environment or not. When estimates of missing information did not match the environment, the accuracy was significantly reduced when option imbalance increased and cue balance increased; however, when estimates did match the environment, the accuracy was almost exclusively dependent on total information. Therefore, the conclusions were that when the estimates do not match the environment, heuristic information acquisition and restriction rules are quite useful and should try to reduce option imbalance. However, when estimates do match the environment, the heuristic information acquisition and restriction rules were much less useful and information should only be acquired to maintain cue balance.

If the effectiveness of the heuristic information acquisition and restriction rules are less effective when the estimates of missing information are accurate, the general usefulness of the rules would be limited. The results in Garcia-Retamero and Rieskamp (2009) and here, show that decision makers quickly and correctly adapt their estimates of missing information when provided feedback on their accuracy. If decision makers adapt their estimates, then the heuristic information acquisition and restriction methods would theoretically lose much of their effectiveness.

In human-subjects, however, these initial results show that certain distributions of missing information may be friendly or unfriendly even when decision makers have accurate estimates of missing information. If these results can be generalized, they would greatly increase the environments in which these heuristic information acquisition and restriction rules can be applied. Preliminarily, the results of this study suggest that regardless of the estimates of missing information, reducing option imbalance and increasing cue balance might be the best way to increase accuracy and decrease time required.

8.4.3 Generalizing the Results

Now that this reality gap between the mathematical and computational models and the human subjects has been found, there is much future work to be done to determine breadth of these results. This experiment examined only two-option, four-cue decision tasks with binary cue scores presented as continuous cue scores. These task attributes were chosen because they have been used frequently by others to examine decision making performance (e.g., Garcia-Retamero and Rieskamp, 2008, 2009; Bröder, 2000; Gigerenzer and Goldstein, 1996; Newell et al., 2003). Specifically, this study intended to link the studies of estimates of missing information by Garcia-Retamero and Rieskamp (2008, 2009) to the study of distributions of incomplete information in Chap. 5-7.

There are opportunities to rerun this type of experiment in ways that will better generalize these results. Only 13 distributions of incomplete information for this task size were examined. To achieve a complete analysis, future work should examine all 22 combinations of incomplete information. This would make comparisons to the computational results in Chap. 7 easier. Then, similar to the computational studies, future experiments could use larger decision tasks with more than two options and distinguish between participants who use WADD or TTB so that results can be analyzed within and between strategies.

Lastly, more focus should be taken to understand how this can be accounted for within the general linear model. Participants seemed to be stable in their estimates of missing information but some unfriendly distributions of incomplete information seemed to account almost exclusively for inaccuracy decisions.

8.5 Conclusion

The comprehensive computational study in Chap. 7 made a convincing argument that there were two types of relationships between decision making accuracy, estimates of missing information, and distributions of incomplete information. When the estimates matched the environment, only total information affected accuracy. When the estimates did not match the environment, option imbalance and cue balance significantly affected accuracy and more so than total information. The results of this human-subjects study suggests that, in reality, there might only be one classification: regardless of whether a decision maker's estimates of missing information matches the environment, their accuracy will be affected by the distribution of incomplete information. Decreasing option imbalance and increase cue balance generally increases accuracy. The disagreement between the computational and human-subjects studies highlights a potential reality gap for decision making models predicting performance of decision makers with incomplete information. Future work will be required to determine the extent of this singular classification and will have implications for the potential of heuristic information acquisition and restriction rules.

CHAPTER 9

DISCUSSION

To summarize, this dissertation has created a unified general linear model model of judgment and decision making (Chap. 3), a contextual model of decision making with incomplete information (Chap. 4), completed three computer simulation studies (Chap. 5-7), and a human-subjects study (Chap. 8) of decision making with incomplete information. This chapter considers how all of this work comes together to identify the key contributions (Table 9.1).

Table 9.1: Important contributions with further discussion in this chapter.

Section	Topic
9.1	Unifying judgment and decision making strategies in a single mathematical model.
9.2	Shedding the methodological bias toward examining only "well-studied" strategies.
9.3	Revaluing the bias of overweighting common attributes as ecologically rational.
9.4	Showing that distributions of incomplete information matter, often more than the total amount of information.
9.5	Defining heuristics for information acquisition and restriction.
9.6	Discovering a potential reality gap between what makes distributions difficult for people versus mathematical models of their strategies.

9.1 Unifying Judgment and Decision Making Strategies in a Single Mathematical Model

The general linear model provides three major benefits to the research community. First, the general linear model allows for specificity in model selection and description. It is common to refer to strategies using proper names (take-the-best, tallying, equal-weighting, etc.) which can leave ambiguity in how to appropriately represent and model the strategy, and limit intuition as to how two strategies may or may not be different. This results in further ambiguity as to whether various studies are even comparable to each other. Second, for the engineers, designers, and decision analysts, the model enables the simulation of a wide range of judgment and decision making strategies quickly and easily – and transparently – in a single equation. Rather than implementing an algorithmic tool or specialty code, it is an equation that can be looked-up and implemented. It can be used to perform sensitivity studies on the robustness of overall designs or the interaction of components across ranges of judgment and decision making strategies. Third, the simple, representative form can link together human-subjects, computational, and mathematical studies completed by various research programs such as fast-and-frugal heuristics and naturalistic decision making by linking models of experience and time pressure.

One of the most significant open questions within the domain of behavioral decision making is how to integrate the two perspectives of judgment and decision making (Katsikopoulos, 2013). A survey of developments in judgment and decision making research shows that this integration has already begun, but remains unfinished (Katsikopoulos, 2011; Martignon et al., 2008; Hogarth and Karelaia, 2005a, 2007). The general linear model from Chap. 3 attempts to achieve this goal by conceptually and mathematically integrating judgment and decision making to make the argument that, from the perspective of an information process model, they are equivalent. The general linear model can represent almost any combination of components for both judgment and decision making: cue weights, utility functions, estimates of missing information, cutoff values, thresholds, and cue directions, for m-alternatives, n-attributes and any distribution of incomplete information.

9.2 Shedding the Methodological Bias Toward Examining Only "Well-Studied" Strategies

The main theoretical tool of the fast-and-frugal heuristics program (a major influence on this dissertation) is to define "precise, computational models of heuristics that can be submitted to extensive testing and analysis" (Keller et al., 2010, p. 264). The issue is that only a few strategies have become ubiquitous and are now essentially the only "well-studied" strategies: weighted-additive, equal-weighting, take-the-best, and tallying. They are included as the main strategies numerous computational and human-subjects studies (e.g., Czerlinski et al., 1999; Martignon and Hoffrage, 2002; Katsikopoulos and Martignon, 2006; Mata et al., 2007), including the first two computational studies in this dissertation (Chap. 5 and 6).

While it is important to acknowledge and leverage the prior research on these strategies, the near-exclusive focus limits the generalizability of the research results. What if people are not using these strategies in these few, precise, computational forms? Hilbig (2010) surveyed the literature showing that there are still questions as to whether and when decision makers actually use these strategies in their standard form. Recent efforts have tried to address the generalizability question by modeling the well-studied strategies with various cue score scales (Katsikopoulos et al., 2010) or various thresholds (Luan et al., 2014). However, these studies just raise more questions about what heuristics have and have not been studied, and which decision makers they represent. Is take-the-best with continuous cue scores and non-zero thresholds instead of the original binary cue scores with thresholds at zero still take-the-best? Within the context of decision making with incomplete information the original descriptions of weighted-additive and equal-weighting make no mention of how to estimate missing information, so what is their accurate computational model? Can the current precise models accurately capture the difficulty of certain distributions of incomplete information as identified in Chap. 8?

These concerns are analagous to the debates over medical eponyms, the practice of naming diseases or conditions based on the name of a person or place (Rashid and Rashid, 2007; Mora and Bosch, 2010; Ma and Chung, 2012; Fargen and Hoh, 2014). There is growing support to reduce and even abandon medical eponyms because they lack medical accuracy and universality. Eponyms serve well as initial placeholders for describing a large set of conditions but as the understanding evolves, they often become too opaque to help patients or even medical professionals clearly understand the disease or condition.

The general linear model presented in Chap. 3 enables researchers to better realize the goal of precise computational models of decision making strategies. The model parameterizes strategies with up to six parameters: utility function, cue weights, estimates of missing information, and if necessary, cutoff values, thresholds, and cue directions. Therefore, future references to decision making strategies in research articles should also present the mathematical model, or if not possible, present the computational form. Specifying the mathematical models will also make it clear to researchers and readers how general the results of the studies are. Beyond the communication issues, the general linear model allows researchers to quickly and easily study a broad range of decision making strategies.

9.3 Revaluing the Bias of Overweighting Common Cues (Attributes) as Ecologically Rational

Marketing researchers Slovic and MacPhillamy (1974) and Kivetz and Simonson (2000) showed that decision makers in paired comparison tasks tend to overweight 'common' cues (cues that have information known for both options) and underweight 'unique' cues (cue that have information known for one option and not the other). This overweighting tendency was robust to various debiasing techniques (Slovic and MacPhillamy, 1974): (1) measuring all cue scores on the same scale, (2) prewarning subjects about the bias in favor of common cues, (3) providing feedback (after each judgment) that promotes equal weighting of dimensions, (4) providing monetary rewards for equal weighting of dimensions, and (5) providing detailed information about the distributions of the values of the three cues. In their conclusion, Slovic and MacPhillamy (1974, p. 194) argued that "further work [would] be needed to determine how to minimize these biases."

The computational and human-subjects studies in this dissertation show that focusing on common attributes, termed in this dissertation as balanced cues, is ecologically rational.

A decision making process is ecological rationality when it successfully exploits the structure of the information in natural environments (Goldstein and Gigerenzer, 2002; Todd et al., 2012). The computer simulations in Chap. 5 and Chap. 7 showed that for strategies that estimated missing information incorrectly, distributions of incomplete information with only balanced cues resulted in higher accuracy than distributions with some balanced cues and some unbalanced cues. Similar simulations showed that restricting information toward only having balanced cues could increase decision making accuracy of strategies that estimated missing information incorrectly (Chap. 6 and Chap. 7). Then the human-subjects studies in Chap. 8 broadened these results by showing that even decision makers who did estimate missing information correctly would also benefit from overweighting common attributes, achieving increased accuracy and decreased time required.

It is quite clear with this research that the empirical human tendency to overweight common attributes is likely ecologically rational, adding another entry to the list of biases which have been revalued as ecologically rational (see Table 4.1, Gigerenzer, 2004).

9.4 Showing that Distributions of Incomplete Information Matter, Often More than the Total Amount of Information

The major contribution of this dissertation is to prove that distributions of incomplete information matter. Measuring incomplete information simply as total information – as has been done in many prior studies (Martignon and Hoffrage, 2002; Garcia-Retamero and Rieskamp, 2008) – is insufficient to characterize its effect on decision making performance. This dissertation introduced two new measures of distributions of incomplete information (Chap. 4): option imbalance, the difference in number of cue scores known between the options with the most and least information; and cue balance, the number of cues which have cue scores known for each option. Through multiple computer simulations (Chap. 5 and 7) and one human-subjects study (Chap. 8), these additional measures were shown to be critical to understanding decision making performance with incomplete information: accuracy, effort, and time required. Ultimately, distributions of incomplete information should be measured via a three-variable triple (total information, option imbalance, cue balance).

Chap. 5 provided the initial computational results showing that option imbalance and cue balance mattered, but only really with respect to heuristic strategies (take-the-best and tallying) and not the analytic strategies (weighted-additive and equal-weighting). For heuristics, decreasing option imbalance and increasing cue balance, ignoring total information, would increase accuracy. Analytic strategies only had increased accuracy as total information increased.

Chap. 7 used the general linear model to computationally study 18 strategies with varying cue weights, estimates of missing information, and cutoff types. Across many real-world environments, it was shown that cue balance and option imbalance can significantly affect decision making accuracy, but only when the estimates of missing information do not match the environment. This result is explained by the fact that decreasing option imbalance and increasing cue balance reduces the strategies' reliance on the incorrect estimates of missing information. When the estimates did match the environment, the accuracy of strategies were almost entirely dependent on total information. These results generalized the findings from Chap. 5 because the heuristics had used incorrect estimates while the analytic strategies had used correct estimates.

These results raised one simple question: did these conclusions from computer simulations about the interaction between estimates, incomplete information, and performance extend to human decision makers? The human-subjects study in Chap. 8 showed some surprising, though preliminary, results. Participants successfully adapted to use the correct estimates for the environment, but their accuracy and time required were significantly affected by option imbalance and cue balance, almost as if they were not using the correct estimates.

These results suggested that some distributions of incomplete information may just simply be difficult, even when a decision maker's estimate of missing information matches the environment. Leveraging the language of Shanteau and Thomas (2000), this implies that there may be ‘friendly’ and ‘unfriendly’ distributions of incomplete information. Friendly distributions have low levels of option imbalance and high levels of cue balance, whereas unfriendly distributions have high levels of option imbalance and low levels of cue balance. These definitions leave out total information because at this point it seems to be a secondary mediator of decision making performance as compared to option imbalance and cue balance. Increased total information only increased the difference in accuracy and in time required between the friendly and unfriendly distributions in the human-subjects study (Chap. 8). The human-subjects results match the results of the prior computational studies in Chap. 5 and 7 for strategies with incorrect estimates: friendly distributions of incomplete information result in better performance (high accuracy and less time required), whereas unfriendly distributions result in worse performance (low accuracy and more time required).

9.5 Defining Heuristics for Information Acquisition and Restriction

Once the distributions of incomplete information were shown to affect decision making performance, it followed that altering the distributions of incomplete information could be a useful method of decision support. Specifically, information could be acquired (adding one cue score) or restricted (removing one cue score) to make the distribution of incomplete information more friendly. This dissertation initially defined four specific heuristic information acquisition and restriction rules that were posited to potentially increase decision making accuracy (Chap. 6): option imbalance acquisition (increase total information while decreasing option imbalance), option imbalance restriction (decrease total information while decreasing option imbalance), cue balance acquisition (increase total information while increasing cue balance), and cue balance restriction (decrease total information while keeping cue balance constant).

Importantly, these information acquisition and restriction rules do not require reliable assessments of probabilities, cue weights, and cue scores like the Bayesian analytic methods (Nelson, 2005, 2008; Meder and Nelson, 2012). By only relying on knowledge of distribution of incomplete information, the heuristic information acquisition and restriction rules should be able to overcome some of the limitations of their analytic counterparts (Katsikopoulos and Fasolo, 2006; Katsikopoulos et al., 2008). First, real-world problems can exhibit statistical dependencies which create computational issues in calculating probabilities and cue weights. Second, psychologically, an operator may not be able to provide accurate information to a decision support system regarding the probabilities and cue weights. Third, the analytic process may seem too complex and too opaque to the operator causing the operator to be reluctant to use the DSS or accept its suggestions.

Two computer simulations tested these rules and others across hundreds of environments, millions of decision tasks with incomplete information, and almost 20 different decision making strategies to determine their effectiveness (Chap. 6 and Chap. 7). Unsurprisingly, just as the effect of distributions of incomplete information depended on the accuracy of the estimates of missing information, so did the heuristic information acquisition and restriction rules.

Focusing on the more comprehensive results in Chap. 7, when the estimates did not match the environment, heuristic information acquisition and restriction rules increased accuracy when they reduced option imbalance. However, when estimates did match the environment, only heuristic information acquisition increased accuracy and only when information was acquired without affecting cue balance. In general, the acquisition rules were shown to significantly increase the accuracy of strategies: over 1% per acquisition for strategies with correct estimates of missing information and up to 4% per acquisition for strategies with incorrect estimates. The restriction rules showed more mixed results. For strategies that estimated missing information correctly, those that used prior cutoffs can restrict information to reduce cue balance with negligible effects on accuracy, while those that used relative cutoffs should not restrict information.

The results of the computational simulations suggested that the rules might be limited in real-world application because they are much less effective when the estimates of missing information were accurate. The results in Garcia-Retamero and Rieskamp (2009) and Chap. 8, have shown that decision makers quickly and correctly adapt their estimates of missing information when provided feedback on their accuracy. If decision makers adapt their estimates, then the heuristic information acquisition and restriction methods would theoretically lose much of their effectiveness.

In the human-subjects study in Chap. 8, however, these initial results show that certain distributions of missing information may be friendly or unfriendly even when decision makers have the accurate estimates of missing information. Decreasing option imbalance and increasing cue balance increased accuracy and decreased time required of participants using accurate estimates. The corresponding heuristic information acquisition rules would then seem to work for human decision makers regardless of their estimates.

Therefore, if the human-subjects results in Chap. 8 can be generalized, they would greatly increase the environments in which these heuristic information acquisition and restriction rules can be applied. Preliminarily, the results of this study suggest that regardless of the estimates of missing information, reducing option imbalance and increasing cue balance might be the best way to increase accuracy and decrease time required.

9.6 Discovering a Potential Reality Gap Between What Makes Distributions Difficult for People Versus their Computational Representations

A reality gap occurs when mathematical models struggle to reproduce the stochasticity of the real world. Though this term is pulled from the robotics community (Jakobi et al., 1995), reality gaps have consistently been the motivational hurdles in the field of decision making. Bounded rationality was posited by Simon (1955) as a counterpoint to the prevailing models of humans as a probabilistically rational 'economic man.' Tversky and Kahneman's "heuristics-and-biases" research then characterized just how unlike economic models humans are, finding that people do not necessarily judge uncertainty according to the rules of probability and statistics because they are influenced by framing and other cognitive biases (e.g., Tversky and Kahneman, 1974; Kahneman et al., 1982; Kahneman and Tversky, 1984). Since then, two additional decision making paradigms have been developed to address what other researchers believed were the heuristics-and-biases' limited ability to explain how people used the bounds on their rationality in successful ways. The naturalistic decision making framework, showed that people faced with time pressure, uncertainty, ill-defined goals, and other complexities in familiar and meaningful environments, can perform very well by using various strategies such as matching strategies to situations, experience-based knowledge modeling, and implementable-based strategy selection (e.g., Klein et al., 1988; Orasanu and Connolly, 1993; Klein and Calderwood, 1996; Klein, 2008; Lipshitz et al., 2001). Furthermore, the fast and frugal heuristics framework has shown that people use the bounds on rationality of simplicity, speed, and frugality as a mechanism for simple, robust, and accurate strategies that adapt to the environment and ecology (e.g., Gigerenzer and Goldstein, 1996; Gigerenzer et al., 1999; Gigerenzer and Gaissmaier, 2011; Todd et al., 2012).

Following in the succession of reality gaps between mathematical models, computational simulations, and human-subjects studies, the disagreement in results between the computer simulation studies (Chap. 5 and Chap. 7) and the human-subjects experiment (Chap. 8) is simply another one. In a comprehensive study of the interaction of estimates of missing information, cue weights, and cutoff types, Chap. 7 showed that correct estimates of missing information make decision making accuracy only a function of total information. For strategies with incorrect estimates of missing information, their accuracy was more influenced by option imbalance and cue balance. In contrast, in the human-subjects experiment in Chap. 8, decision makers were shown to be using correct estimates of missing information but their accuracy was typical of decision makers with incorrect estimates – having increased accuracy as option imbalance decreased and cue balance increased.

This was a surprising result. The mathematical model of decision making used in this dissertation (Chap. 3), and particularly in Chap. 7, accounts for a very broad range of decision making strategies that have been shown to generally match participant's strategies. Additionally, the human-subjects study by Garcia-Retamero and Rieskamp (2009) explicitly studying estimates of missing information, did not report accuracy or time required as a function of estimates of missing information and thus did not identify the difference between their computer studies (Garcia-Retamero and Rieskamp, 2008) and their human-subjects studies (Garcia-Retamero and Rieskamp, 2009). Future work will be required to address this reality gap.

Narrowing the reality gap will likely require a mix of qualitative and mathematical methods. It may be that the participants in the human-subjects study had their estimates of missing information directly affected by the distribution of incomplete information. For example, when there is a imbalanced distribution of incomplete information, participants may be second-guessing their accurate estimates of missing information resulting in the use of inaccurate estimates (i.e., their estimates were not stable). To measure this, participants could be asked during each decision task to input the estimated value of missing information or participants could report their overall estimate of missing information after the experiment (Garcia-Retamero and Rieskamp, 2009). Once the relationship between estimates of missing information and the distribution of incomplete information is found, then a mathematical alteration to the general linear model could be introduced.

CHAPTER 10

SUMMARY AND FUTURE RESEARCH DIRECTIONS

10.1 Summary

Decision makers are continuously required to make choices in environments with incomplete information. This dissertation sought to understand and, ultimately, support the wide range of decision making strategies used in environments with incomplete information.

Prior to this dissertation, incomplete information was measured almost exclusively by the total amount of information (e.g., Martignon and Hoffrage, 2002; Garcia-Retamero and Rieskamp, 2008). This dissertation conclusively showed that total information alone is insufficient for understanding and supporting decision making accuracy and effort with incomplete information. How the information was distributed was often a more important determinant of decision making performance.

Two new measures of the distribution of incomplete information were introduced (option imbalance and cue balance, Chap. 4) and tested alongside total information across three computer simulations of 18 variations of decision making strategies within hundreds of environments and millions of decision tasks with incomplete information (Chap. 5-7). The simulations were powered by a new general linear model of judgment and decision making which is capable of efficiently and transparently modeling a wide range of strategies beyond the traditional set in the literature (Chap. 3). Of the many potential mediators of the relationship between the distributions of incomplete information and performance, only the strategies' estimates of missing information were shown to be significant in the computational studies. For strategies with accurate estimates, only total information determined accuracy, while for strategies with inaccurate estimates, accuracy was maximized when option imbalance was low and cue balance was high.

Based on these results, heuristic rules were defined to guide information acquisition and restriction based on altering the distribution of incomplete information to decrease option imbalance and increase cue balance (Chap. 6-7), from unfriendly distributions to friendly distributions. These rules do not require probability and cue weight information, unlike the prominent Bayesian information acquisition methods (Nelson, 2005; Nelson et al., 2010; Meder and Nelson, 2012), and thus allow the heuristic rules to overcome some major environmental limitations of the Bayesian methods (Katsikopoulos and Fasolo, 2006). Similar to the studies of incomplete information, the heuristic rules were simulated and shown to increase accuracy for strategies with inaccurate estimates of missing information, and did not significantly affect strategies with accurate estimates (Chap. 6-7).

The simulation results were partially contradicted in a human-subjects study (Chap. 8) which showed that human decision makers with accurate estimates were strongly affected by option imbalance and cue balance – suggesting that some distributions that might simply be difficult or unfriendly regardless of the estimates of missing information. This reality gap between the human decision makers and their computational representations is a new open question in the decision making literature.

If the human-subjects study results can be generalized through repeated studies and different task environments, there are significant implications on decision support. The goal of decision support systems would then be to modify the presentation of information from unfriendly to friendly distributions. Within most unfriendly distributions, there is a friendly distribution with less total information. Similarly, when starting from an unfriendly distribution, there is a clear path toward a friendly distribution through information acquisition.

10.2 Future Research Directions

There are numerous directions for future research based on this dissertation:

10.2.1 Expanding the General Linear Model

Create an R-package of the general linear model

Creating an computational package of the general linear model in the open-source R-statistical software will increase is accessibility within the research community. The R-package *FFTrees*[1] is popular among the heuristics community for constructing and testing fast-and-frugal trees but the general linear model developed in this dissertation should be able do everything that package is capable of, and much more. Given that the mathematical model and its matrix-representation have already been defined, this should be a fairly straightforward task.

Incorporate measures of time and effort

The primary limitation of the model is caused by the same mechanism that makes it fast and transparent: each option is evaluated independently, even in decision making tasks. Therefore, in decision tasks with multiple options it is not clear how to incorporate the measures of effort or time into the model since many strategies – particularly heuristics – will not process all the information provided in each option.

Develop a regression-fitting version of the general linear model

The variables of the current general linear model are actively selected by the user constructing the model based on knowledge of the decision maker. Future work should focus on enabling the major variables in the model to be statistically fit to a decision maker without explicit knowledge of the decision maker's process, but with only the information about

[1] *FFTrees* is available at https://cran.rstudio.com/web/packages/FFTrees/index.html.

the environment, decision task, and decisions. This capability would enable the model to be more easily incorporated into machine learning algorithms and experiments in which the task cannot sufficiently constrained to elicit the major variables.

Incorporate variables to account for dynamic strategy change or selection

The current general linear model is a static representation of decision making, in conflict with our knowledge that participants adapt their decision making strategies as a function of context variables such as time pressure, environment, etc. (Hammond, 1988; Hollnagel, 1993; Marewski and Link, 2014). Clarifying how the model can adapt and represent dynamic decision making would be an important contribution but it is not yet clear how to incorporate that into the model. Potentially the mathematics of strategy selection models (e.g., Rieskamp, 2006a) or phase-change financial models (e.g., Howison, 1995) can be used to model specific archetypes of decision makers and their transitions between strategies.

Modeling naturalistic decision making with the general linear model

Keller et al. (2010) argued for the development of naturalistic heuristics, an integration of the naturalistic decision making program with the fast-and-frugal heuristics program. Since the general linear model was built upon the framework of the fast-and-frugal heuristics, there is a potential to use the model to assist in mathematically representing naturalistic decision making processes, especially with their focus on time pressure and experience. Canellas and Feigh (2016b) already started to use the general linear model to integrate the two programs formally, showing that the naturalistic decision making quick test process could be modeled as a fast-and-frugal tree equivalent to the tallying heuristic. That successful attempt, however, should only be viewed as evidence that much more integration is possible.

10.2.2 Judgment and Decision Making with Incomplete Information

Expand the capabilities of the general linear model for judgment

The general linear model has already been shown in Chap. 3 to be linear classifier capable of calculating a score, or criterion, based on many components (cue weights, utility functions, incomplete information, estimates of missing information, cue directions, thresholds, and cutoff values). Specifically, it was shown that the general linear model can represent classification trees with up to 3-edges per node. Future conceptual work will be required to discuss implications of the model for other theories and models of judgment, e.g., the Lens Model. Whereas the Lens Model characterizes various types of agreement between two judges' linear models (human and environment, human and automation, etc.), given that the general linear model has many more subcomponents, agreements could be measured with respect to additional parameters such as cutoffs, thresholds, estimates of missing information, etc. Furthermore, given the general linear model's formulation of fast-and-frugal trees (essentially non-linear, heuristic judgments), agreements could be measured between a fast-and-frugal tree and the human participants, especially when the Lens Model suggests that there may be a non-linear relationship between the cues and criterion.

Construct a cascading judgment-and-decision making process model

The general linear model represents judgment and decision making as the same information process. Therefore, the two domains are represented using the same parameters, inputs, and outputs, just with different meanings. Judgment processes one option and outputs a classification of that one option. Decision making processes two or more options and outputs the option with the highest criterion. Since the two processes share the same variables, they should be able to be linked together in a meaningful way in order to model cascading decision events: 1) many judgment processes are used to determine the cue values for multiple options and cues, 2) the decision process chooses the 'best' option, 3) based on the decision, the environment is acted upon, resulting in a new set of judgments to be made, and the process is repeated again and again.

Examine fast-and-frugal trees with incomplete information

Fast-and-frugal trees (FFTs) have become a popularly studied tool for decision support (e.g., Jenny et al., 2013; Luan et al., 2011). They have been shown to be quick to use, easy to remember, and accurate across many domains: medical (Green and Mehr, 1997; Fischer et al., 2002; Katsikopoulos et al., 2008), military (Keller and Katsikopoulos, 2016), and financial (Aikman et al., 2014). Despite the increased adoption of FFTs there has not been a clear accounting of incomplete information in the standard FFT form. As a result there has not been a study of how incomplete information affects FFT performance despite the prevalence of incomplete information in medial, military, and financial decision making. The results of this dissertation should be able to advise the development of FFTs which can support classification with incomplete information. The general linear model can be used to easily and rapidly test FFTs and examine how their performance with incomplete information relates to the results of this dissertation on the interaction between estimates of missing information, incomplete information, and decision making performance.

Complete a more comprehensive human-subjects study of decision making with incomplete information

The reality gap between the computer simulations in Chap. 7 and the human-subjects study in Chap. 8, requires more investigation. The human-subjects study in Chap. 8 should be repeated with all 22 instead of 13 distributions of incomplete information. That would make the human-subjects study results more easily comparable to the computational studies. If the same results are found, it would also be a second independent set of participants indicating that some distributions of incomplete information are difficult regardless of whether the estimates of missing information are accurate.

10.2.3 Heuristic Information Acquisition and Restriction Rules

Compare the heuristic information acquisition and restriction rules to Bayesian optimal experimental design

The heuristic rules were tested across many environments and strategies in this dissertation (Chap. 6 and Chap. 7) but their baseline comparison was whether accuracy increased by following the rules as opposed to not following the rules. In the interest of competitive testing, the performance of the rules should be compared to the Bayesian optimal experimental methods of information acquisition (Nelson, 2005, 2008; Meder and Nelson, 2012). The Bayesian methods require reliable assessments of probabilities, cue weights, and cue scores, while the heuristic methods do not, making the comparison an interesting general test of heuristic versus analytic methods.

Complete a human-subject study of heuristic information acquisition and restriction

The human-subjects study of decision making with incomplete information showed that some distributions were difficult, or 'unfriendly,' regardless of the accuracy of the participants' estimates of missing information (Chap. 8). This shows that the heuristic information acquisition and restriction rules have the potential to improve accuracy of a wide range of decision makers. A human-subjects study is needed to test the extent to which the rules apply to human, not computational, decision makers. The human-subjects study could also acquire information via the heuristic and the Bayesian methods to determine if participants have difference in performance and a preference for one method or the other.

Develop and test a decision support tool implementing the heuristic information acquisition and restriction rules

Implementation of these heuristic information acquisition and restriction rules is the ultimate goal of this project. However, it is not clear at this point how the heuristic information acquisition and restriction rules manifest themselves in real decision support tools. There are multiple options available to system and interface designers such as modifying content within interfaces by changing the saliency of information or by providing some information automatically while requiring active interaction to reveal other information (Feigh et al., 2012).

Appendices

APPENDIX A

DISSERTATION APPENDICES

A.1 General Linear Model of Judgment and Decision Making

A.1.1 Extending Non-Compensatory Weights Beyond Strictly Binary Cues

Martignon and Hoffrage ("Proof of Theorem 3," 2002, pp. 65–66) provided the original proof that TTB decisions can be represented as a linear model with non-compensatory weights. The cue weights had to be selected such that any cue's weight was larger than the sum of all the cue weights of lower ranked cues (Eq. A.1). By considering *strictly* binary cues (e.g. a binary utility function producing cue scores of either 0 or 1), they solved Eq. A.1 using a geometric series with $1/2$ as the common ratio, as shown in Eq. A.2.

$$w_j > \sum_{k>j}^{\infty} w_k \tag{A.1}$$

$$w_j = \left(\frac{1}{2}\right)^{j-1} = 2^{1-j} \tag{A.2}$$

Here I show that a cue weight generating function can be created for a binary utility function with three states for cue scores: $\{0, 0.5, 1\}$, with the 0.5 representing situations where it is unclear whether the cue value should be represented as a 0 or 1. This is particularly useful for tasks with incomplete information where a decision maker may want to estimate missing cue values somewhere between 0 and 1 (Garcia-Retamero and Rieskamp, 2008, 2009).

First, to construct such a set of non-compensatory cue weights, three preferences between two alternatives (A and B) with incomplete information must be satified as shown in Conditions 1–3. The arguments of A and B refer to the cue scores.

$$\text{Condition 1:} \quad A(1,0,...,0), \quad \text{preferred to } B(0,1,...,1) \tag{A.3}$$

$$\text{Condition 2:} \quad A(1,0,...,0), \quad \text{preferred to } B(0.5,1,...,1) \tag{A.4}$$

$$\text{Condition 3:} \quad A(0.5,0,...,0), \quad \text{preferred to } B(0,1,...,1) \tag{A.5}$$

Using the weighted utility function, the three conditions can be restated as inequalities that must be satisfied:

$$\text{Condition 1:} \quad w_1 \quad > \sum_{k=2}^{n} w_j \tag{A.6}$$

$$\text{Condition 2:} \quad w_1 \quad > w_1 \cdot 0.5 + \sum_{k=2}^{n} w_j \tag{A.7}$$

$$\text{Condition 3:} \quad w_1 \cdot 0.5 \quad > \sum_{k=2}^{n} w_j \tag{A.8}$$

Condition 1 can be satisfied by the use of geometric series of at least common ratio of $1/2$ (Table A.1, assuming that there are less than infinite cues). To determine which common ratios can account for the three-state binary cue scores, Conditions 2 and 3 (Eqs. A.7 and A.8) are combined into the following equation with $w_1 = 1$ as shown in Table A.1:

Table A.1: Potential geometric series for cue weights (w_j) for various common ratios (r): $w_j = r^{j-1}$.

r	w_1	$\sum_{j=2}^{n=\infty} w_j$
1	1	∞
$1/2$	1	1
$1/3$	1	$1/2$
$1/4$	1	$1/3$

$$1 > 0.5 > \sum_{j=2}^{n} w_j \tag{A.9}$$

Therefore, of the standard common ratios, $r = 1/4$ is the largest that satisfies Conditions 1, 2, and 3. The final form for the non-compensatory cue weight generating function that can account for cue scores of $\{0, 0.5, 1\}$ is presented in final form in Eq. A.10.

$$w_j = \sum_j^n \left(\frac{1}{4}\right)^{j-1} = \sum_j^n 4^{1-j} \tag{A.10}$$

A.1.2 Cue weights to categories

Tables A.2 and A.3 show the specific criterion values for two-state and three-state cue scores, respectively, for non-compensatory and equal cue weights.

Table A.2: Criterion values for all combinations of 2-state binary cue scores for three cues. Non-compensatory cue weights generate 8 unique criterion values while equal cue weights generate 4 unique criterion values.

$\mathbf{a_1^s}$	$\mathbf{a_2^s}$	$\mathbf{a_3^s}$	**Non-Compensatory** $\mathbf{C : w_j = 4^{1-j}}$	**Equal** $\mathbf{C : w_j = 1}$
0	0	0	0	0
0	0	1	0.0625	1
0	1	0	0.25	1
0	1	1	0.3125	2
1	0	0	1	1
1	0	1	1.0625	2
1	1	0	1.25	2
1	1	1	1.3125	3

Table A.3: Criterion values for all combinations of 3-state binary cue scores for three cues. Non-compensatory cue weights generate 27 unique criterion values while equal cue weights generate 7 unique criterion values.

$\mathbf{a_1^s}$	$\mathbf{a_2^s}$	$\mathbf{a_3^s}$	**Non-Compensatory** $\mathbf{C : w_j = 4^{1-j}}$	**Equal** $\mathbf{C : w_j = 1}$
1	1	1	1.3125	3
0	1	1	0.3125	2
0.5	1	1	0.8125	2.5
1	0	1	1.0625	2
0	0	1	0.0625	1
0.5	0	1	0.5625	1.5
1	0.5	1	1.1875	2.5
0	0.5	1	0.1875	1.5
0.5	0.5	1	0.6875	2
1	1	0	1.25	2
0	1	0	0.25	1
0.5	1	0	0.75	1.5
1	0	0	1	1
0	0	0	0	0
0.5	0	0	0.5	0.5
1	0.5	0	1.125	1.5
0	0.5	0	0.125	0.5
0.5	0.5	0	0.625	1
1	1	0.5	1.281 25	2.5
0	1	0.5	0.281 25	1.5
0.5	1	0.5	0.781 25	2
1	0	0.5	1.031 25	1.5
0	0	0.5	0.031 25	0.5
0.5	0	0.5	0.531 25	1
1	0.5	0.5	1.156 25	2
0	0.5	0.5	0.156 25	1
0.5	0.5	0.5	0.656 25	1.5

A.1.3 Prior Formalization of Fast-and-Frugal Trees

Martignon et al. (2008) formalized FFTs with two parameters: cue profiles (x) and splitting profiles (S). Cue profiles are equivalent to a single alternative's binary cue scores used in the general linear model with one notable exception: the values in the cue profiles must be coded so that a cue score of 1 indicates that the alternative (according to that cue) should be assigned to Category 1 and a cue score of 0 indicates that the alternative (according to that cue) should be assigned to Category 0. A splitting profile is a cue profile that describes the FFT structure. An FFT assigns an alternative to Category 1 if and only if the cue profile is lexicographically larger than the splitting profile. In simpler terms, the elements of the splitting profile indicate the following: if $S_j = 0$, the FFT has an exit "yes" to Category 1, and if $S_j = 1$, the FFT has an exit "no" to Category 0.

While requiring a relationship between cue directions and categorizations allows for a compact notation – a single vector of 0's and 1's (the splitting profile) can specify an entire FFT – it restricts the ability of modelers to specify their cue directions independently of the categorizations and restricts the number of categories to 2 (though that is the standard number in signal detection theory which has been used to formalize FFTs, Luan et al., 2011). Therefore, for the general linear model presented in this thesis, there is no required relationship between cue directions and categorizations; only together with the cue directions D can the splitting profile specify an FFT. The cue directions specify whether the exit occurs because the cue value is less than the cutoff (-1) or more than the cutoff (1). The splitting profile specifies the categorization of that exit. See Sec. 3.2.4 for an example of FFTs using the general linear model.

A.1.4 Dataset Description for Fast-and-Frugal Trees Example

The dataset used for the two examples in Sec. 3.2.4, was accessed via the Adaptive Toolbox Online. The original dataset is available at the UCI Machine Learning Database (Lichman,

2013) which has a complete text description of the dataset[1] as shown below. The Adaptive Toolbox Online used "Contraceptive Method Used" (`cm_bin`) as the criterion and converted it to binary: 0, no-use, and 1, long-term or short-term use.

1. Title: Contraceptive Method Choice
2. Sources:
   (a) Origin: This dataset is a subset of the 1987 National Indonesia Contraceptive Prevalence Survey
   (b) Creator: Tjen-Sien Lim (limt@stat.wisc.edu)
   (c) Donor: Tjen-Sien Lim (limt@stat.wisc.edu)
   (d) Date: June 7, 1997
3. Past Usage: Lim, T.-S., Loh, W.-Y. & Shih, Y.-S. (1999). A Comparison of Prediction Accuracy, Complexity, and Training Time of Thirty-three Old and New Classification Algorithms. Machine Learning. Forthcoming.

   (ftp://ftp.stat.wisc.edu/pub/loh/treeprogs/quest1.7/mach1317.pdf)

   or (http://www.stat.wisc.edu/ limt/mach1317.pdf)
4. Relevant Information: This dataset is a subset of the 1987 National Indonesia Contraceptive Prevalence Survey. The samples are married women who were either not pregnant or do not know if they were at the time of interview. The problem is to predict the current contraceptive method choice (no use, long-term methods, or short-term methods) of a woman based on her demographic and socio-economic characteristics.
5. Numer of Instances: 1473
6. Number of Attributes (Cues): 10 (including the class attribute)

[1] https://archive.ics.uci.edu/ml/datasets/Contraceptive+Method+Choice

Table A.4: Environmental parameters for the 5-cue decision environments.

Dataset	Full Information Accuracy					Pred.	Redun.	Var.
	WADD	EW	Tallying	Take Two	TTB			
Biodiversity	0.84	0.83	0.83	0.84	0.86	0.53	0.26	0.47
Boys	0.78	0.78	0.78	0.78	0.82	0.53	0.58	0.15
Car	0.75	0.78	0.78	0.75	0.72	0.57	0.21	0.16
Cities	0.81	0.83	0.83	0.81	0.81	0.86	0.22	0.23
Dropout	0.68	0.69	0.69	0.68	0.69	0.20	0.75	0.03
Fat	0.58	0.58	0.58	0.58	0.59	0.10	0.34	0.09
Fuel	0.79	0.83	0.83	0.79	0.79	0.51	0.25	0.16
Girls	0.74	0.76	0.76	0.74	0.76	0.44	0.51	0.18
Homeless	0.69	0.69	0.69	0.69	0.70	0.32	0.17	0.37
House	0.91	0.92	0.92	0.91	0.91	0.49	0.57	0.17
Oxygen	0.92	0.92	0.92	0.92	0.92	0.22	0.89	0.12
Pollution	0.77	0.79	0.79	0.78	0.77	0.48	0.34	0.17
Professor	0.81	0.81	0.81	0.81	0.82	0.70	0.30	0.43
Rainfall	0.69	0.73	0.73	0.69	0.68	0.33	0.12	0.13
Sleep	0.81	0.81	0.81	0.81	0.84	0.42	0.52	0.22

7. Attribute (Cue) Information[2]:

    2. `edu`, Wife's education (categorial) 1=low, 2, 3, 4=high

    4. `n_child`, Number of children ever born (numerical)

    6. `work`, Wife's now working? (binary) 0=Yes, 1=No

    10. `cm_bin`, Contraceptive method used (categorical) 1=No-use, 2=Long-term, 3=Short-term

A.2 Computer Simulation Studies

A.2.1 Study 2: Heuristic Information Acquisition and Restriction Rules for Decision Support

Environmental Parameters for Each Dataset

[2]Only the three cues and the criterion used in the example are described here.

Table A.5: Environmental parameters for the 4-cue decision environments.

Dataset	**Full Information Accuracy**					**Pred.**	**Redun.**	**Var.**
	WADD	**EW**	**Tallying**	**Take Two**	**TTB**			
Biodiversity	0.87	0.86	0.86	0.87	0.88	0.52	0.35	0.44
Boys	0.79	0.80	0.80	0.79	0.82	0.51	0.59	0.14
Car	0.73	0.77	0.77	0.73	0.72	0.53	0.25	0.14
Cities	0.84	0.86	0.86	0.84	0.83	0.86	0.19	0.22
Dropout	0.69	0.69	0.69	0.69	0.69	0.19	0.78	0.02
Fat	0.59	0.59	0.59	0.59	0.60	0.10	0.28	0.09
Fuel	0.79	0.81	0.81	0.79	0.79	0.49	0.35	0.05
Girls	0.78	0.78	0.78	0.78	0.80	0.44	0.65	0.12
Homeless	0.72	0.74	0.74	0.72	0.71	0.31	0.16	0.32
House	0.93	0.94	0.94	0.93	0.94	0.45	0.67	0.12
Oxygen	1.00	1.00	1.00	1.00	1.00	0.20	1.00	0.00
Pollution	0.78	0.78	0.78	0.78	0.78	0.47	0.44	0.10
Professor	0.86	0.87	0.87	0.86	0.86	0.69	0.37	0.30
Rainfall	0.70	0.74	0.74	0.70	0.68	0.32	0.11	0.06
Sleep	0.84	0.84	0.84	0.84	0.84	0.42	0.55	0.17

Table A.6: Environmental parameters for the 3-cue decision environments.

Dataset	**Full Information Accuracy**					**Pred.**	**Redun.**	**Var.**
	WADD	**EW**	**Tallying**	**Take Two**	**TTB**			
Biodiversity	0.90	0.90	0.90	0.90	0.90	0.52	0.49	0.28
Boys	0.81	0.82	0.82	0.81	0.82	0.50	0.54	0.04
Car	0.72	0.75	0.75	0.72	0.72	0.40	0.36	0.13
Cities	0.90	0.90	0.90	0.90	0.90	0.85	0.25	0.13
Dropout	0.69	0.70	0.70	0.69	0.69	0.19	0.72	0.01
Fat	0.60	0.60	0.60	0.60	0.60	0.09	0.27	0.07
Fuel	0.81	0.81	0.81	0.81	0.81	0.46	0.39	0.03
Girls	0.81	0.82	0.82	0.81	0.84	0.43	0.69	0.06
Homeless	0.73	0.77	0.77	0.73	0.73	0.30	0.24	0.32
House	0.96	0.96	0.96	0.96	0.97	0.41	0.82	0.06
Oxygen	1.00	1.00	1.00	1.00	1.00	0.20	1.00	0.00
Pollution	0.80	0.80	0.80	0.80	0.79	0.41	0.74	0.04
Professor	0.89	0.91	0.91	0.89	0.90	0.69	0.53	0.10
Rainfall	0.69	0.73	0.73	0.69	0.68	0.24	0.07	0.04
Sleep	0.88	0.87	0.87	0.88	0.88	0.39	0.66	0.08

A.2.2 Study 3: Determinants of Decision Making with Incomplete Information

Strategy Equations

$$C_{PPE} = \sum_{j=1}^{n} H\Big[d_j\Big(\max a_j^v + (a_{i,j}^v - \max a_j^v) z_{i,j}\Big) - (d_j \cdot \tilde{a}_j^v)\Big] \quad \text{(A.11)}$$

$$C_{PME} = \sum_{j=1}^{n} H\Big[d_j\Big(\tilde{a}_j^v + (a_{i,j}^v - \tilde{a}_j^v) z_{i,j}\Big) - (d_j \cdot \tilde{a}_j^v)\Big] \quad \text{(A.12)}$$

$$C_{PNE} = \sum_{j=1}^{n} H\Big[d_j\Big(\min a_j^v + (a_{i,j}^v - \min a_j^v) z_{i,j}\Big) - (d_j \cdot \tilde{a}_j^v)\Big] \quad \text{(A.13)}$$

$$C_{RPE} = \sum_{j=1}^{n} H\Big[d_j\Big(\max a_j^v + (a_{i,j}^v - \max a_j^v) z_{i,j}\Big) - (d_j \cdot \texttt{maxmin} a_j)\Big] \quad \text{(A.14)}$$

$$C_{RME} = \sum_{j=1}^{n} H\Big[d_j\Big(\tilde{a}_j^v + (a_{i,j}^v - \tilde{a}_j^v) z_{i,j}\Big) - (d_j \cdot \texttt{maxmin} a_j)\Big] \quad \text{(A.15)}$$

$$C_{RNE} = \sum_{j=1}^{n} H\Big[d_j\Big(\min a_j^v + (a_{i,j}^v - \min a_j^v) z_{i,j}\Big) - (d_j \cdot \texttt{maxmin} a_j)\Big] \quad \text{(A.16)}$$

(A.17)

$$C_{PPN} = \sum_{j=1}^{n} 4^{1-j} \cdot H\Big[d_j\Big(\max a_j^v + (a_{i,j}^v - \max a_j^v)z_{i,j}\Big) - (d_j \cdot \tilde{a}_j^v)\Big] \quad (A.18)$$

$$C_{PMN} = \sum_{j=1}^{n} 4^{1-j} \cdot H\Big[d_j\Big(\tilde{a}_j^v + (a_{i,j}^v - \tilde{a}_j^v)z_{i,j}\Big) - (d_j \cdot \tilde{a}_j^v)\Big] \quad (A.19)$$

$$C_{PNN} = \sum_{j=1}^{n} 4^{1-j} \cdot H\Big[d_j\Big(\min a_j^v + (a_{i,j}^v - \min a_j^v)z_{i,j}\Big) - (d_j \cdot \tilde{a}_j^v)\Big] \quad (A.20)$$

$$C_{RPN} = \sum_{j=1}^{n} 4^{1-j} \cdot H\Big[d_j\Big(\max a_j^v + (a_{i,j}^v - \max a_j^v)z_{i,j}\Big) - (d_j \cdot \texttt{maxmin} a_j)\Big] \quad (A.21)$$

$$C_{RMN} = \sum_{j=1}^{n} 4^{1-j} \cdot H\Big[d_j\Big(\tilde{a}_j^v + (a_{i,j}^v - \tilde{a}_j^v)z_{i,j}\Big) - (d_j \cdot \texttt{maxmin} a_j)\Big] \quad (A.22)$$

$$C_{RNN} = \sum_{j=1}^{n} 4^{1-j} \cdot H\Big[d_j\Big(\min a_j^v + (a_{i,j}^v - \min a_j^v)z_{i,j}\Big) - (d_j \cdot \texttt{maxmin} a_j)\Big] \quad (A.23)$$

(A.24)

$$C_{PPC} = \sum_{j=1}^{n} v \cdot H\Big[d_j\Big(\max a_j^v + (a_{i,j}^v - \max a_j^v)z_{i,j}\Big) - (d_j \cdot \tilde{a}_j^v)\Big] \quad (A.25)$$

$$C_{PMC} = \sum_{j=1}^{n} v \cdot H\Big[d_j\Big(\tilde{a}_j^v + (a_{i,j}^v - \tilde{a}_j^v)z_{i,j}\Big) - (d_j \cdot \tilde{a}_j^v)\Big] \quad (A.26)$$

$$C_{PNC} = \sum_{j=1}^{n} v \cdot H\Big[d_j\Big(\min a_j^v + (a_{i,j}^v - \min a_j^v)z_{i,j}\Big) - (d_j \cdot \tilde{a}_j^v)\Big] \quad (A.27)$$

$$C_{RPC} = \sum_{j=1}^{n} v \cdot H\Big[d_j\Big(\max a_j^v + (a_{i,j}^v - \max a_j^v)z_{i,j}\Big) - (d_j \cdot \texttt{maxmin} a_j)\Big] \quad (A.28)$$

$$C_{RMC} = \sum_{j=1}^{n} v \cdot H\Big[d_j\Big(\tilde{a}_j^v + (a_{i,j}^v - \tilde{a}_j^v)z_{i,j}\Big) - (d_j \cdot \texttt{maxmin} a_j)\Big] \quad (A.29)$$

$$C_{RNC} = \sum_{j=1}^{n} v \cdot H\Big[d_j\Big(\min a_j^v + (a_{i,j}^v - \min a_j^v)z_{i,j}\Big) - (d_j \cdot \texttt{maxmin} a_j)\Big] \quad (A.30)$$

(A.31)

Tradeoff Figures

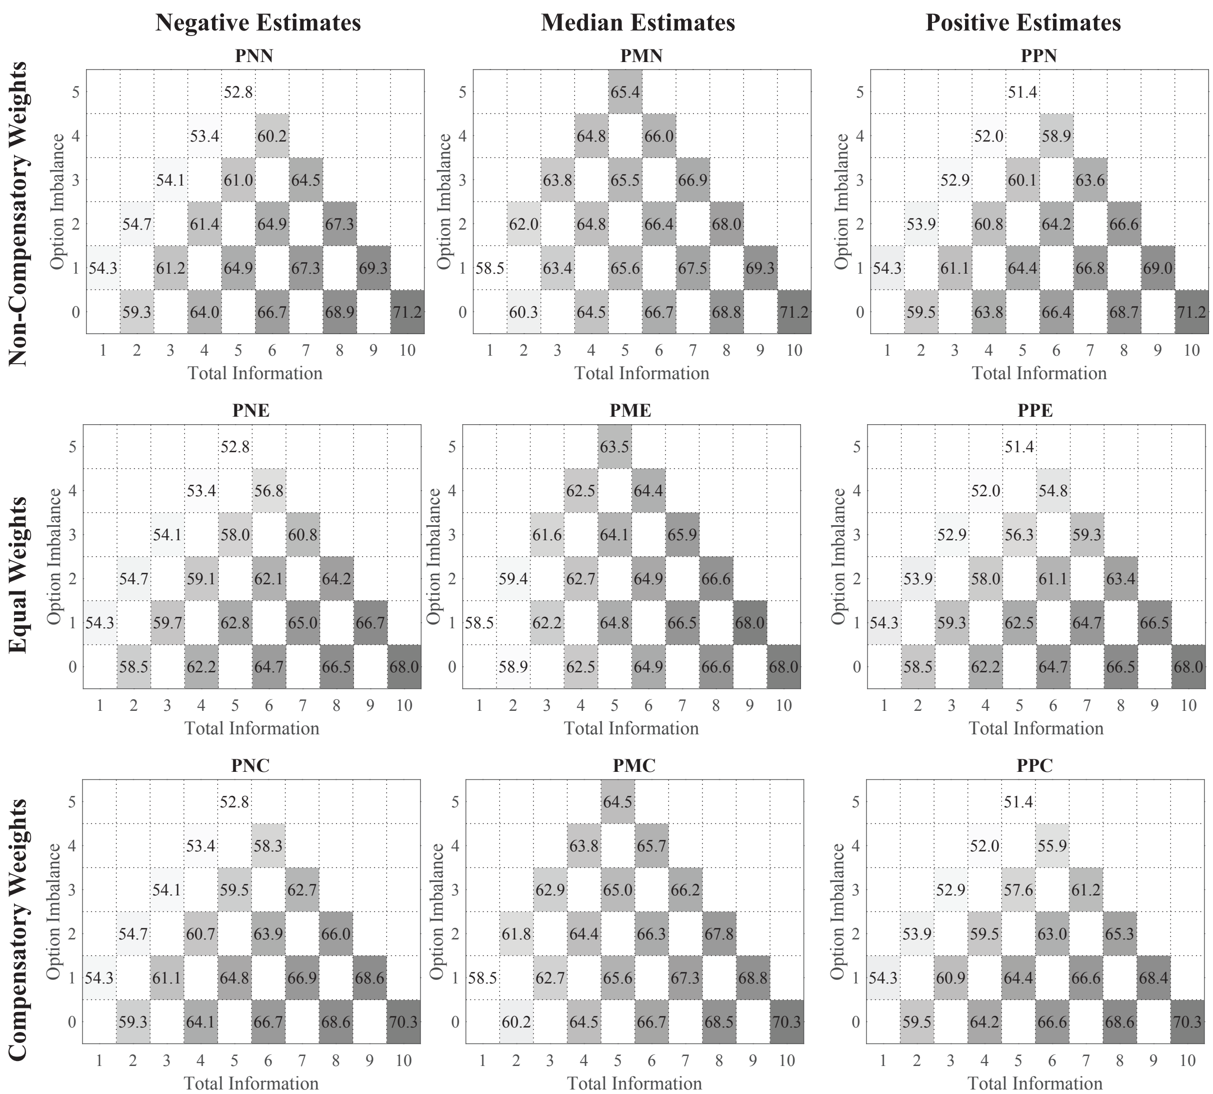

Figure A.1: Two-way interaction between total information and option imbalance for strategies with prior cutoff values and each combination of estimates and weights.

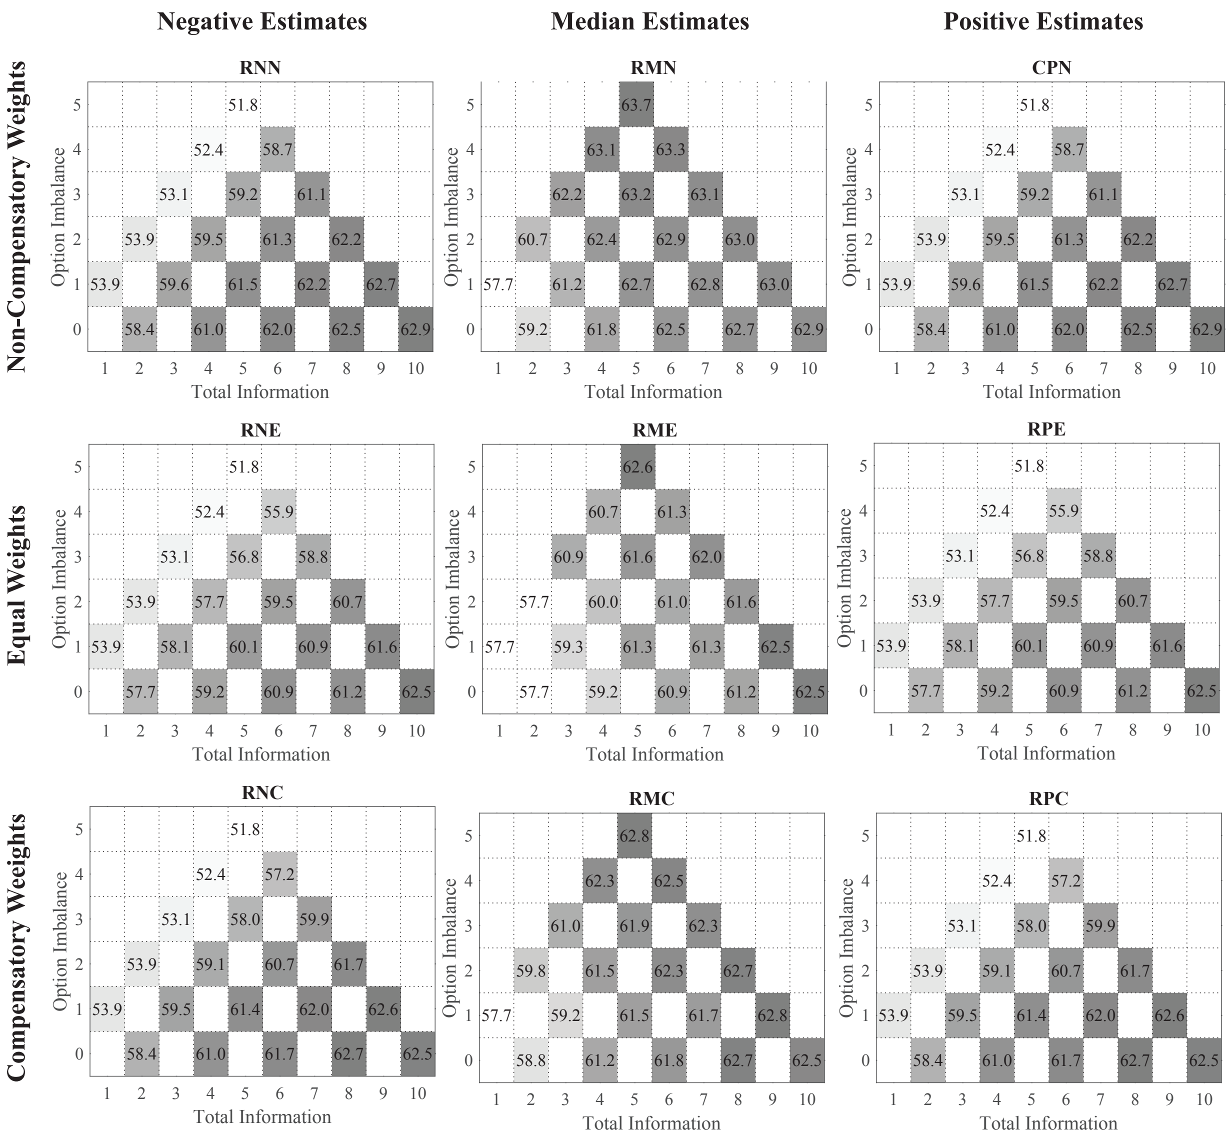

Figure A.2: Two-way interaction between total information and option imbalance for strategies with relative cutoff values and each combination of estimates and weights.

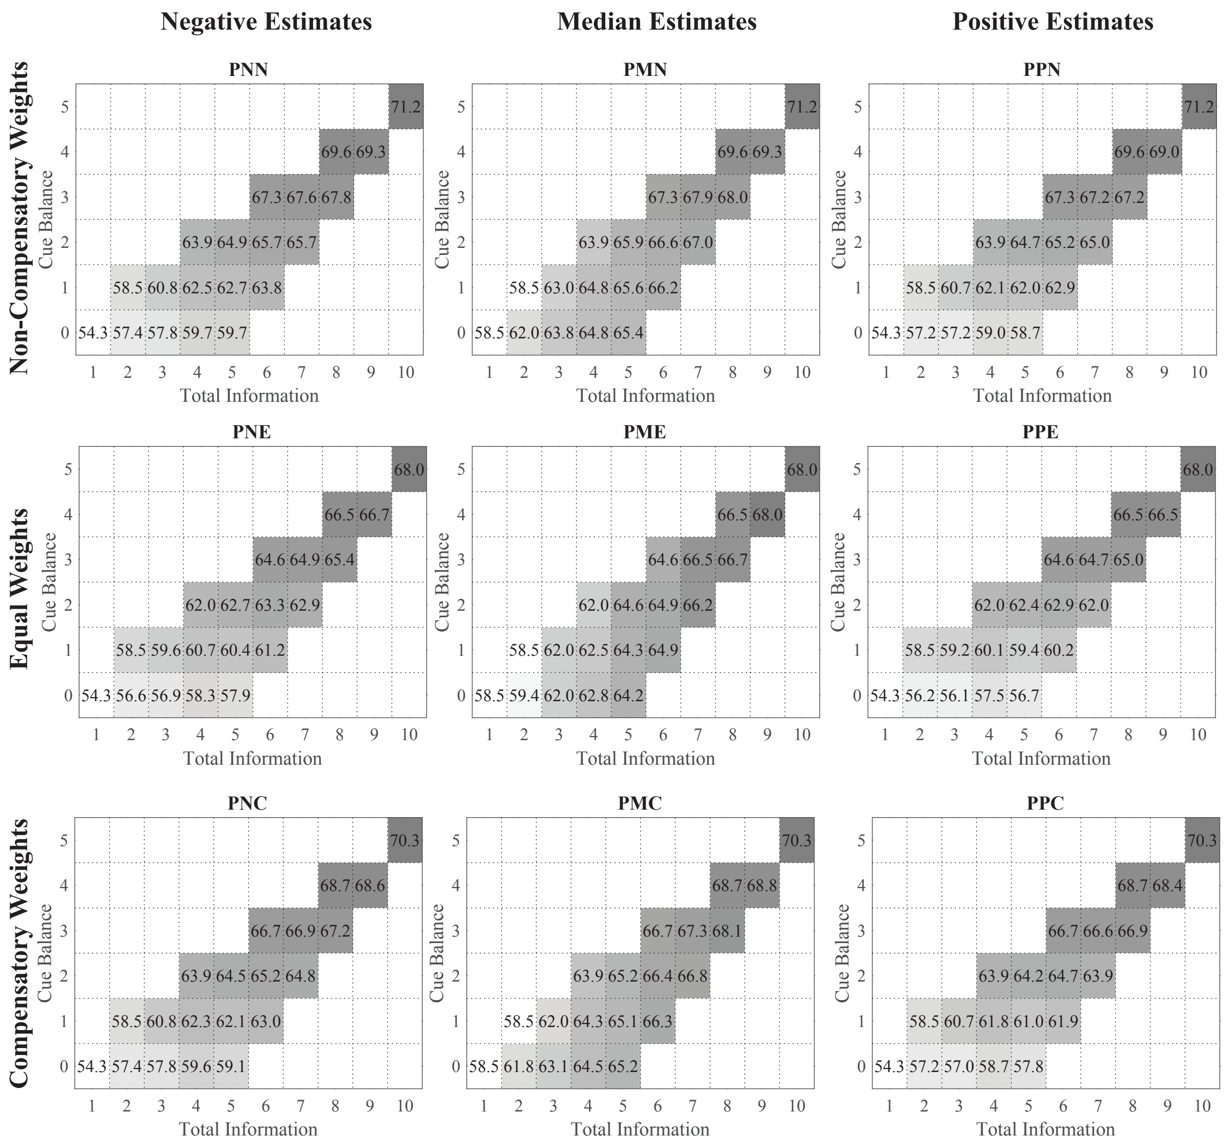

Figure A.3: Two-way interaction between total information and cue balance for strategies with prior cutoff values and each combination of estimates and weights.

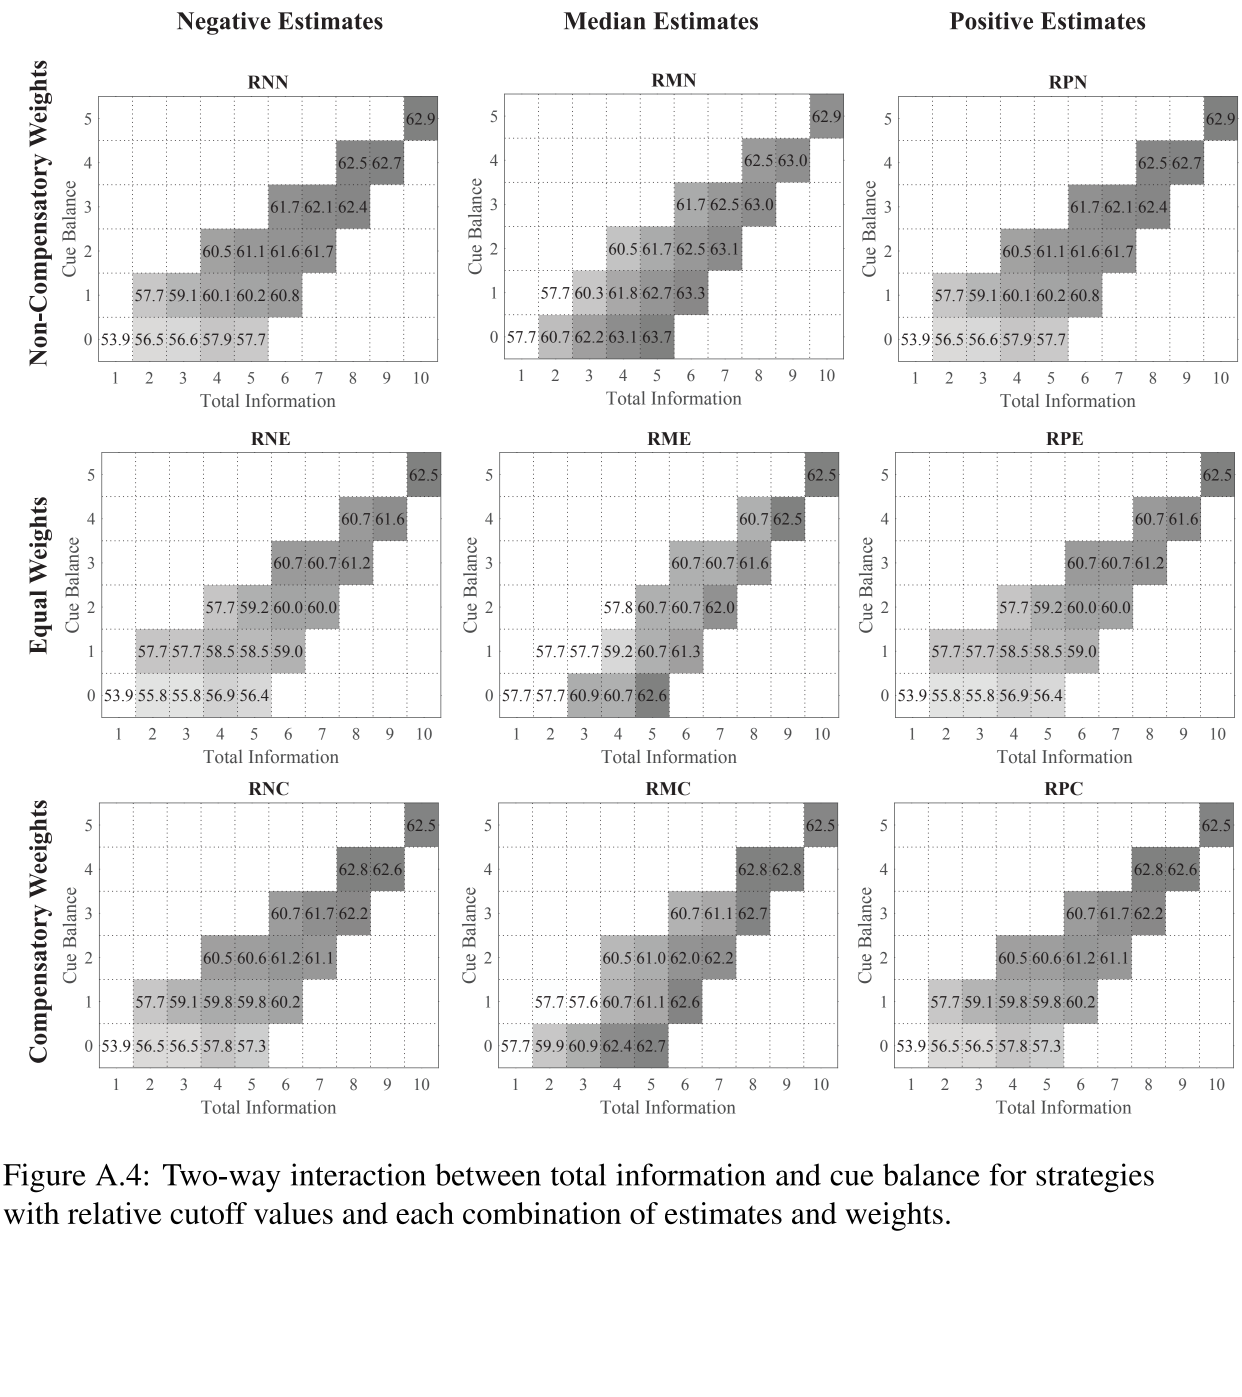

Figure A.4: Two-way interaction between total information and cue balance for strategies with relative cutoff values and each combination of estimates and weights.

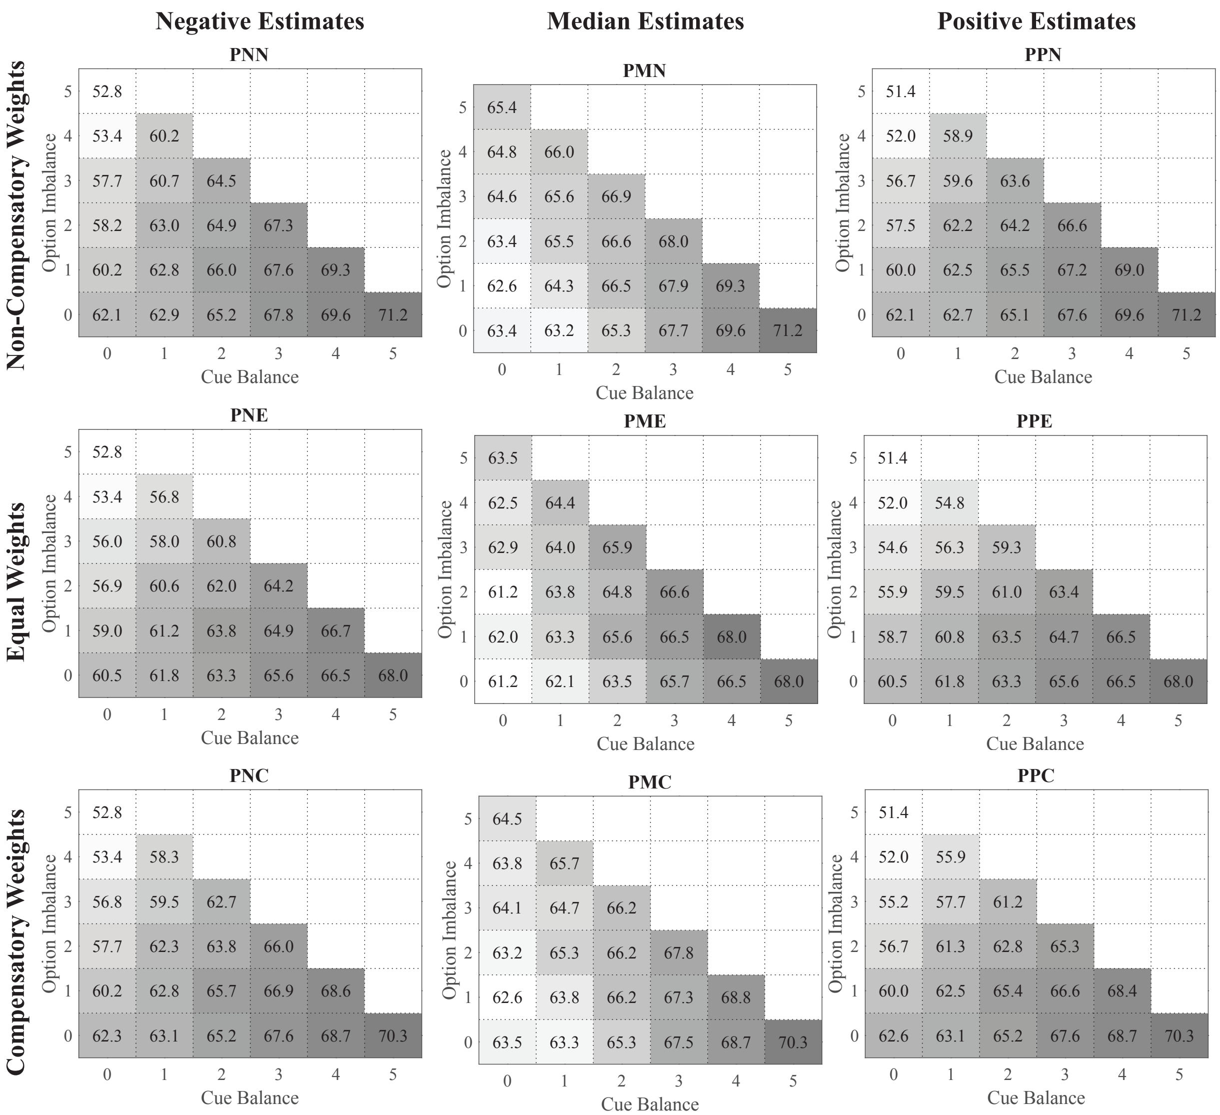

Figure A.5: Two-way interaction between cue balance and option imbalance for strategies with prior cutoff values and each combination of estimates and weights.

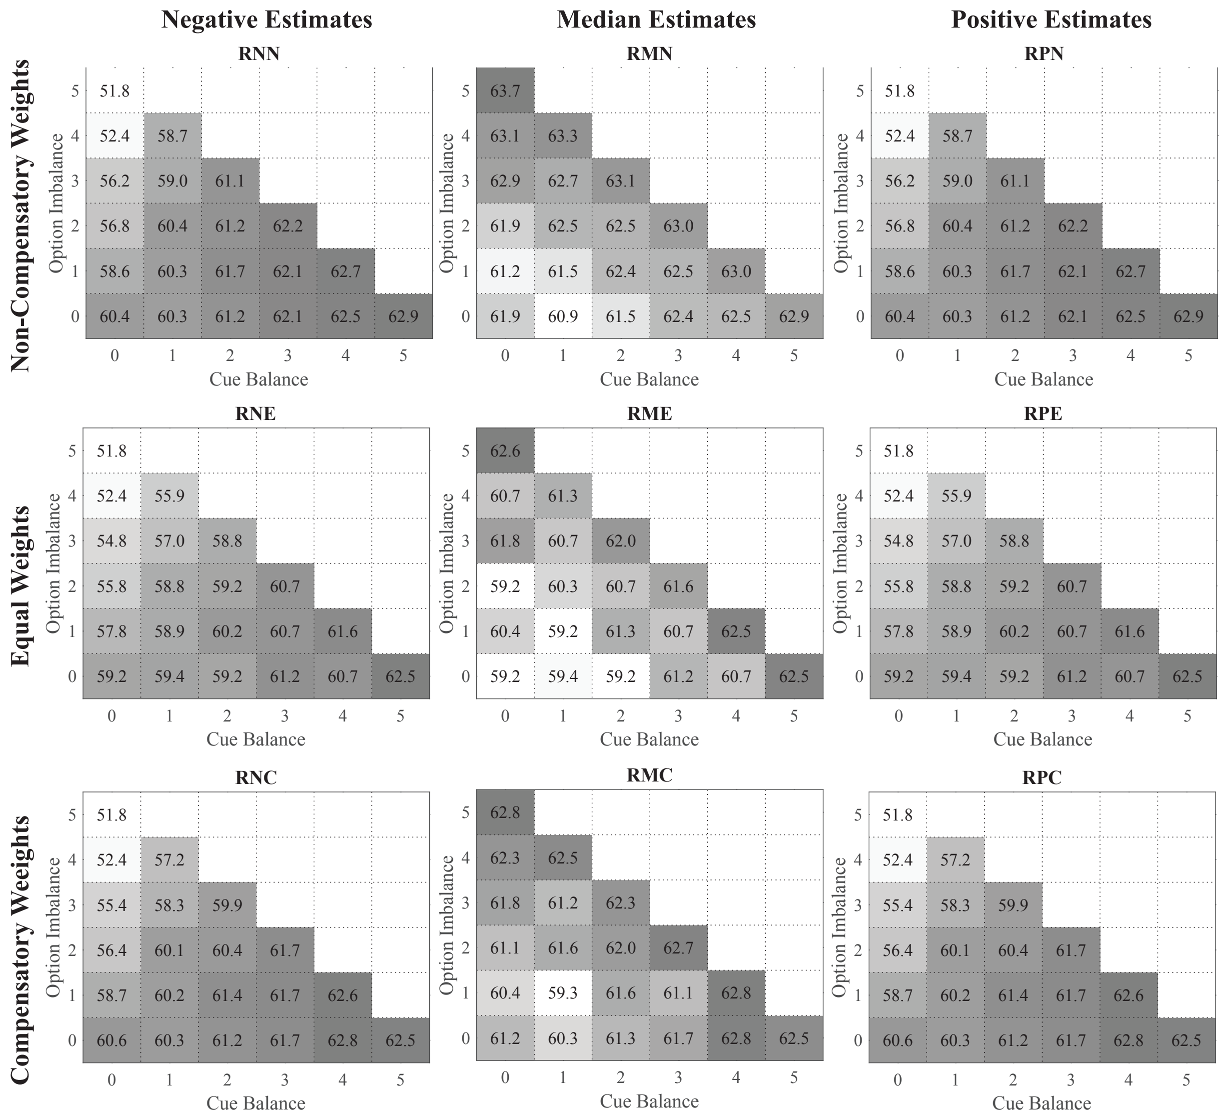

Figure A.6: Two-way interaction between cue balance and option imbalance for strategies with relative cutoff values and each combination of estimates and weights.

Table A.7: The average accuracy change of the information acquisition and restriction types.

Strategy	**Cue Scale**	**Estimate**	**Weights**	**-4** **(-1, -1)**	**-3** **(-1, 1)**	**-2** **(0, -1)**	**-1** **(0, 1)**	**1** **(0, 1)**	**2** **(0, -1)**	**3** **(1, 1)**	**4** **(1, -1)**
RMC	Continuous	Median	Compensatory	0.33	-0.09	-0.99	-1.38	0.99	1.38	-0.33	0.09
PMC	Binary	Median	Compensatory	-0.52	-1.11	-1.17	-1.65	1.17	1.65	0.52	1.11
RMN	Continuous	Median	Non-CF	0.20	0.19	-1.05	-0.98	1.05	0.98	-0.20	-0.19
PMN	Binary	Median	Non-CF	-0.98	-1.04	-1.13	-1.04	1.13	1.04	0.98	1.04
RME	Continuous	Median	Equal	-0.15	0.10	-1.28	-0.89	1.28	0.89	0.15	-0.10
PME	Binary	Median	Equal	-1.13	-0.90	-1.46	-1.10	1.46	1.10	1.13	0.90
RNN	Continuous	Negative	Non-CF	0.14	-1.53	-0.19	-2.21	0.19	2.21	-0.14	1.53
PNN	Binary	Negative	Non-CF	-0.43	-2.99	-0.35	-3.21	0.35	3.21	0.43	2.99
RNE	Continuous	Negative	Equal	0.43	-1.59	-0.05	-2.28	0.05	2.28	-0.43	1.59
PNE	Binary	Negative	Equal	0.13	-2.92	0.11	-3.06	-0.11	3.06	-0.13	2.92
RNC	Continuous	Negative	Compensatory	0.48	-1.66	0.00	-2.42	0.00	2.42	-0.48	1.66
PNC	Binary	Negative	Compensatory	0.20	-3.18	0.03	-3.54	-0.03	3.54	-0.20	3.18
RPN	Continuous	Positive	Non-CF	0.14	-1.53	-0.19	-2.21	0.19	2.21	-0.14	1.53
PPN	Binary	Positive	Non-CF	-0.27	-3.21	0.04	-3.28	-0.04	3.28	0.27	3.21
RPE	Continuous	Positive	Equal	0.43	-1.59	-0.05	-2.28	0.05	2.28	-0.43	1.59
PPE	Binary	Positive	Equal	0.54	-3.43	0.60	-3.49	-0.60	3.49	-0.54	3.43
RPC	Continuous	Positive	Compensatory	0.48	-1.66	0.00	-2.42	0.00	2.42	-0.48	1.66
PPC	Binary	Positive	Compensatory	0.64	-3.69	0.61	-3.96	-0.61	3.96	-0.64	3.69

Table A.8: The positive accuracy count of the information acquisition and restriction types.

Strategy	Cue Scale	Estimate	Weights	-4 (-1, -1)	-3 (-1, 1)	-2 (0, -1)	-1 (0, 1)	1 (0, 1)	2 (0, -1)	3 (1, 1)	4 (1, -1)
PMC	Binary	Median	Compensatory	48%	47%	46%	46%	54%	54%	52%	53%
PME	Binary	Median	Equal	46%	47%	45%	46%	55%	54%	54%	53%
PMN	Binary	Median	Non-CF	46%	46%	46%	47%	54%	54%	54%	54%
RMC	Continuous	Median	Compensatory	51%	50%	46%	45%	54%	55%	49%	50%
RME	Continuous	Median	Equal	49%	50%	45%	46%	55%	54%	51%	50%
RMN	Continuous	Median	Non-CF	51%	51%	46%	47%	54%	53%	49%	49%
PNC	Binary	Negative	Compensatory	52%	37%	49%	37%	51%	63%	48%	63%
PNE	Binary	Negative	Equal	51%	38%	50%	39%	50%	61%	49%	62%
PNN	Binary	Negative	Non-CF	50%	36%	48%	37%	52%	63%	50%	64%
RNC	Continuous	Negative	Compensatory	53%	43%	50%	40%	50%	60%	47%	57%
RNE	Continuous	Negative	Equal	52%	43%	49%	41%	51%	59%	48%	57%
RNN	Continuous	Negative	Non-CF	54%	41%	51%	38%	49%	62%	46%	59%
PPC	Binary	Positive	Compensatory	53%	37%	51%	37%	49%	63%	47%	63%
PPE	Binary	Positive	Equal	51%	37%	51%	39%	49%	61%	49%	63%
PPN	Binary	Positive	Non-CF	51%	36%	49%	37%	51%	63%	49%	64%
RPC	Continuous	Positive	Compensatory	53%	43%	50%	40%	50%	60%	47%	57%
RPE	Continuous	Positive	Equal	52%	43%	49%	41%	51%	59%	48%	57%
RPN	Continuous	Positive	Non-CF	54%	41%	51%	38%	49%	62%	46%	59%

A.3 Validation Human-in-the-Loop Experiment

A.3.1 Institutional Review Board Documents

Georgia Tech Office of Research Integrity Assurance

Protocol Number: H16325
Funding Agency: Navy/ONR
Review Type: Exempt, Category 2
Title: Impact of Incomplete Information on Performance in a Naval Defense Task
Number of Subjects: 120

August 9, 2016
Karen Feigh
Aerospace Engineering
0150

Dear Dr. Feigh:

The Institutional Review Board (IRB) has carefully considered the referenced protocol. Your approval is effective as of **08/09/2016**. The proposed procedures are exempt from further review by the Georgia Tech Institutional Review Board.

> *Minimal risk research qualified for exemption status under 45 CFR 46 101b. 2.*
>
> ***DOD COMPLIANCE CONCURRENCE MUST BE OBTAINED BEFORE WORK WITH HUMAN SUBJECTS MAY BEGIN, DESPITE GEORGIA TECH IRB APPROVAL BEING ISSUED.***
> *Obtaining DOD compliance concurrence is the responsibility of the Principal Investigator.*
> *DOD compliance concurrence must be documented in the Georgia Tech IRB record.*

Thank you for allowing us the opportunity to review your plans. If any complaints or other evidence of risk should occur, or if there is a significant change in the plans, the IRB must be notified.

If you have any questions concerning this approval or regulations governing human subject activities, please feel free to contact Dennis Folds, IRB Chair, at 404/407-7262, or me at 404/385-5208.

Sincerely,

Scott S. Katz, MS, CIP
Compliance Officer
Georgia Tech Office of Research Integrity Assurance

cc: Dr. Dennis Folds, IRB Chair

Unit of the University System of Georgia An Equal Education and Employment Opportunity Institution

Figure A.7: Georgia Tech Institute Review Board approval.

School of Aerospace Engineering
Cognitive Engineering Center
Georgia Institute of Technology
Human Subject Consent

1. **Project Title**: Impact of Incomplete Information on Performance in a Naval Defense Task

2. **Principal Investigator**: Dr. Karen Feigh, (404) 385-7686, karen.feigh@gatech.edu
   Graduate Students: Marc Canellas (marc.c.canellas@gatech.edu)
   Rachel Haga (rachel.haga@gatech.edu)

3. **Protocol and Consent Title:** Impact of Incomplete Information on Performance in a Naval Defense Task – Experiment 1

4. **Introduction**: You are being asked to participate in a research study. The purpose of this study is to examine how different combinations of environmental factors can affect your decision making process in a naval defense task. As a participant of this study, you will be interacting with a computer interface in order to make decisions about how to engage with various targets in order to defend your ship(s). The scenarios presented are purely hypothetical.

5. **Procedures**:
   Introduction: The introductory briefing:
   - Seeks informed consent
   - Explains the experiment.
   - Details the schedule of events.
   - Explains the interface.

   Training Sets: This training will familiarize you with the interface and the decision you will need to make. At the end of the training sets, you should feel comfortable with the interface and decisions.
   Experimental Sets: Prior to each experiment you will be given a brief description of the task you will perform. You will be required to make decisions within the time allotted. For each decision task presented, there is a correct answer.

   The entire procedure will last approximately 4 hours. You are free to request a break at any time.

6. **Foreseeable Risks or Discomforts**: Every study involves some risk. This study is considered to have low risk. There is the same possibility of discomfort or fatigue while interacting with the interface that you would find interacting with a computer.

   You will be provided breaks during the course of the experiment.

7. **Benefits**: There are no direct benefits to you for participating in this research study. The study may help us understand how decision makers perform in different environments and to design decision support systems.

8. **Compensation/Costs**: There is no cost to you. All participants will receive $10 for showing up as compensation for your time. Then for each correct decision made in each task, you will receive $0.10.

   Please note that U.S. Tax Law requires a mandatory withholding of 30% for nonresident alien payments of any type. Your address and citizenship/visa status may be collected for compensation purposes only. This information will be shared only with the Georgia Tech department that issues compensation, if any, for your participation.

9. **Confidentiality**: The information that you give in the study will be handled confidentially. Personal information about you will not be published or made available to any third party in any form. Your name will

APPROVED *Consent Form Approved by Georgia Tech IRB: August 09, 2016 - Indefinite*

Figure A.8: Human-subjects study consent form approved by Georgia Tech Institute Review Board - page 1.

not be used in any report (pseudonyms such as ‘Participant X’ will be used), and we will be careful to ensure that no identifiable patterns in your actions and questionnaire responses (including demographic information such as your experience) are released in a way that would allow for any outsider to guess your identity. The raw data will only be shared with the research team. Once the analysis and documentation of this experiment are complete, the video and audio files will be destroyed; electronic and paper stores of results will be archived in a locked facility within the principal investigator’s Georgia Tech office or laboratory. To make sure that this research is being carried out in the proper way, the Georgia Institute of Technology Institute Review Board (IRB) will review study records. The Office of Human Research Protections may also look at study records.

10. **Injury/Adverse Reactions**: Reports of injury or reaction should be made to the Principal Investigator of this research study. Neither the Georgia Institute of Technology nor the principal investigator has made provision for payment of costs associated with any injury resulting from participation in this study.

11. **Contact Person**: If you have questions about the research, call or write Dr. Karen Feigh at (404) 894-0199, Montgomery Knight Building, Room 419, Georgia Institute of Technology, 270 Ferst Drive, Atlanta GA 30332-0150.

12. **Voluntary Participation/Withdrawal**: You have the right to withdraw from the study at any time without penalty. The audio recordings along with all questionnaires and transcripts will be destroyed upon your withdrawal from the study.

13. **Participant’s Rights:**
    - Your participation in this study is voluntary. You do not have to be in this study if you don't want to be.
    - You have the right to change your mind and leave the study at any time without giving any reason and without penalty.
    - Any new information that may make you change your mind about being in this study will be given to you.
    - You will be given a copy of this consent form to keep.
    - You do not waive any of your legal rights by signing this consent form.

If you have any questions about your rights as a research volunteer, call or write:

Ms. Melanie Clark
Office of Research Integrity Assurance
Georgia Institute of Technology
Atlanta, GA 30332-0420
Voice (404) 894-6942

Your signature below indicates that the researchers have answered all of your questions to your satisfaction, and that you consent to volunteer for this study.

Subject’s Signature: ______________________________________ Date: _______________

Subject’s Name: ______________________________________

Investigator’s Signature: ___________________________________ Date: _______________

Investigator’s Name: ______________________________

APPROVED *Consent Form Approved by Georgia Tech IRB: August 09, 2016 - Indefinite*

Figure A.9: Human-subjects study consent form approved by Georgia Tech Institute Review Board - page 2.

A.3.2 General Experiment Information

A complete definition of a decision task requires not just a description of the cue scores, but also a description of which cue scores are known or unknown to the decision maker. Combinations of incomplete information were generated via a full factorial combination of ones (denoting that a cue value is presented to the participant) and zeros (denoting that a cue value is not presented to the participant), as shown in Table A.9. Combining the 16 options with 16 incomplete information combinations resulted in 32,640 decision tasks. Each specific decision task had a unique decision task identity (ID) consisting of a concatenation of option numbers and incomplete information numbers. For example, task ID $[1.2.3.4]$ is the decision task with Option 1 consisting of Option No. 1 and Incomplete Information No. 2, and Option 2 consisting of Option No. 3 and Incomplete Information No. 4.

Table A.9: Incomplete information combinations.

Incomplete Information No.	**Cue 1, z_1**	**Cue 2, z_2**	**Cue 3, z_3**	**Cue 4, z_4**
	Altitude	**Speed**	**Distance from Corridor**	**Size**
1	0	0	0	0
2	0	0	0	1
3	0	0	1	0
4	0	0	1	1
5	0	1	0	0
6	0	1	0	1
7	0	1	1	0
8	0	1	1	1
9	1	0	0	0
10	1	0	0	1
11	1	0	1	0
12	1	0	1	1
13	1	1	0	0
14	1	1	0	1
15	1	1	1	0
16	1	1	1	1

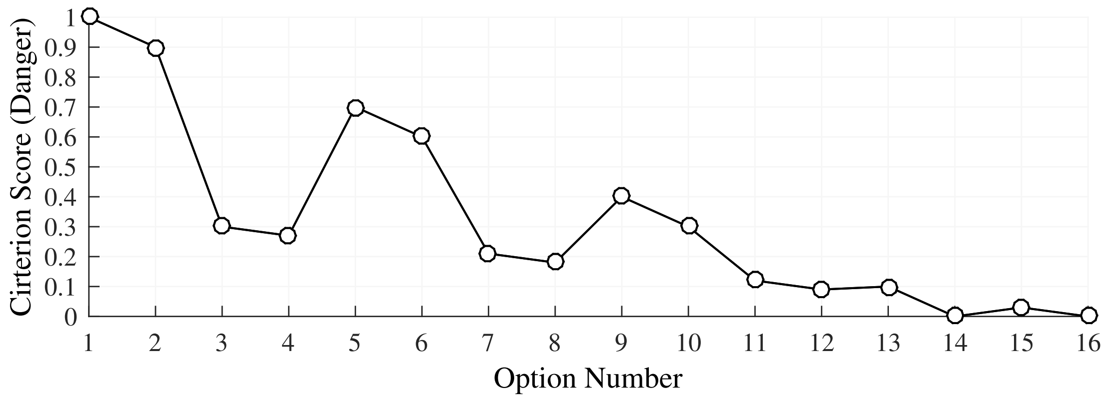

Figure A.10: Criterion scores of the 16 options.

Table A.10: Transformations of experimental cue scores to presented cue scores for the participants.

Cue	**Cue Score**	**Cue Value Lower Bound**	**Cue Value Upper Bound**	**Increment**
Altitude (ft)	Low (1)	500	5,000	500
	High (0)	30,000	40,000	500
Speed (kts)	Fast (1)	1,000	1,200	50
	Slow (0)	200	400	50
Distance from Corridor (nm)	Far (1)	40	50	2
	Near (0)	10	0	2
Size (m^2)	Small (1)	0.1	1	0.1
	Large (0)	80	100	5

Table A.11: Description of the decision task "Type" in Tables A.12 and A.13

Type		**A**	**B-P**	**B-N**	**E-P**	**E-N**
Measures Bias		N	Y	Y	Y	Y
Measures Estimates		N	N	N	Y	Y
Unbiased, Accurate	Average Estimates	Y	Y	Y	Y	Y
Unbiased, Inaccurate	Average Estimates	N	N	N	N	N
Positive Biased	Negative Estimates	Y	Y	N	N	Y
Negative Biased	Positive Estimates	Y	N	Y	Y	N

Table A.12: Description the three blocks of easy difficulty tasks: incomplete information, type, and decision task ID.

Total Information	Information Imbalance	Complete Attribute Pairs	Type	Trial A	Trial B	Trial C
2	0	0	A	5.16.9.2	1.4.3.2	1.4.5.2
2	0	0	A	5.16.2.9	1.7.2.3	2.12.9.3
2	0	1	A	1.4.2.2	5.12.3.3	1.7.3.3
2	0	1	A	1.4.3.3	5.16.2.2	5.14.2.2
2	2	0	B-N	2.7.1.4	1.10.1.11	2.11.1.13
2	2	0	B-P	2.13.10.1	5.15.7.1	2.4.4.1
4	0	0	E-N	3.15.11.6†	4.15.11.6**	3.13.11.6**
4	0	0	E-P	2.13.6.11	2.10.6.11	2.14.6.11‡
4	0	1	E-N	4.5.6.13†	4.6.6.13†	3.14.11.7†
4	0	1	E-P	3.14.7.11†	4.14.7.11†	3.13.7.11**
4	0	2	A	1.7.7.7	2.13.7.7	2.7.4.4
4	0	2	A	2.12.7.7	1.8.10.10	1.7.6.6
4	2	0	E-N	2.11.2.15	5.12.2.15	2.13.5.12
4	2	0	E-P	5.14.15.2	5.16.15.2	2.11.12.5
4	2	1	E-N	1.4.9.15	5.14.2.12	5.15.5.14
4	2	1	E-P	2.8.12.9	2.4.12.9	5.15.14.2
4	4	0	E-N	5.14.1.16	2.11.1.16	5.12.1.16
4	4	0	E-P	2.11.16.1	5.12.16.1	2.4.16.1
6	0	2	E-N	4.14.12.8†	3.13.15.8†	3.13.12.8**
6	0	2	E-P	4.14.8.15†	4.14.8.12†	3.13.8.12**
6	0	3	A	5.12.14.14	5.14.15.15	1.10.14.14
6	0	3	A	2.13.8.8	1.7.12.12	1.8.15.15
6	2	2	E-N	2.4.10.16	1.4.10.16	5.12.6.16
6	2	2	E-P	2.7.16.10	2.4.16.13	5.14.16.4
8	0	4	A	1.7.16.16	2.11.16.16	5.16.16.16
8	0	4	A	2.13.16.16	1.3.16.16	1.10.16.16

*Extra-hard difficulty tasks: 0.02; **Medium-hard difficulty tasks: [0.17,0.24]
† Medium difficulty tasks: [0.27,0.43]; ‡ Extra-easy difficulty tasks: [0.90]

Table A.13: Description the three blocks of hard difficulty tasks: incomplete information, type, and decision task ID.

Total Information	Information Imbalance	Complete Attribute Pairs	Type	Trial A	Trial B	Trial C
2	0	0	A	5.6.2.5	9.10.2.9	4.7.5.3
2	0	0	A	3.8.9.3	13.16.2.9	3.7.2.3
2	0	1	A	8.10.5.5	3.7.5.5	7.13.9.9
2	0	1	A	4.7.5.5	11.14.5.5	13.16.3.3
2	2	0	B-N	12.14.1.11	7.13.1.11	11.15.1.10
2	2	0	B-P	13.15.7.1	8.13.13.1	12.15.6.1
4	0	0	E-N	4.7.11.6	7.13.11.6	8.13.11.6
4	0	0	E-P	12.14.6.11	4.13.6.11**	10.13.6.11**
4	0	1	E-N	8.13.11.7	7.13.11.7	4.13.11.7**
4	0	1	E-P	12.14.7.11	11.14.7.11	11.13.7.11*
4	0	2	A	13.14.6.6	12.15.7.7	8.12.11.11
4	0	2	A	11.15.7.7	7.11.11.11	11.15.13.13
4	2	0	B-N	12.14.5.12	7.13.3.14	11.15.3.14
4	2	0	B-P	3.8.15.2	3.8.12.5	4.8.12.5
4	2	1	B-N	7.12.2.8	12.15.3.8	8.11.3.12
4	2	1	B-P	13.15.15.9	4.8.15.5	4.7.14.2
4	4	0	B-N	12.15.1.16	13.14.1.16	7.11.1.16
4	4	0	B-P	3.8.16.1	5.6.16.1	3.7.16.1
6	0	2	E-N	7.13.12.8	4.13.15.8**	8.14.12.8**
6	0	2	E-P	12.14.8.12	12.14.8.15	11.13.8.12*
6	0	3	A	9.10.8.8	5.6.12.12	13.14.14.14
6	0	3	A	1.2.14.14	5.6.8.8	5.6.14.14
6	2	2	B-N	4.7.10.16	11.14.4.16	12.15.10.16
6	2	2	B-P	13.16.16.10	5.6.16.10	5.6.16.4
8	0	4	A	7.11.16.16	3.7.16.16	7.13.16.16
8	0	4	A	13.15.16.16	13.14.16.16	12.15.16.16

APPENDIX B
PUBLICATIONS

B.1 Published Articles

B.1.1 Journal Articles

1. Canellas, M. C. and Feigh, K. M. (2017). Heuristic information acquisition and restriction rules for decision support. IEEE Transactions on Human-Machine Systems. (In press)

2. Canellas, M. C. and Feigh, K. M. (2016b). Toward simple representative mathematical models of naturalistic decision making through fast-and-frugal heuristics. Journal of Cognitive Engineering and Decision Making, 10(3):255–267

3. Canellas, M. C., Feigh, K. M., and Chua, Z. K. (2015). Accuracy and effort of decision-making strategies with incomplete information: Implications for decision support system design. IEEE Transactions on Human-Machine Systems, 45(6):686–701

B.1.2 Conference Papers with Podium Presentations

1. Canellas, M. C. and Feigh, K. M. (2014). Heuristic decision making with incomplete information: Conditions for ecological rationality. In Systems, Man and Cybernetics (SMC), 2014 IEEE International Conference on, pages 1963–1970

2. Canellas, M. C., Feigh, K. M., and Chua, Z. K. (2014). Accuracy and effort of decision making strategies with incomplete information. In Cognitive Methods in Situation Awareness and Decision Support (CogSIMA), 2014 IEEE International Inter-Disciplinary Conference on, pages 7–13

B.1.3 Podium Presentations (Only)

1. Canellas, M. C., Feigh, K. M., and Haga, R. A. (2016). Mathematical representations of human judgment and decision making in military contexts. In Military Operations Research Society Emerging Techniques Special Meeting (MORS METSM) in Washington, D.C.

B.1.4 Poster Presentations (Only)

1. Canellas, M. C. and Feigh, K. M. (2016a). A general linear model of fast-and-frugal judgment and decision making. In Summer Institute on Bounded Rationality, in Berlin, Germany

B.2 Planned

B.2.1 Journal Articles

- **A General Linear Model of Judgment and Decision Making.** Presents the development of the general linear model of judgment and decision making strategies in Chap. 3. Discusses the mathematical, computational, and theoretical contributions of the model. Planned for *Psychological Review* or *Journal of Mathematical Pscyhology*.

- **Determinants of Decision Making Accuracy with Incomplete Information.** Presents the results in Chap. 7 of decision making strategies with incomplete information extended to decision tasks with three options instead of two. Planned for *Judgment and Decision Making*.

- **The Reality Gap: What makes distributions of incomplete information difficult?** Presents the human-subjects study results in Chap. 8 as a contrast to the computer simulation studies (Chap. **??**). Planned for *Psychological Science* or *Journal of Experimental Psychology: General*.

- **As they are: Representing and supporting the heuristics of military decision makers.** Review of the capabilities of the general linear model for modeling and simulating military decision makers. Additional review of the prior work analyzing military heuristics from the naturalistic decision making and fast-and-frugal heuristics programs. Planned for *Military Operations Research Journal*.

B.2.2 Magazine Articles

- **As they are: Representing and supporting the heuristics of military decision makers.** Review of the capabilities of the general linear model for modeling and simulating heuristic decision makers. Additional review of the prior work analyzing military heuristics from the naturalistic decision making and fast-and-frugal heuristics programs. Planned for *Phalanx*, the magazine for the Military Operations Research Society.

BIBLIOGRAPHY

Aikman, D., Galesic, M., Gigerenzer, G., Kapadia, S., Katsikopoulos, K., Kothiyal, A., Murphy, E., and Neumann, T. (2014). Taking uncertainty seriously: simplicity versus complexity in financial regulation. Financial Stability Paper 28, Bank of England.

Anderson, J. R. (2007). How can the human mind occur in the physical universe? Oxford University Press, Inc.

Arbuckle, J. L. (1996). Full Information Estimation in the Presence of Incomplete Data.

Baron, J. (1985). Rationality and Intelligence. Cambridge University Press, Cambridge, England.

Bass, E. J. (2002). Human-Automated Judgment Learning: A Research Paradigm Based on Interpersonal Learning to Investigate Human Interaction with Automated Judgments of Hazards. PhD thesis, Georgia Institute of Technology.

Beattie, J., Baron, J., Hershey, J., and Spranca, M. (1994). Determinants of decision attitude. Journal of Behavioral Decision Making, 7:129–144.

Bertuccelli, L., Choi, H., Cho, P., and How, J. (2009). Real-time multi-uav task assignment in dynamic and uncertain environments.

Bettman, J. R., Johnson, E. J., and Payne, J. W. (1990). A componential analysis of cognitive effort in choice. Organizational Behavior and Human Decision Processes, 45:111–139. Negative utility of expending effort is combined with the expected utility of decision accuracy to predict the DM method.

Brighton and Gigerenzer, G. (2012). Ecological rationality: Intelligence in the world, chapter How Heuristics Handle Uncertainty, pages 3–30. Oxford University Press.

Bröder, A. (2000). Assessing the empirical validity of the" take-the-best" heuristic as a model of human probabilistic inference. Journal of Experimental Psychology: Learning, Memory, and Cognition, 26(5):1332.

Brunswik, E. (1943). Organismic achievement and environmental probability. Psychological Review, 50(3):255 – 272.

Canellas, M. C. and Feigh, K. M. (2014). Heuristic decision making with incomplete information: Conditions for ecological rationality. In Systems, Man and Cybernetics (SMC), 2014 IEEE International Conference on, pages 1963–1970.

Canellas, M. C. and Feigh, K. M. (2016a). A general linear model of fast-and-frugal judgment and decision making. In Summer Institute on Bounded Rationality, in Berlin, Germany.

Canellas, M. C. and Feigh, K. M. (2016b). Toward simple representative mathematical models of naturalistic decision making through fast-and-frugal heuristics. Journal of Cognitive Engineering and Decision Making, 10(3):255–267.

Canellas, M. C. and Feigh, K. M. (2017). Heuristic information acquisition and restriction rules for decision support. IEEE Transactions on Human-Machine Systems. (In press).

Canellas, M. C., Feigh, K. M., and Chua, Z. K. (2014). Accuracy and effort of decision making strategies with incomplete information. In Cognitive Methods in Situation Awareness and Decision Support (CogSIMA), 2014 IEEE International Inter-Disciplinary Conference on, pages 7–13.

Canellas, M. C., Feigh, K. M., and Chua, Z. K. (2015). Accuracy and effort of decision-making strategies with incomplete information: Implications for decision support system design. IEEE Transactions on Human-Machine Systems, 45(6):686–701.

Canellas, M. C., Feigh, K. M., and Haga, R. A. (2016). Mathematical representations of human judgment and decision making in military contexts. In Military Operations Research Society Emerging Techniques Special Meeting (MORS METSM) in Washington, D.C.

Cohen, M. S., Freeman, J. T., and Wolf, S. (1996). Metarecognition in time-stressed decision making: Recognizing, critiquing, and correcting. Human Factors: The Journal of the Human Factors and Ergonomics Society, 38:206–219.

Cooksey, R. W. (1996). Judgment Analysis: Theory, Methods, and Applications. Academic.

Czerlinski, J., Gigerenzer, G., and Goldstein, D. G. (1999). Simple Heuristics that Make Use Smart, chapter How good are simple heuristics?, pages 97–118. New York: Oxford University Press.

Dawes, R. M. (1979). The robust beauty of improper linear models in decision making. American psychologist, 34(7):571.

Dawes, R. M. and Corrigan, B. (1974). Linear models in decision making. Psychological bulletin, 81(2):95.

Dick, A., Chakravarti, D., and Biehal, G. (1990). Memory-based inferences during consumer choice. Journal of Consumer Research, 17(1):pp. 82–93.

Dieckmann, A. and Rieskamp, J. (2007). The influence of information redundancy on probabilistic inferences. Memory & Cognition, 35(7):1801–1813.

Dudey, T. and Todd, P. M. (2001). Making good decisions with minimal information: Simultaneous and sequential choice. Journal of Bioeconomics, 3(2-3):195–215.

Edwards, W. (1954). The theory of decision making. Psychological Bulletin, 51:380–417.

Edwards, W. and Fasolo, B. (2001). Decision technology. Annual Review of Psychology, 52(1):581.

Einhorn, H. J., Kleinmuntz, D. N., and Kleinmuntz, B. (1979). Linear regression and process-tracing models of judgment. Psychological Review, 86(5):465.

Elwyn, G., Edwards, A., Eccles, M., and Rovner, D. (2001). Decision analysis in patient care. The Lancet, 358(9281):571–574.

Fargen, K. M. and Hoh, B. L. (2014). The debate over eponyms. Clinical Anatomy, 27(8):1137–1140.

Fasolo, B., McClelland, G. H., and Todd, P. M. (2007). Escaping the tyranny of choice: When fewer attributes make choice easier. Marketing Theory, 7(1):13–26.

Feigh, K. M., Dorneich, M. C., and Hayes, C. C. (2012). Toward a characterization of adaptive systems: A framework for researchers and system designers. Human Factors, 54(6):1008–1024.

Fischer, J. E., Steiner, F., Zucol, F., Berger, C., Martignon, L., Bossart, W., Altwegg, M., and Nadal, D. (2002). Use of simple heuristics to target macrolide prescription in children with community-acquired pneumonia. Archives of pediatrics & adolescent medicine, 156(10):1005–1008.

Garcia-Retamero, R. and Dhami, M. K. (2009). Take-the-best in expert-novice decision strategies for residential burglary. Psychonomic Bulletin & Review, 16(1):163–169.

Garcia-Retamero, R. and Rieskamp, J. (2008). Adaptive mechanisms for treating missing information: A simulation study. The Psychological Record, 58.

Garcia-Retamero, R. and Rieskamp, J. (2009). Do people treat missing information adaptively when making inferences? The Quarterly Journal of Experimental Psychology, 62(10):1991–2013.

Gigerenzer, G. (2004). Blackwell handbook of judgment and decision making, chapter

Fast and Frugal Heuristics: The Tools of Bounded Rationality, pages 62–88. Blackwell, Oxford, UK.

Gigerenzer, G. (2007). Gut Feelings: The Intelligence of the Unconscious. Viking.

Gigerenzer, G. and Gaissmaier, W. (2011). Heuristic decision making. Annual review of psychology, 62:451–482.

Gigerenzer, G. and Goldstein, D. G. (1996). Reasoning the fast and frugal way: models of bounded rationality. Psychological review, 103(4):650.

Gigerenzer, G., Hoffrage, U., and Kleinbölting, H. (1991). Probabilistic mental models: A brunswikian theory of confidence. Psychological Review, 98(4):506–528.

Gigerenzer, G., Todd, P. M., Group, A. R., et al. (1999). Simple heuristics that make us smart. Oxford University Press New York.

Goldstein, D. G. and Gigerenzer, G. (2002). Models of ecological rationality: the recognition heuristic. Psychological review, 109(1):75.

Goldstein, W. M. and Hogarth, R. M. (1997). Research on Judgment and Decision Making: Currents, Connections, and Controversies, chapter Judgment and decision research: Some historical context, pages 3–68. Cambridge Series on Judgment and Decision Making.

Gonzalez, C. (2005). Decision support for real-time, dynamic decision-making tasks. Organizational Behavior and Human Decision Processes, 96(2):142 – 154.

Green, L. and Mehr, D. R. (1997). What alters physicians' decisions to admit to the coronary care unit? The Journal of Family Practice, 45:219–226.

Grüne-Yanoff, T. and Hertwig, R. (2016). Nudge versus boost: How coherent are policy and theory? Minds and Machines, 26(1):149–183.

Hammond, K. (1955). Probabilistic functionalism and the clinical method. Psychological Review, 62:255–262.

Hammond, K. R. (1988). Judgement and decision making in dynamic tasks. Technical report, Colorado Univ At Boulder Center For Research On Judgment And Policy.

Hammond, K. R., Hamm, R. M., Grassia, J., and Pearson, T. (1987). Direct comparison of the efficacy of intuitive and analytical cognition in expert judgment. IEEE Transactions on Systems Man and Cybernetics, 17:753–770.

Hammond, K. R., Hursch, C. J., and Todd, F. J. (1964). Analyzing the components of clinical inference. Psychological Review, 71:438–456.

Heller, C. M. (2013). A computational model of engineering decision making. Master's thesis, Georgia Institute of Technology, School of Aerospace Engineering.

Hertwig, R. and Todd, P. (2003). More is not Always Better: The Benefits of Cognitive Limits, chapter 10, pages 213–231. John Wiley & Sons Ltd.

Hilbig, B. E. (2010). Reconsidering "evidence" for fast-and-frugal heuristics. Psychonomic Bulletin & Review, 17(6):923–930.

Hogarth, R. M. and Karelaia, N. (2005a). Ignoring information in binary choice with continuous variables: When is less more? Journal of Mathematical Psychology, 49(2):115 – 124.

Hogarth, R. M. and Karelaia, N. (2005b). Simple models for multiattribute choice with many alternatives: When it does and does not pay to face trade-offs with binary attributes. Management Science, 51(12):1860–1872.

Hogarth, R. M. and Karelaia, N. (2006). Take-the-best? and other simple strategies: Why and when they work well with binary cues. Theory and Decision, 61(3):205–249.

Hogarth, R. M. and Karelaia, N. (2007). Heuristic and linear models of judgment: Matching rules and environments. Psychological Review, 114:733–758.

Hollnagel, E. (1993). Human reliability analysis: context and control. Academic Press.

Howison, S. D. (1995). Applied mathematics and finance.

Jacoby, J., Speller, D. E., and Kohn, C. A. (1974a). Brand choice behavior as a function of information load. Journal of Marketing Research, 11(1):pp. 63–69.

Jacoby, J., Speller, D. E., and Kohn, C. A. (1974b). Brand choice behavior as a function of information load: Replication and extension. Journal of Consumer Research, 1:33–42.

Jakobi, N., Husbands, P., and Harvey, I. (1995). Noise and the reality gap: The use of simulation in evolutionary robotics, pages 704–720. Springer Berlin Heidelberg, Berlin, Heidelberg.

Jenny, M. (2016). Improving medical decisions with cognitive data science. In Summer Institute on Bounded Rationality 2016.

Jenny, M., Hertwig, R., Ackermann, S., Messmer, A., Karakoumis, J., Nickel, C., and Bingisser, R. (2015). Are mortality and acute morbidity in ppatient prepresent with nonspecific complaints predictable using routine variables? 22:1155–1163.

Jenny, M. A., Pachur, T., Lloyd Williams, S., Becker, E., and Margraf, J. (2013). Simple rules for detecting depression. Journal of Applied Research in Memory and Cognition, 2(3):149–157.

Jones, D., Schonlau, M., and Welch, W. J. (1998). Efficient global optimization of expensive black-box functions. 13:455–492.

Juslin, P. and Persson, M. (2002). Probabilities from exemplars (probex): a "lazy" algorithm for probabilistic inference from generic knowledge. pages 563–607.

Kahneman, D., Slovic, P., and Tversky, A. (1982). Judgment under uncertainty: Heuristics and biases. Cambridge University. Biases/processes of DM (e.g. heuristics).

Kahneman, D. and Tversky, A. (1984). Choices, values, and frames. American Psychologist, 39:341–350.

Karelaia, N. (2006). Thirst for confirmation in multi-attribute choice: Does search for consistency impair decision performance? Organizational Behavior and Human Decision Processes, 100(1):128–143.

Karelaia, N. and Hogarth, R. M. (2008). Determinants of linear judgment: a meta-analysis of lens model studies. Psychological Bulletin, 134(3):404.

Katsikopoulos, K. V. (2010). The less-is-more effect: Predictions and tests. Judgment and Decision Making, 5(4):244 – 257.

Katsikopoulos, K. V. (2011). Psychological heuristics for making inferences: Definition, performance, and the emerging theory and practice. Decision Analysis, 8(1):10–29.

Katsikopoulos, K. V. (2013). Why do simple heuristics perform well in choices with binary attributes? Decision Analysis, 10(4):327–340.

Katsikopoulos, K. V., Egozcue, M., and Garcia, L. F. (2014). Cumulative dominance in multi-attribute choice: benefits and limits. EURO Journal on Decision Processes, pages 1–11.

Katsikopoulos, K. V. and Fasolo, B. (2006). New tools for decision analysts. IEEE Transactions on Systems, Man, and Cybernetics–Part A: Systems and Humans, 36(5):960–967.

Katsikopoulos, K. V. and Martignon, L. (2006). Naive heuristics for paired comparisons: Some results on their relative accuracy. Journal of Mathematical Psychology, 50(5):488–494.

Katsikopoulos, K. V., Pachur, T., Machery, E., and Wallin, A. (2008). From meehl to fast and frugal heuristics (and back): New insights into how to bridge the clinical–actuarial divide. Theory & Psychology, 18(4):443–464.

Katsikopoulos, K. V., Schooler, L. J., and Hertwig, R. (2010). The robust beauty of ordinary information. Psychological Review, 117(4):1259.

Kattah, J. C., Talkad, A. V., Wang, D. Z., Hsieh, Y.-H., and Newman-Toker, D. E. (2009). Hints to diagnose stroke in the acute vestibular syndrome three-step bedside oculomotor examination more sensitive than early mri diffusion-weighted imaging. Stroke, 40(11):3504–3510.

Keller, N., Cokely, E. T., Katsikopoulos, K. V., and Wegwarth, O. (2010). Naturalistic heuristics for decision making. Journal of Cognitive Engineering and Decision Making, 4(3):256–274.

Keller, N., Czienskowski, U., and Feufel, M. A. (2014). Tying up loose ends: a method for constructing and evaluating decision aids that meet blunt and sharp-end goals. Ergonomics, 57(8):1127–1139.

Keller, N. and Katsikopoulos, K. (2016). On the role of psychological heuristics in operational research; and a demonstration in military stability operations. European Journal of Operational Research, 249(3):1063–1073.

Kivetz, R. and Simonson, I. (2000). The effects of incomplete information on consumer choice. Journal of Marketing Research, 37(4):427–448.

Klein, G. (2008). Naturalistic decision making. Human Factors, 50:456–460.

Klein, G. A. (1993). Decision making in action: Models and Methods, chapter A recognition-primed decision (RPD) model of rapid decision making, pages 138–147. Ablex Publishing Corporation.

Klein, G. A. (1998). Sources of power: How people make decisions. MIT press.

Klein, G. A. and Calderwood, R. (1996). Investigations of naturalistic decision making and the recognition-primed decision model. Technical report, U.S. Army Research Institute for Behavioral and Social Sciences.

Klein, G. A., Calderwood, R., and Clinton-Cirocco, A. (1988). Rapid decision making on the fire ground. Technical report, U.S. Army Research Institute for the Behavioral and Social Sciences.

Kleinmuntz, B. (1990). Why we still use our heads instead of formulas: Toward an integrative approach. Psychological Bulletin, 107:296–310.

Kleinmuntz, D. N. and Schkade, D. A. (1993). Information displays and decision processes. Psychological Science, 4(4):221–227.

Kohli, R. and Jedidi, K. (2007). Representation and inference of lexicographic preference models and their variants. Marketing Science, 26(3):380–399.

Lichman, M. (2013). UCI machine learning repository.

Lipshitz, R., Klein, G., Orasanu, J., and Salas, E. (2001). Taking stock of naturalistic decision making. Journal of Behavioral Decision Making, 14(5):331–352.

Lipshitz, R. and Strauss, O. (1997). Coping with uncertainty: A naturalistic decision-making analysis. Organizational Behavior and Human Decision Processes, 69(2):149–163.

Luan, S., Schooler, L. J., and Gigerenzer, G. (2011). A signal-detection analysis of fast-and-frugal trees. Psychological Review, 118(2):316–338.

Luan, S., Schooler, L. J., and Gigerenzer, G. (2014). From perception to preference and on to inference: An approach–avoidance analysis of thresholds. Psychological review, 121(3):501–525.

Ma, L. and Chung, K. (2012). In defense of eponyms. Plastic and Reconstructive Surgery, 129:896–898.

Malhotra, N. K. (1982). Information load and consumer decision making. Journal of Consumer Research, 8(4):pp. 419–430.

Malhotra, N. K., Jain, A. K., and Lagakos, S. W. (1982). The information overload controversy: An alternative viewpoint. Journal of Marketing, 46(2):pp. 27–37.

Marewski, J. N. and Link, D. (2014). Strategy selection: An introduction to the modeling challenge. Wiley Interdisciplinary Reviews - Cognitive Science, 5:39–59.

Martignon, L. and Hoffrage, U. (2002). Fast, frugal, and fit: Simple heuristics for paired comparison. Theory and Decision, 52(1):29–71.

Martignon, L., Katsikopoulos, K. V., and Woike, J. K. (2008). Categorization with limited resources: A family of simple heuristics. Journal of Mathematical Psychology, 52(6):352–361.

Mata, R., Schooler, L. J., and Rieskamp, J. (2007). The aging decision maker: cognitive aging and the adaptive selection of decision strategies. Psychology and aging, 22(4):796.

Maule, A. J. (1994). A componential investigation of the relation between structural modelling and cognitive accounts of human judgement. Acta Psychologica, 87:199–216.

McCammon, I. and Hägeli, P. (2007). An evaluation of rule-based decision tools for travel in avalanche terrain. Cold Regions Science and Technology, 47(1):193–206.

McLeod, P. and Dienes, Z. (1996). Do fielders know where to go to catch the ball or only how to get there? 22:531–543.

Meder, B. and Nelson, J. D. (2012). Information search with situation-specific reward functions. Judgment and Decision Making, 7(2):119–148.

Meehl, P. E. (1954). Clinical versus statistical prediction; a theoretical analysis and a review of the evidence. Minneapolis, University of Minnesota Press.

Mora, B. and Bosch, X. (2010). Medical eponyms: Time for a name change. Archives of Internal Medicine, 170(16):1499–1500.

Mosier, K. L. and Fischer, U. M. (2010). Judgment and decision making by individuals and teams: issues, models, and applications. Reviews of Human factors and Ergonomics, 6(1):198–256.

Nelson, J. D. (2005). Finding useful questions: On bayesian diagnosticity, probability, impact, and information gain. Psychological Review, 112(4):979–999.

Nelson, J. D. (2008). The probabilistic mind: prospects for rational models of cognition, chapter Towards a rational theory of human information acquisition, pages 144–163. Number 7. Charter and Oaksford.

Nelson, J. D., McKenzie, C. R., Cottrell, G. W., and Sejnowski, T. J. (2010). Experience matters: Information acquisition optimizes probability gain. Psychological Science, 21(7):960–969.

Newell, B. R. (2005). Re-visions of rationality? 9(1):11–15.

Newell, B. R. and Shanks, D. R. (2003). Take the best or look at the rest? factors influencing one-reason decision making. Journal of Experimental Psychology: Learning, Memory & Cognition, 29(1):53–65.

Newell, B. R., Weston, N. J., and Shanks, D. R. (2003). Empirical tests of a fast-and-frugal heuristic: Not everyone “takes-the-best”. Organizational Behavior and Human Decision Processes, 91:82–96.

Olsson, A.-C., Enkvist, T., and Juslin, P. (2006). Go with the flow: How to master a nonlinear multiple-cue judgment task. Journal of Experimental Psychology: Learning, Memory, and Cognition, 32(6):1371 – 1384.

Orasanu, J. and Connolly, T. (1993). Decision making in action: Models and methods, chapter The Reinvention of Decision Making, pages 3–20. Ablex Publishing.

Pachur, T. and Olsson, H. (2012). Type of learning task impacts performance and strategy selection in decision making. Cognitive Psychology, 65:207–240.

Park, K. S. (2004). Mathematical programming models for characterizing dominance and potential optimality when multicriteria alternative values and weights are simultaneously incomplete. Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on, 34(5):601–614.

Payne, J., Johnson, E., Bettman, J., and Coupey, E. (1990). Understanding contingent choice: a computer simulation approach. Systems, Man and Cybernetics, IEEE Transactions on, 20(2):296–309.

Payne, J. W., Bettman, J. R., and Johnson, E. J. (1988). Adaptive strategy selection in decision making. Journal of Experimental Psychology: Learning, Memory, and Cognition, 14(3):534–552.

Payne, J. W., Bettman, J. R., and Johnson, E. J. (1993). The Adaptive Decision Maker. Cambridge University Press.

Payne, J. W., Bettman, J. R., and Luce, M. F. (1996). When time is money: Decision behavior under opportunity-cost time pressure. Organizational Behavior and Human Decision Processes, 66(2):131–152.

Prengaman, R. J., Wetzlar, E. C., and Bailey, R. J. (2001). Integrated ship defense. Johns Hopkins APL Technical Digest, 22:523–535.

Raiffa, H. (1968). Decision Analysis: Introductory Lectures on Choice Under Certainty. Addison-Wesley.

Rashid, R. M. and Rashid, R. M. (2007). Medical eponyms: our past, present, and future. International Journal of Dermatology, 46(9):996–996.

Rasmussen, J. (1983). Skill, rules and knowledge:signals, signs, and symbols, and other distinctions in human performance models. IEEE Transactions on Systems, Man and Cybernetics, 13(3):257–266.

Rieskamp, J. (2006a). Perspectives of probabilistic inference: Reinforcement and adaptive network models compared. Journal of Experimental Psychology: Lerning, Memory, and Cognition, 32(6):1355–1370.

Rieskamp, J. (2006b). Positive and negative recency effects in retirement savings decisions. Journal of Experimental Psychology: Applied, 12:233–250.

Rieskamp, J., Busemeyer, J. R., and Laine, T. (2003). How do people learn to allocate resources? comparing two learning theories. Journal of Experimental Psychology: Learning, Memory & Cognition, 29:1066–1081.

Rieskamp, J. and Hoffrage, U. (1999). When do people use simple heuristics, and how can we tell?, chapter When do people use heuristics and how can we tell?, pages 141–167. New York: Oxford University Press.

Rieskamp, J. and Hoffrage, U. (2008). Inferences under time pressure: How opportunity costs affect strategy selection. Acta psychologica, 127(2):258–276.

Ross, William T., J. and Creyer, E. H. (1992). Making inferences about missing information: The effects of existing information. Journal of Consumer Research, 19(1):pp. 14–25.

Ross, K. G., Klein, G. A., Thunholm, P., Schmitt, J. F., and Baxter, H. C. (2004). The recognition-primed decision model. Military Review.

Rothrock, L. (1995). Performance Measures and Outcome Analyses of Dynamic Decision Making in Real-Time Supervisory Control. PhD thesis, Georgia Institute of Technology.

Rothrock, L. and Kirlik, A. (2003). Inferring rule-based strategies in dynamic judgment tasks: toward a noncompensatory formulation of the lens model. Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on, 33(1):58–72.

Russo, J. E. (1974). More information is better: A reevaluation of jacoby, speller and kohn. Journal of Consumer Research, 1(3):pp. 68–72.

Savage, L. J. (1954). The Foundations of Statistics. Wiley, New York, NY.

Schwartz, B. (2003). The Paradox of Choice: Why More Is Less. HarperCollins Publishers Inc., New York, NY.

Shanteau, J. and Thomas, R. (2000). Fast and frugal heuristics: What about unfriendly environments? Behavioral and Brain Sciences, 23:762–763.

Silver, M. S. (1988). Descriptive analysis for computer-based decision support: Special focus article. Operations Research, 36(6):904–916.

Silver, M. S. (1990). Decision support systems: directed and nondirected change. Information Systems Research, 1(1):47–70.

Simmons, C. J. and Lynch, John G., J. (1991). Inference effects without inference making? effects of missing information on discounting and use of presented information. Journal of Consumer Research, 17(4):pp. 477–491.

Simon, H. A. (1955). A behavioral model of rational choice. Quarterly Journal of Economics, 69(1):99 – 118.

Simon, H. A. (1956). Rational choice and the structure of the environment. Psychological review, 63(2):129.

Şimşek, Ö. (2013). Linear decision rule as aspiration for simple decision heuristics. In Advances in Neural Information Processing Systems, pages 2904–2912.

Slegers, D. W., Brake, G. L., and Doherty, M. E. (2000). Probabilistic mental models with continuous predictors. Organizational Behavior and Human Decision Processes, 81(1):98 – 114.

Slovic, P. and MacPhillamy, D. (1974). Dimensional commensurability and cue utilization in comparative judgment. Organizational Behavior and Human Performance, 11(2):172–194.

Stabell, C. B. (1987). Decision support systems: Alternative perspectives and schools. Decision Support Systems, 3(3):243 – 251.

Summers, J. O. (1974). Less information is better? Journal of Marketing Research, 11(4):pp. 467–468.

Tetlock, P. E. (1983). Accountability and the perseverance of first impressions. Social Psychology Quarterly, 46(4):285–292.

Tetlock, P. E. and Kim, J. I. (1987). Accountability and judgment processes in a personality prediction task. Journal of personality and social psychology, 52(4):700–709.

Thaler, R. H. and Sunstein, C. R. (2009). Nudge: Improving Decisions about Health, Wealth, and Happiness. Penguin Books.

Thunholm, P. (2003). Decision making under time pressure: To evaluate or not to evaluate three options before the decision is made?

Todd, P. and Benbasat, I. (1999). Evaluating the impact of DSS, cognitive effort, and incentives on strategy selection. Information Systems Research, 10(4):356 – 374.

Todd, P. M. (2007). How much information do we need? European Journal of Operational Research, 177:1317–1332.

Todd, P. M. and Gigerenzer, G. (2000). Précis of simple heuristics that make us smart. Behavioral and brain sciences, 23(05):727–741.

Todd, P. M., Gigerenzer, G., and the ABC Research Group (2012). Ecological rationality: Intelligence in the world. Oxford University Press.

Tversky, A. (1969). Intransitivity of preferences. Psychological review, 76(1):31.

Tversky, A. and Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. Science, 185:1124–1131.

von Helversen, B., Karlsson, L., Mata, R., and Wilke, A. (2013). Why does cue polarity information provide benefits in inference problems? the role of strategy selection and knowledge of cue importance. Acta Psychologica, 144:73–82.

von Helversen, B. and Rieskamp, J. (2009). Predicting sentencing for low-level crimes: Comparing models of human judgment. Journal of Experimental Psychology: Applied, 15(4):375–395.

Wbben, M. and v. Wangenheim, F. (2008). Instant customer base analysis: Managerial heuristics often ”get it right”. Journal of Marketing, 72(3):82–93.

Wilkie, W. L. (1974). Analysis of effects of information load. Journal of Marketing Research, 11(4):pp. 462–466.

Zeelenberg, M. and Beattie, J. (1997). Consequences of regret aversion 2: Additional evidence for effects of feedback on decision making. Organizational Behavior and Human Decision Processes, 72(1):63–78.

Zeelenberg, M., Beattie, J., van der Plight, J., and de Vries, N. K. (1996). Consequences of regret aversion: Effects of expected feedback on risky decision making. Organizational Behavior and Human Decision Processes, 65(2):148–158.

VITA

Marc Canellas is a Ph.D. candidate in aerospace engineering at Georgia Tech having completed his M.S. in aerospace engineering from Georgia Tech and B.S. in mechanical engineering from the University of Missouri. His work combines the principles and techniques of cognitive engineering, human-automation interaction, and behavioral decision making. His Ph.D. research, supported by the Office of Naval Research, uses mathematical, computational, and human-subjects studies to design decision support tools for command and control in degraded and denied information environments. He has also leveraged this cognitive engineering background to address questions of how to govern human-automation systems, especially autonomous weapons systems. He has published this research in 4 peer-reviewed journals and 4 conference proceedings including *IEEE Transactions on Human-Machine Systems*, the *Journal of Cognitive Engineering and Decision Making*, the *IEEE Technology and Society Magazine*, and *WeRobot*. From 2015-2016, he served as the President of the Graduate Student Body at Georgia Tech, responsible for advocating on behalf of over 9,500 graduate students, and authorizing and allocating a $5.1M budget. For his contributions in research and to his communities, he has earned multiple awards, including the NSF Graduate Research Fellowship, the Best Poster Award at the Summer Institute on Bounded Rationality at the Max Planck Institute for Human Development, the Campus Life and Community at Georgia Tech Scholarship, the Sam Nunn Security Program Fellowship, and the University of Missouri Unsung Hero Award. He currently serves as a member of the IEEE-USE Ad-Hoc Committee on Artificial Intelligence Policy and as the IEEE Standards Association P7000 Working Group - Model Process for Addressing Ethical Concerns During System Design.