

Cognitive and Neural Mechanisms of Adaptive Satisficing Decision Making

by

Hanna Oh

Department of Psychology & Neuroscience
Duke University

Date: _____

Approved:

Tobias Egner, Supervisor

Scott A. Huettel

Jennifer M. Groh

Marc A. Sommer

Dissertation submitted in partial fulfillment of
the requirements for the degree of Doctor
of Philosophy in the Department of
Psychology & Neuroscience in the Graduate School
of Duke University

2017

ABSTRACT

Cognitive and Neural Mechanisms of Adaptive Satisficing Decision Making

by

Hanna Oh

Department of Psychology & Neuroscience
Duke University

Date: _____

Approved:

Tobias Egner, Supervisor

Scott A. Huettel

Jennifer M. Groh

Marc A. Sommer

An abstract of a dissertation submitted in partial
fulfillment of the requirements for the degree
of Doctor of Philosophy in the Department of
Psychology & Neuroscience in the Graduate School of
Duke University

2017

Copyright by
Hanna Oh
2017

Abstract

Much of our real-life decision making is bounded by uncertain information, limitations in cognitive resources, and a lack of time to allocate to the decision process. To mitigate these pressures, people *satisfice*, foregoing a full evaluation of all available evidence to focus on a subset of cues that allow for fast and “good-enough” decisions. Although this form of decision-making likely mediates many of our everyday choices, very little is known about the manner in which satisficing is spontaneously triggered and accomplished. The aim of this dissertation, therefore, is to characterize cognitive and neural mechanisms underlying human satisficing behavior via tasks that closely model real-life challenges in decision making. Specifically, the empirical studies presented here examine how people solve a novel multi-cue probabilistic classification task under various external and internal pressures, using a set of strategy analyses based on variational Bayesian inference, which can track and quantify shifts in strategies. Results from these behavioral and computational approaches are then applied to model human functional magnetic resonance imaging (fMRI) data to investigate neural correlates of satisficing. The findings indicate that the human cognitive apparatus copes with uncertainty and various pressures by adaptively employing a “Drop-the-Worst” heuristic that minimizes cognitive time and effort investment while preserving the consideration of the most diagnostic cue information.

Dedication

To my parents, Oh Chang Hwan and Kim Hae Hyun.

Table of Contents

| | |
|---|------|
| Abstract | iv |
| List of Tables | xii |
| List of Figures | xiii |
| Acknowledgements | xvi |
| 1. Introduction to Satisficing Decision Making..... | 1 |
| 1.1 Decision making in small worlds..... | 3 |
| 1.1.1 Homo economicus, a rational decision maker | 3 |
| 1.1.2 Rational decision making in the real world..... | 4 |
| 1.2 Decision making in large worlds..... | 5 |
| 1.2.1 Bounded rationality: Coping with computational limits | 5 |
| 1.2.2 Satisfice = Satisfy + Suffice..... | 6 |
| 1.3 Classic models of heuristics | 8 |
| 1.3.1 Memory-based heuristics | 8 |
| 1.3.2 Cue-based heuristics | 10 |
| 1.4 Heuristics, good or bad?..... | 13 |
| 1.4.1 Satisficing under uncertainty: Insights from cognitive neuroscience | 16 |
| 1.5 Neural basis of satisficing decision making | 19 |
| 1.5.1 Insights from rational choice models..... | 19 |
| 1.5.2 Neural mechanisms of heuristic decision making..... | 22 |
| 1.6 Towards understanding satisficing decision making | 23 |

| | |
|---|----|
| 2. Multi-Cue Probabilistic Classification Task and Decision Strategy Analysis | 25 |
| 2.1 Multi-cue probabilistic classification task..... | 26 |
| 2.1.1 Stimuli..... | 27 |
| 2.1.2 Task | 28 |
| 2.1.3 Post-experimental survey..... | 31 |
| 2.2 Decision strategy analysis using variational Bayesian inference | 31 |
| 2.2.1 Behavioral performance analysis | 32 |
| 2.2.2 Sum of Evidence (SoE)..... | 33 |
| 2.2.3 Ideal observer model | 33 |
| 2.2.4 Subjective cue weights..... | 34 |
| 2.2.5 Decision strategy models | 37 |
| 2.2.6 Strategy model selection..... | 40 |
| 3. Satisficing Decision Making Under Time Pressure | 41 |
| 3.1 Experiment 1: Satisficing under moderate time pressure..... | 44 |
| 3.1.1 Methods | 44 |
| 3.1.1.1 Participants | 44 |
| 3.1.1.2 Procedure | 45 |
| 3.1.1.3 Survey | 47 |
| 3.1.1.4 Data analysis..... | 47 |
| 3.1.2 Results | 48 |
| 3.1.2.1 Task performance..... | 48 |
| 3.1.2.2 Subjective cue weights | 51 |

| | |
|--|----|
| 3.1.2.3 Decision strategy model comparison and selection..... | 53 |
| 3.1.2.4 Survey results | 58 |
| 3.1.3 Discussion..... | 60 |
| 3.2 Experiment 2: Satisficing under severe time pressure | 62 |
| 3.2.1 Methods | 62 |
| 3.2.1.1 Participants | 62 |
| 3.2.1.2 Procedure | 62 |
| 3.2.2 Results | 63 |
| 3.2.2.1 Task performance..... | 63 |
| 3.2.2.2 Subjective cue weights | 64 |
| 3.2.2.3 Decision strategy model comparison and selection..... | 65 |
| 3.2.2.4 Survey results | 68 |
| 3.2.3 Discussion..... | 69 |
| 3.3 Experiment 3: Satisficing using non-compound cues under time pressure..... | 70 |
| 3.3.1 Methods | 71 |
| 3.3.1.1 Participants | 71 |
| 3.3.1.2 Stimuli..... | 71 |
| 3.3.1.3 Procedure | 72 |
| 3.3.1.4 Survey | 74 |
| 3.3.1.5 Data analysis..... | 75 |
| 3.3.2 Results..... | 75 |
| 3.3.2.1 Task performance..... | 75 |

| | |
|--|-----|
| 3.3.2.2 Subjective cue weights | 76 |
| 3.3.2.3 Decision strategy model comparison and selection..... | 78 |
| 3.3.2.4 Survey results | 81 |
| 3.3.3 Discussion..... | 82 |
| 3.4 General discussion | 84 |
| 4. Satisficing Decision Making Under Anticipatory Anxiety | 96 |
| 4.1 Methods | 99 |
| 4.1.1 Participants..... | 99 |
| 4.1.2 Stimuli..... | 100 |
| 4.1.3 Procedure..... | 100 |
| 4.1.4 Skin conductance level (SCL) analysis | 102 |
| 4.1.5 Participant clustering..... | 102 |
| 4.1.6 Behavioral performance analysis | 103 |
| 4.1.7 Decision strategy model selection..... | 103 |
| 4.1.8 Explicit knowledge of the cue structure..... | 104 |
| 4.2 Results | 106 |
| 4.2.1 Physiological response to threat-of-shock | 106 |
| 4.2.2 Task performance | 108 |
| 4.2.3 Decision strategy model comparison | 109 |
| 4.2.4 Post-experimental survey..... | 112 |
| 4.3 Discussion..... | 115 |
| 5. Neural Mechanisms of Satisficing Decision Making Under Time Pressure..... | 121 |

| | |
|--|-----|
| 5.1 Methods | 123 |
| 5.1.1 Participants..... | 123 |
| 5.1.2 Procedure..... | 123 |
| 5.1.3 Behavioral performance analysis | 125 |
| 5.1.3.1 Sum of Evidence (SoE) | 125 |
| 5.1.3.2 Subjective cue weights | 125 |
| 5.1.3.3 Comparison between decision strategies adopted under low and high time pressure | 125 |
| 5.1.4 fMRI data acquisition..... | 126 |
| 5.1.5 Image preprocessing | 126 |
| 5.1.6 fMRI data analyses | 127 |
| 5.1.6.1 Neural representation of subjective cue weights..... | 128 |
| 5.1.6.2 Modulation of neural activity by SoE | 130 |
| 5.1.6.3 Region-of-Interest analysis of objective SoE | 132 |
| 5.1.6.4 Practice effect analysis..... | 133 |
| 5.2 Results | 134 |
| 5.2.1 Behavioral data | 134 |
| 5.2.1.1 Task performance..... | 134 |
| 5.2.1.2 Subjective cue weights | 136 |
| 5.2.1.3 Decision strategy model selection | 137 |
| 5.2.2 Neuroimaging data | 140 |
| 5.2.2.1 Neural representation of subjective cue weights..... | 140 |

| | | |
|------------|---|-----|
| 5.2.2.2 | Modulation of neural activity by SoE | 144 |
| 5.2.2.3 | Ruling out practice effects..... | 151 |
| 5.3 | Discussion..... | 155 |
| 6. | General Discussion | 162 |
| 6.1 | Adaptive satisficing decision making..... | 163 |
| 6.1.1 | Near-optimal decision making under low pressure..... | 164 |
| 6.1.2 | Satisficing under externally and internally induced pressures | 167 |
| 6.1.3 | Limitations of the current approach | 168 |
| 6.2 | Theoretical and methodological implications | 170 |
| 6.3 | Practical implications | 172 |
| 6.4 | Directions for future research | 175 |
| 6.5 | Conclusion..... | 177 |
| Appendix A | | 178 |
| Appendix B | | 180 |
| References | | 184 |
| Biography | | 206 |

List of Tables

| | |
|--|-----|
| Table 1: Compound Stimulus Organization | 28 |
| Table 2: Cue Weight Assignment | 28 |
| Table 3: Summary of Task Performance. Numbers in parenthesis are SEM. | 51 |
| Table 4: Mean Spearman's Rank Correlation Coefficients Between Ideal Weights, Subjective Weights, and Explicit Survey Ratings. | 60 |
| Table 5: Clusters Showing Significant Decoding Accuracy of Subjective Cue Weights. | 143 |
| Table 6: Activation Clusters for Positive Subjective SoE Modulation. | 150 |
| Table 7: Activation Clusters for Negative Subjective SoE Modulation. | 151 |
| Table 8: Results from 2 (LowP/HighP Phase) x 2 (First/Second Set) Repeated-Measures ANOVAs in Selected Regions of Interest. | 181 |
| Table 9: Activation Clusters for Subjective SoE Modulation Using the First Set of Each Phase. | 183 |

List of Figures

- Figure 1: Compound Stimuli and Multi-Cue Probabilistic Classification Task. (A) Stimulus organization with four cue dimensions. (B) Sixteen compound stimuli used in the task. (C) Schematic of the multi-cue probabilistic task; participants were asked to predict a stimulus that is more likely to win, and received a probabilistic outcome (“win” or “lose”), followed by an intertrial interval (ITI). 30
- Figure 2: Decision Strategy Models. (A) Sixteen strategy models, including Model *Opt*, an optimal cue-integration model, and Model TTB, the Take-the-Best heuristic model. Filled circles denote the cues that are included in a given model. (B) Expected accuracy per strategy model. Dotted lines separate strategy models according to the number of cues used. 39
- Figure 3: Task and Performance of Experiments 1 and 2. (A) Probabilistic classification task under time pressure; participants performed the task under varying response windows of 2 s (learning and LowP phases), 750 ms (HighP phase in Experiment 1), and 500 ms (HighP phase in Experiment 2). (B) Average percent optimal responses in Experiment 1 (750 ms time pressure) and Experiment 2 (500 ms time pressure) during the learning, LowP, and HighP phases. Error bars are standard errors (SEM). 50
- Figure 4: Logistic Regression Fit of Individual Performance and Estimation of Subjective Cue Weights. (A) Performance of a single example participant in the LowP phase (2 s response window). (B) Performance of the same participant in the HighP phase (750 ms response window). An individual data point reflects percentage of left choices made given the sum of evidence. Curves are logistic regression fits to the data. (C) Average subjective cue weights for Experiment 1 (750 ms HighP). (D) Average subjective cue weights for Experiment 2 (500 ms HighP). Error bars are SEM. 53
- Figure 5: Bayesian Strategy Model Selection. (A) Experiment 1 (750 ms HighP) using uninformative priors. (B) Experiment 1 using informative priors. (C) Experiment 2 (500 ms HighP) using uninformative priors. (D) Experiment 2 using informative priors. Model 15 represents the optimal cue integration model (Model *Opt*) and Model 16 represents the Take-the-Best (TTB) model. 55
- Figure 6: Cue Organization and Task of Experiment 3. (A) Non-compound stimulus organization using four independent cue dimensions (weather, transportation, activity, and building), each with binary states of varying weights. (B) Schematic of the non-compound multi-cue probabilistic classification task. 73

Figure 7: Task Performance, Subjective Cue Weights, and Bayesian Strategy Model Selection of Experiment 3. (A) Average percent optimal responses in the learning, LowP (3s response window) and HighP (750 ms response window) phases. (B) Average subjective cue weights as a function of LowP and HighP phases. Error bars are SEM. (C) Strategy model selection group results using informative priors. (D) Strategy model selection group results using informative priors. 77

Figure 8: Schematic of the Experimental Procedure. (A) Schematic of the standard multi-cue probabilistic classification task in the initial learning and neutral phases. (B) Schematic of the threat-of-shock phase; participants additionally experienced randomly occurring shocks at the time of feedback along with a shock symbol..... 101

Figure 9: Participant Clustering and Decision Performance. (A) Average skin conductance level (range-corrected) per set in each experimental phase. (B) Participant clustering based on changes in SCL (ΔSCL). Dots represent individual participants. (C) Percent correct choices per group. (D) Average response time (ms) per group. Error bars indicate SEM. 108

Figure 10: Strategy Model Selection per Responder Group. (A) Protected exceedance probabilities for the low responder group. (B) Protected exceedance probabilities for the mid responder group. (C) Protected exceedance probabilities for the high responder group..... 111

Figure 11: The effect of anticipatory anxiety on task knowledge. Solid line represents a significant quadratic fit ($p = 0.011$) between changes in SCL (ΔSCL) and the implicit-explicit strategy concordance (Δr_{sw}). Colored dots indicate individual participant per group. Dashed lines represent the 95% CI for the regression line. 114

Figure 12: Behavioral Results. (A) Performance throughout the task runs. Plotted is the percentage of correct choices favored by the sum of cue weights, regardless of outcome feedback. (B) Percentage of correct choices as a function of objective SoE. (C) Response time (ms) as a function of SoE. Error bars and shaded area represent SEM. 136

Figure 13: Subjective Cue Weights and Strategy Model Selection. (A) Average subjective cue weights as a function of objective, pre-assigned cue weights. (B) Percentage of correct choices when only the least important cue is different between the stimulus pair. Error bars indicate SEM; * $p < 0.01$. (C) Bayesian model selection group results presented in exceedance probabilities under low and high time pressure conditions. (D) Decision strategy models used in model comparison. Filled circles denote the cue dimensions that are included in a given model..... 139

Figure 14: Neural Representation of Subjective Cue Weights. (A) Brain regions significantly predictive of subjective cue weights under low time pressure. (B) Brain regions significantly predictive of subjective cue weights under high time pressure. All maps are $p < 0.05$ corrected. Image display according to neurological convention (left is left). (C) Average prediction accuracy (Fisher's Z-transformed correlation coefficients) of each cluster. Accuracy of zero represents at chance decoding performance. dStr = dorsal striatum; PG = postcentral gyrus; Error bars indicate SEM; * $p < 0.05$ 143

Figure 15: Modulation of Neural Activity by SoE. (A) Significant clusters under LowP. (B) Significant clusters under HighP. Positive modulation (red) indicates activity increasing with accumulation of evidence, whereas negative modulation (blue) represents activity increasing with decreasing SoE, and hence, greater uncertainty or decision difficulty. (C) Brain areas demonstrating significantly greater positive SoE modulation in the HighP compared to the LowP phase. STG = superior temporal gyrus; MTG = middle temporal gyrus; LG = lingual gyrus. All maps were whole-brain corrected to $p < 0.05$. (D) Mean activity as a function of objective SoE in the putamen. (E) Mean activity as a function of objective SoE in the vmPFC. Objective SoE of 1.6* includes trials with SoE levels ranging from 1.6 to 2. Shaded area indicates SEM. 149

Figure 16: Changes in Trial-Induced BOLD Activity over Time in the dlPFC, IPL, Putamen, and Thalamus/Midbrain ROIs. (A) Average percent signal changes. (B) Average parameter estimates, θ_{Trial} . A full statistical report is included in Table 8. Numbers on x-axis corresponds to the set number within the low (L) and high (H) time pressure phases. Error bars indicate SEM. 180

Figure 17: Strategy Model Selection and fMRI Data Analyses per Set. (A) Exceedance probabilities of each strategy model per set. (B) Brain regions significantly modulated by subjective SoE under the low time pressure (LowP). (C) Brain regions significantly modulated by subjective SoE under the high time pressure (HighP). (D) Areas demonstrating significantly greater positive SoE modulation in the HighP compared to the LowP phase. All neuroimaging analyses were performed using only the first set of each phase. All maps were whole-brain corrected to $p < 0.05$ 182

Acknowledgements

I am especially grateful to my advisor, Tobias Egner, for his brilliant mentorship, constant support, and generosity. I truly enjoyed working with him and he indeed deserves the title, “Best Boss in the World”. I would also like to thank my collaborators, Jeff Beck, Marc Sommer, and Silvia Ferrari, who enabled the satisficing project. I also wish to acknowledge my committee members, Scott Huettel and Jenni Groh, for their helpful insights. I owe a great deal of gratitude to Randolph Blake, who did not let me give up. I also thank the current and former Egner lab members, who shared invaluable wisdom and entertainment. Thank you to my awesome family and friends for their love, support, and encouragement, especially to Nik Descher, for always being there for me.

1. Introduction to Satisficing Decision Making

Much of our real-life decision making involves combining multiple pieces of information, each associated with some degrees of uncertainty in predicting an outcome. The process of identifying, evaluating, and integrating decision-relevant information is not a computationally simple problem, yet in our daily lives, we carry it out seamlessly, and often quite effectively, without investing much time or cognitive effort. Hence, understanding how humans make choices in a complex real-world environment has been a central issue in a wide field of research, which have led to diverse theories of cognitive processes underlying decision making. Specifically, many prominent choice models that are widely used in decision neuroscience assume humans as rational agents, who select an option that maximizes the outcome (e.g., Kahneman & Tversky, 1979; Von Neumann & Morgenstern, 1944). Achieving such optimality, however, often requires exhaustive computations based on a complete knowledge of all possible choices, their consequences, and probabilities (Simon, 1990). This basic assumption, therefore, could lead to failures when applied to real world situations, in which some of the relevant information is unknown or highly uncertain.

Consider following cases: 1) Playing a game of single-deck Blackjack and 2) purchasing a house. Both cases require making a series of choices but the bases for arriving at decisions are quite different. In a game of Blackjack, a player can make fairly accurate inferences about the cards that are remaining to be dealt based on careful

observations of the exposed cards, which then can be used to adjust betting strategies to maximize winnings. This so-called card counting is possible because the player has a complete knowledge of the cards that comprise a deck, and probabilities of their occurrences. In case of purchasing a house, however, decision making is not as simple as card counting. A buyer has to examine and compare various aspects of houses such as location, condition, price, and possible appreciations in the future, some of which can be highly uncertain or even unattainable. In addition, even if the buyer has obtained sufficient information, it becomes computationally very expensive (or impossible) to take all available information, possible variability, and probabilities of occurrence into account to arrive at an optimal choice.

The failure of rational choice theory in characterizing decision behavior in the real world initiated the search for decision rules that are simpler, less rational, and hence, lead to solutions that are not necessarily optimal but “good-enough”, a process which is referred to as *satisficing* (Simon, 1955). This line of research has established that when faced with such complex problems (e.g., purchasing a house), people readily adopt heuristics, a set of strategies that ignore part of the available information to produce fast and close-to-optimal decisions, while preserving cognitive effort (Gigerenzer & Gaissmaier, 2011). Even though this type of bounded decision making likely underpins most decisions we make in everyday life, a large literature in cognitive neuroscience has focused on applying rational models to make inferences on underlying

processes (for a discussion, see Volz & Gigerenzer, 2014). This approach limits our understanding of cognitive and neural mechanisms associated with real-world decision making, in which rationality is bounded by uncertain information, limitations in cognitive resources, and a lack of time to allocate to the decision process. The overall goal of this dissertation, therefore, is to characterize cognitive and neural mechanisms underlying human satisficing behavior under various choice scenarios that closely model real-life challenges in decision making.

1.1 Decision making in small worlds

1.1.1 Homo economicus, a rational decision maker

Perhaps one of the most influential contributions to current decision science is Von Neumann & Morgenstern (1947)'s Expected Utility (EU) Theory. Building on from Bernoulli (1738/1954)'s notion of *utility* – an individual's preferences or *subjective values* over a set of outcomes– the EU theory provides a set of behavioral rules (axioms) that delineate necessary and sufficient conditions that a rational decision maker should obey. Hence, the most important assumption underlying the EU theory is that there *exists* a rational decision maker, the homo economicus or an economic human (Edwards, 1954; Persky, 1995). Homo economicus is rational, completely informed about all the available courses of action and their outcome probabilities, and therefore makes choices so as to maximize (expected) utility. Von Neumann & Morgenstern (1947) proved that the

preferences of a rational economic agent has a monotonic and concave utility function, which emphasizes internal *consistency* of preferences.

Although the EU theory provides a simple mathematical tool that can be used to infer hidden subjective values in a decision process, there is substantial evidence indicating that most people violate these axioms (e.g., Allais, 1953; Kahneman & Tversky, 1979). To account for experimentally observed behavioral inconsistencies, Kahneman & Tversky (1979) later expanded the EU theory by making the valuation process reference-dependent. Regardless of these differences, these utility-based decision theories share the same core conclusion: A decision maker assigns subjective values to the available options, weighs them with the relevant probability functions, and chooses the option with the maximum value.

1.1.2 Rational decision making in the real world

The study of rational decision-making has focused on decision strategies that maximize utility, a process which involves exhaustive computations based on perfect knowledge of decision-relevant information, possible choices, and their outcome probabilities and consequences, referred to as *substantive* (or *unbounded*) *rationality* (Simon, 1955, 1990). Savage (1954) described this as a *small world* scenario, where the future is completely predictable and hence, the optimal solution to a decision problem can be determined from statistical methods. In other words, decisions in small worlds are made under *risk* (Knight, 1921), in which a rational economic agent can take the best

available action based on accurate knowledge of a distribution underlying possible outcomes for all choice alternatives.

Most real-world decision-making, however, entails dealing with varying degrees of ambiguity in information. For instance, compare the bases for making a decision in a game of Blackjack with those involved in deciding where to buy a house, which college to attend, or whom to marry. To address the intrinsic differences underlying various choice problems, Knight (1921) proposed a conceptual distinction between decisions under risk and under uncertainty. Contrary to risk, *uncertainty* refers to situations where the probabilities or likelihood of different outcomes cannot be specified with any mathematical precision, making decision problems computationally intractable (for additional details, see Platt & Huettel, 2008; Tobler & Weber, 2013; Volz & Gigerenzer, 2012). In such *large worlds*, where some decision-relevant information is unknown or incomplete, and the future is uncertain, rational models often fail to provide the “correct” answer (Savage, 1954). Therefore, optimal strategies based on small-world assumption can lead to failures or even disasters when applied to large-world situations.

1.2 Decision making in large worlds

1.2.1 Bounded rationality: Coping with computational limits

In large worlds, human decision-making ability is constrained not only by incomplete knowledge of the decision-relevant information but also by limitations placed upon cognitive resources and available computation time (Simon, 1955, 1956,

1990). Specifically, Simon (1990) conjectured that computational limits of the human brain (i.e., computing speed and power) make an exact estimation of almost any real-world problem impossible, which in turn renders the optimizing strategy rarely feasible. To deal with large-world situations, therefore, we must find strategies that can approximately solve the problem. Consistent with this notion, recent neuroscience findings suggest that human suboptimal choice behavior could be explained by the efficient coding scheme of the brain (Barlow, 1961), which enables most decision diagnostic (or expected) information to be processed with higher gain than low-valued (or unexpected) inputs (for a review, see Summerfield & Tsetsos, 2015). Efficient coding, therefore, maximizes the information encoding and transmission within the limits of a neuronal system, although the mechanism may not lead to optimal choices, especially within a volatile task environment (e.g., Fischer & Whitney, 2014; Michael, de Gardelle, Nevado-Holgado, & Summerfield, 2015). Simon (1955, 1990) explained such suboptimality in human choice behavior through the concept of *bounded rationality*. That is, our rationality is bounded by the computational capabilities of the brain and the structure of decision environment. To efficiently deal with real-world problems, therefore, we oftentimes have to rely on identifying approximate and feasible solutions.

1.2.2 Satisfice = Satisfy + Suffice

To overcome the aforementioned limitations, humans are thought to *satisfice* through the use of heuristic strategies that simplify decision-making problems by

prioritizing some sources of information (cues) while ignoring others (Simon, 1955, 1956, 1990). This approach leads to solutions that are not precisely optimal but *suffice to satisfy* some specified criterion level, thereby facilitating fast and “good-enough” (i.e., approximate) decision-making. Specifically, a *heuristic* refers to a strategy that forgoes part of the available information to enable fast, frugal, and/or near-optimal decision making (Gigerenzer & Gaissmaier, 2011; Shah & Oppenheimer, 2008). Performance of satisficing decision making, therefore, crucially depends on *adaptive* exploitation of the environmental structure (Gigerenzer & Todd, 1999; Simon, 1990). That is, the key to achieving good decision accuracy through heuristics lies with utilizing prior knowledge of the information structure to identify the most predictive set of cues to consider. Indeed, in an environment with high information redundancy (Dieckmann & Rieskamp, 2007), people tend to satisfice by focusing on a subset of the most informative cues to guide their choices (see Section 1.3 for additional details). The same heuristic, however, would not be as effective, if applied to a different (i.e., non-redundant cues) or a volatile (i.e., changing cue values) task environment.

In sum, decades of theoretical and experimental studies have demonstrated that, faced with complex large world problems, people satisfice to cope with bounded rationality. Specifically, adaptive employment of heuristics can often yield comparable decision performance to maximization approaches, while effectively preserving time and cognitive effort (see Section 1.4). Although the applicability of a given heuristic is

highly specific to the decision context, common properties of satisficing strategies that people adopt in various large-world situations have emerged, some of which have been formalized into the models of heuristic decision behavior.

1.3 Classic models of heuristics

The main goal of heuristics is to reduce the effort associated with decision making by processing less information, thereby carrying out less computations than necessary for maximization strategies (e.g., multivariate regression, Bayesian inference), without sacrificing too much accuracy. Accordingly, Shah & Oppenheimer (2008) proposed that heuristics rely on one or more of the following to reduce effort: 1) examining fewer cues, 2) reducing the difficulty of retrieving and storing cue information, 3) simplifying the cue weighting principles, 4) integrating less cues, and 5) considering fewer alternatives. Studies have shown that the application of these principles are robust and reliable, especially under various task pressures such as time (e.g., Payne, Bettman, & Johnson, 1988; Rieskamp & Hoffrage, 2008), information cost (e.g., Newell & Shanks, 2003), and negative affect (e.g., Pachur, Hertwig, & Wolkewitz, 2014). In this section, two prominent concepts of heuristics and corresponding examples will be discussed.

1.3.1 Memory-based heuristics

The most basic heuristic shortcut to decision making is memory: when faced with familiar decision problems or choices, people often make memory-based decisions

using strategies such as the recognition heuristic (Goldstein & Gigerenzer, 2002; see also Newell & Shanks, 2004; Pachur, Todd, Gigerenzer, Schooler, & Goldstein, 2011; Pohl, 2006), the fluency heuristic (Jacoby & Brooks, 1984; Schooler & Hertwig, 2005; Whittlesea, 1993), and the exemplar-based approach (Juslin, Olsson, & Olsson, 2003; Juslin & Persson, 2002; Nosofsky & Palmeri, 1997). For instance, in case of the recognition/fluency heuristics, the heuristics for a two-alternative choice task with options A and B can be defined as follows: if A is recognized and B is novel (recognition) or if both A and B are recognized but A is recognized faster (fluency), then the decision is made based on A (Gigerenzer & Gaissmaier, 2011; Schooler & Hertwig, 2005).

Although the use of memory-based heuristics would not always yield correct inferences, this type of intuitive judgment can be advantageous when there is a strong correlation between recognition and decision criterion. For example, when U.S. students were asked to pick a city with the larger population from a pair of German cities, the mean proportion of choices in accordance with the recognition heuristic was 90% (Goldstein & Gigerenzer, 2002). When this question was reformulated to select a Swiss city that is close to the geographical center of Switzerland, however, only 54% of subjects chose the recognized city (Newell & Shanks, 2004; Pohl, 2006), suggesting that people are able to determine when the recognition cue is useful and adaptively employ the heuristic.

Another advantage of adopting recognition heuristics is to reduce the set size and form a “consideration set” comprised of recognized options, especially when

presented with a large set of alternatives to consider (Marewski, Gaissmaier, Schooler, Goldstein, & Gigerenzer, 2010). This phenomenon is well represented in consumer's choice behavior. For instance, when asked to choose a peanut butter from a set of three different brands, one of which is a heavily advertised national brand and the others are not, over 90% of participants selected the familiar brand despite of the mediocre quality of the product (Hoyer & Brown, 1990; Macdonald & Sharp, 2000). A similar strategy is applied when faced with a novel decision problem constructed with recognizable cues. Instead of examining and combining individual pieces of new information, the exemplar-based heuristic model proposes that people retrieve exemplars (i.e., patterns of cue combinations) from their memory, and make a choice by analogy (Juslin et al., 2003; Kruschke, 1992; Nosofsky & Palmeri, 1997). Hence, the judgement is made based on similarity of the set of new information to an old exemplar and its associated choice criterion. Taken together, these findings show that, if deemed applicable, people readily employ memory-based heuristics to simplify decision problems rather than trying to evaluate all possible alternatives to optimize the decision outcome.

1.3.2 Cue-based heuristics

When dealing with unfamiliar situations with a large space of possibilities to explore, the search for solutions can become very selective based on various rules of thumb or heuristics (Simon, 1990). In this case, people appear to take into account the informational structure of the decision problem, such as the relative values and inter-

correlations between different sources of information and the cost of acquiring new information, in adopting decision heuristics (e.g., Bröder, 2000, 2003; Dieckmann & Rieskamp, 2007; Newell & Shanks, 2003; Rieskamp & Otto, 2006). One optimal way of solving this type of decision problem is to identify all possible alternatives and relevant pieces of information, assess the importance (weight) of each cue to generate an overall value or utility for each alternative, and select the one with the highest value (e.g., weighted additive rule; see Shah & Oppenheimer, 2008). This compensatory decision process involves great computational demand, which in some cases, cannot be appropriately handled given the constraints imposed by the task environment and available cognitive resources. One way to simplify such decision problem is to evaluate one piece of information at a time, starting from what is considered the most important to less important cues, based on a subjective rank order (i.e., lexicographic heuristic; Fishburn, 1974). In case of a binary choice problem, one well-known strategy is the "Take-the-Best" heuristic, which sequentially searches through cues in descending order of their assumed values and stops upon finding the first (highest-ranked) cue that discriminates between the two alternatives (Gigerenzer & Goldstein, 1996). A choice is made without further evaluating less valuable available cues, thereby simplifying the decision problem. Therefore, in contrast to compensatory strategies that weight and integrate all decision-relevant information, Take-the-Best is noncompensatory because a

deterministic cue cannot be outweighed by any combination of less valuable cues (Gigerenzer & Goldstein, 1996).

Laboratory studies that have documented participants' spontaneous use of this Take-the-Best strategy assessed the manner in which subjects combine the different cues as a function of satisficing pressures such as redundant cue information, information costs, and time pressure (e.g., Bröder, 2000, 2003; Dieckmann & Rieskamp, 2007; Newell, Weston, & Shanks, 2003; Payne et al., 1988; Payne, Bettman, & Luce, 1996; Rieskamp & Hoffrage, 2008; Rieskamp & Otto, 2006; Rieskamp, 2006). Many of these studies have shown that under low satisficing pressure, a strategy model that integrates some or all of available information performed best at predicting participants' choices, whereas under high pressure, people adapted by using a simple lexicographic heuristic or Take-the-Best strategy, which looks for "one clever cue" to base decisions on. For example, one study asked participants to predict an outcome (e.g., a university with the highest number of future publications) based on six cues with varying validities (Rieskamp & Hoffrage, 2008). Under low time pressure, the weighted additive model which combines all six cues yielded the highest accuracy in predicting participants' choices. By contrast, under high time pressure, the lexicographic heuristic was the best fitting model. These results demonstrate that, although people have a general tendency to maximize under low satisficing pressures and computational demand, they adaptively switch strategies as characteristics of the decision environment change.

In summary, evidence from both memory- and cue-based heuristics research suggests that people are quite adaptive in exploiting the task structure and employing appropriate heuristics, especially when the optimal strategy is not available (e.g., lack of sufficient information) or not applicable (e.g., limited time or computational power). The effectiveness of satisficing decision making, however, remains controversial as it is often considered that heuristics, for the cost of yielding fast and frugal decision processes, would result in worse outcomes than maximization strategies. In the next section, I will examine this common belief and assess situations in which heuristics could thrive. In addition, I will discuss some shortcomings of existing laboratory approaches, and propose an alternative method inspired by experimental paradigms from cognitive neuroscience research.

1.4 Heuristics, good or bad?

Studies on satisficing and heuristics take a fundamentally different stance from utility-based models of economic decision making in interpreting human behavior and its underlying cognitive processes. Rather than focusing on a small-world scenario, in which utility maximization is deemed possible, they consider large-world situations, when conditions for rational models are difficult to meet (see Gigerenzer & Gaissmaier, 2011). In such scenarios, as reviewed above, people readily apply simple heuristics to produce fast and good-enough choices without investing much effort, a process often

considered as effort and accuracy trade-off¹ (Payne et al., 1988; Shah & Oppenheimer, 2008). This conclusion could imply that, within a large-world context, human choice behavior is inherently irrational and therefore, produces suboptimal results. Payne et al. (1988), on the other hand, proposed that this trade-off could achieve beneficial results, if applied to situations where the gain from achieving the best outcome is not worth the costs of effort. This raises the question whether heuristics can ever outperform maximization strategies.

The short answer is, yes: simple heuristics, if applied appropriately, can yield comparable and sometimes even better performance compared to more complex maximization methods (e.g., Czerlinski, Gigerenzer, & Goldstein, 1999; Martignon, Katsikopoulos, & Woike, 2008). This observation, however, is heavily dependent on the selection of a heuristic that can maximally exploit the environmental structure (Simon, 1990). That is, heuristics can succeed only if employed *adaptively* according to the structure of decision-relevant information, accounting for the correlation between cues, distribution of cue weights, and outcome uncertainty. For instance, a noncompensatory strategy such as Take-the-Best can yield optimal results if applied to a noncompensatory decision problem (e.g., skewed cue weight distributions), provided that a decision

¹ This concept is similar to the speed-accuracy trade-off, often characterized based on an abstract race model (e.g., drift diffusion model; Ratcliff, 1978). Here, evidence is independently accumulated in two competing integrators, and the choice is made when the activity of any integrator reaches a specified threshold (for reviews, see Bogacz, 2014; Gold & Shadlen, 2007). Speeded decisions, therefore, require a conservative shift of the decision or response threshold, which in turn decreases accuracy.

maker has strong prior knowledge of the cue organization. Hence, a given heuristic may not generalize well to tackle a new task with different information distributions and a decision criterion. To achieve good performance via satisficing, therefore, a decision maker must be able to adaptively switch strategies in response to changing environment. Thus, to understand mechanisms underlying human bounded rationality, it is critical to examine how satisficing decision making is triggered and how the application of heuristics change according to dynamic choice conditions.

While laboratory studies of heuristics have produced valuable insights into adaptive shifts in decision strategies (see Section 1.3), many have typically provided subjects with a set of cues of varying, explicitly known values (e.g., Bröder, 2003; Dieckmann & Rieskamp, 2007; Newell, Weston, & Shanks, 2003; Payne, Bettman, & Johnson, 1988; Rieskamp & Otto, 2006; Rieskamp, 2006). In addition, a majority of these studies adopted a serial information search paradigm (Payne et al., 1988), which enables easy tracking of choice patterns but also constrains behavior by hindering quick comparisons of multiple pieces of information (Glöckner & Betsch, 2008). Studies that did not employ this type of paradigm have instead often adopted an outcome-oriented approach, in which certain choices are mapped on to certain decision models (e.g., Bergert & Nosofsky, 2007; Lee & Cummins, 2004). This provides a convenient way to infer strategies but limits the number of trial types that could be used as probes. Additionally, to facilitate the learning of multiple cue values, many studies have

provided participants with trial-by-trial deterministic feedback indicating whether they made a correct or incorrect judgment (e.g., Bergert & Nosofsky, 2007; Juslin et al., 2003; Lamberts, 1995; Lee & Cummins, 2004; Pachur & Olsson, 2012; Rieskamp & Otto, 2006). These paradigms, therefore, arguably fall short of simulating satisficing in the real world, where we are often exposed to multiple sources of information simultaneously and must infer their approximate values from experience. For example, in deciding whether and when to merge into an exit lane on a busy highway, we must estimate our own speed, the distance to the exit ramp, and the relative distances and speeds of cars traveling in front of us, behind us, and in the lane we would like to merge into, as well as their drivers' intentions (e.g., indicators). Estimates of all of these cues are uncertain, and their relative importance for avoiding an accident is inferred from previous experience entailing probabilistic feedback (e.g., we do not get in an accident every time we ignore the rear mirror).

1.4.1 Satisficing under uncertainty: Insights from cognitive neuroscience

One promising approach to examine how people acquire knowledge of the probabilistic nature of a decision environment involving multiple cues is through probabilistic classification tasks, which have been widely used in cognitive neuroscience to study learning, memory, and decision making. One famous example is the weather prediction task (Knowlton, Mangels, & Squire, 1996; Knowlton, Squire, & Gluck, 1994). In this task, participants are shown one, two, or three cards from a set of four "cue

cards" (containing sets of abstract symbols), each of which predicts outcome of rain or sunshine with a given probability that is initially unknown to the participants. On each trial, participants are asked to predict the weather based on a specific combination of cards and receive feedback, which enables them to gradually learn the probabilistic cue-outcome associations. Performance is typically evaluated based on the proportion of "correct" (optimal) choices, which are defined as selecting an outcome that is favored by the probability combination of cue cards, independent of actual probabilistic feedback. Numerous studies have confirmed that healthy participants can indeed learn the task through trial-and-error, achieving choice accuracies that are not significantly different from that of an ideal observer toward the end of the experiment (e.g., Gluck, Shohamy, & Myers, 2002; Lagnado, Newell, Kahan, & Shanks, 2006; Meeter, Radics, Myers, Gluck, & Hopkins, 2008; Shohamy, Myers, Onlaor, & Gluck, 2004). These results, thus, suggest that people are able to utilize probabilistic information to achieve near-optimality through repeated experience.

More importantly, several studies have assessed the decision strategies that subjects adopt in solving the classic weather prediction task, and documented possible shifts in strategies over the course of learning (e.g., Gluck et al., 2002; Lagnado et al., 2006; Meeter, Myers, Shohamy, Hopkins, & Gluck, 2006; Meeter et al., 2008; Speekenbrink, Lagnado, Wilkinson, Jahanshahi, & Shanks, 2010). For instance, Gluck et al. (2002) reported that the majority of participants switched from simple, sub-optimal

strategies to more complex maximization strategy toward the end of training. However, due to a small number of possible combinations of cues (e.g., total 14 combinations in the weather prediction task; see Gluck et al., 2002) and/or extensive training of subjects (e.g., Gould, Nobre, Wyart, & Rushworth, 2012; Wheeler et al., 2015; Yang & Shadlen, 2007), probabilistic classification tasks are potentially susceptible to memorization of specific patterns (see Ma & Jazayeri, 2014). That is, once participants are aware of all possible patterns and their likely outcome associations, it is possible that they could rely on memory-based heuristics to carry out the task.

Taken together, the classic probabilistic classification task provides an ideal template to examine learning and decision making under uncertainty. However, the task is potentially susceptible to memory-guided choice processes and hence, upon sufficient learning of cue-outcome relationships, participants may no longer need to actively combine cue information to make an accurate prediction. Therefore, to better model decision-making in large worlds, an innovative approach, which could encourage the adoption of cue-based strategies in a controlled laboratory setting, is necessary. Such a new approach may significantly improve our understanding of cognitive mechanisms underlying satisficing, which would further contribute to elucidating neuro-computational mechanisms enabling real-world decision making.

1.5 Neural basis of satisficing decision making

The traditional utility-based economic decision theories established a basis for explaining human choice behavior assuming that decision makers behaves *as if* they are seeking rationally to maximize expected outcomes (Friedman, 1953).

Neurophysiological studies of decision making have taken these theories one step further to elucidate neural mechanisms by which decision-relevant information is integrated and transformed into choices (see Kable & Glimcher, 2009). Hence, much of initial research focused on risky choice scenarios, which laid the groundwork for understanding neural correlates of decision-making. In this section, I will introduce some of the core concepts and major discoveries in value-based decision neuroscience and examine whether these findings could be extended to heuristic decisions making.

1.5.1 Insights from rational choice models

Studies of value-based decision-making have advanced in large part from the discovery of dopamine neurons in the midbrain whose firing rates appear to encode reward information (Schultz, Dayan, & Montague, 1997). This signal, referred to as a *reward prediction error*, contains information contrasting actual versus expected reward, and thus, is critical for learning by trial-and-error (i.e., reinforcement learning) (Montague, Dayan, & Sejnowski, 1996; for a review, see Bromberg-Martin, Matsumoto, & Hikosaka, 2010). Monkey electrophysiology studies have shown that phasic activation of dopamine neurons is closely correlated with a decision variable that predicts the

magnitude, probability, and timing of reward (Fiorillo, Newsome, & Schultz, 2008; Fiorillo, Tobler, & Schultz, 2003; Morris, Arkadir, Nevet, Vaadia, & Bergman, 2004; Tobler, Fiorillo, & Schultz, 2005). Similarly, human functional magnetic resonance imaging (fMRI) studies have revealed activations in the striatum that are predictive of rewards as well as subjective values assigned to choices (e.g., Knutson, Adams, Fong, & Hommer, 2001; I. Levy, Snell, Nelson, Rustichini, & Glimcher, 2010; McClure, Berns, & Montague, 2003; Tobler, O'Doherty, Dolan, & Schultz, 2007). For example, blood-oxygen-level dependent (BOLD) responses in the ventral striatum were significantly correlated with the anticipation of rewards (Knutson et al., 2001), and the activation of this region covaried with an expected value, independent of various combinations of magnitude and probability (Tobler et al., 2007).

Another key cortical region that has been found to consistently encode subjective value is the ventromedial prefrontal cortex (vmPFC) (for a review, see Kable & Glimcher, 2009). Recent meta-analysis studies demonstrated that both the striatum and vmPFC are reliably involved in encoding subjective values and individual choice preferences in various decision-making tasks (Bartra, McGuire, & Kable, 2013; Clithero & Rangel, 2013). Importantly, orbitofrontal neurons encode subjective values reflective of combined choice-relevant information such as quantity and probability of reward (Raghuraman & Padoa-Schioppa, 2014). Consistent with these findings, numerous human fMRI studies have shown involvement of the vmPFC in value computation (see

Bartra et al., 2013; Clithero & Rangel, 2013), which is independent from value-driven attentional capture (Grueschow, Polania, Hare, & Ruff, 2015).

Furthermore, studies based on probabilistic classification tasks offer additional insights into neural correlates underlying statistical learning and evidence integration. That is, the early studies of the weather prediction task based on amnesic and Parkinson's disease patients have revealed that amnesic patients with damage to the medial temporal lobe were able to learn the probabilistic cue-outcome associations, whereas Parkinson's disease patients with damage to the basal ganglia failed to learn, demonstrating the critical role of the basal ganglia in acquiring probabilistic associations (Knowlton et al., 1996, 1994). Extending these findings, several neuroimaging studies on healthy participants demonstrated that probabilistic classification learning is primarily mediated by interactions between the striatum and hippocampus (Poldrack et al., 2001; Poldrack & Rodriguez, 2004). In addition, electrophysiology recordings of the monkeys have shown that the firing rates of parietal neurons tracked cumulative probabilistic evidence (Yang & Shadlen, 2007) and reached a stereotyped decision threshold prior to choices (Kira, Yang, & Shadlen, 2015). Other recent studies additionally emphasize an important role of the vmPFC (Akaiishi, Kolling, Brown, & Rushworth, 2016) and the frontoparietal attentional control network (Niv et al., 2015) in mediating goal-directed learning in a multi-cue decision environment.

Taken together, neural evidence from value-based and probabilistic decision making research established a crucial involvement of both subcortical (striatum, hippocampus) and cortical (prefrontal and parietal) regions in encoding and integrating multiple dimensions of reward information. Many of these studies, however, were conducted based on an assumption that all subjects employ a uniform strategy throughout the task, and therefore, offer limited insights into the role of these brain regions in association with strategy use, especially when participants shift from adopting optimal to satisficing decision strategies.

1.5.2 Neural mechanisms of heuristic decision making

There are only handful of studies that looked into neural mechanisms underlying the use of heuristics. Volz et al. (2006, 2010) reported that recognition-heuristic-based decision processes were accompanied by enhanced BOLD responses in the anterior medial prefrontal cortex, while activity in the dorsal claustrum was correlated with fluency-heuristic-based decisions. In addition, when human participants were trained to make choices according to the Take-the-Best algorithm, activity of the dorsolateral prefrontal cortex (dlPFC) systematically varied as a function of number of cues needed to be retrieved to reach a decision (Khader et al., 2011; Khader, Pachur, Weber, & Jost, 2015), although it is difficult to dissociate whether the dlPFC mediated the Take-the-Best decision process or reflected other decision-related activities such as working memory load and updating (see Curtis & D'Esposito, 2003). Furthermore, during multisensory

decision making, BOLD responses of sensory areas and their effective connectivity with the ventral striatum depended on selection of a strategy (e.g., multi- versus single-sensory), suggesting that people can adaptively switch strategies by effectively utilizing relevant resources (Gluth, Rieskamp, & Büchel, 2014). This strategy-related control has also been found in the dorsomedial prefrontal cortex, which exhibits enhanced BOLD activation when subjects make a choice that is inconsistent with their overall preferred strategy (Venkatraman, Payne, Bettman, Luce, & Huettel, 2009).

In summary, these findings provide some initial neural evidence related to the use of heuristics, although it is yet difficult to draw connections between the neural bases of probabilistic decision making and the spontaneous adoption of satisficing strategies. That is, much of heuristics neuroscience research has focused on investigating BOLD activations underlying artificially trained heuristics, and thus, does not directly address the question of how the brain spontaneously copes with uncertainty via heuristics. Clearly, more studies are needed to understand how incomplete knowledge of decision-relevant information is integrated in the brain, and how satisficing decision making is achieved at the neural level.

1.6 Towards understanding satisficing decision making

Studies of rational and bounded rational theories of decision making, as discussed throughout this chapter, have progressed into largely separate fields of research, leading to somewhat disparate perspectives on interpreting human choice

behavior. While the rational approach has led to several groundbreaking discoveries of neural mechanisms underlying simple choices, studies of bounded rationality have primarily focused on investigating and modelling cognitive processes mediating heuristics in various complex environments. The overarching goal of this dissertation, therefore, is to characterize how satisficing heuristics are triggered and achieved via tasks that more closely model real-life challenges in decision making. Specifically, based on novel techniques developed by combining approaches from two fields of research (Chapter 2), changes in decision strategies induced by externally (time pressure; Chapter 3) and internally (anticipatory anxiety; Chapter 4) driven pressures will be examined. Then, neural correlates of satisficing under time pressure will be characterized (Chapter 5), with the aim of improving our understanding of mechanisms underlying satisficing decision making in a complex world.

2. Multi-Cue Probabilistic Classification Task and Decision Strategy Analysis

The current chapter introduces a novel multi-cue probabilistic classification task and an accompanying set of decision strategy analyses in detail. This chapter, therefore, serves as a general methods section for three empirical studies included in Chapters 3, 4, and 5. The goal of this dissertation is to quantify the adaptive use of satisficing heuristics in an environment that approximates the uncertainty associated with most real-world decision making. Therefore, we eschewed the use of explicit cue values and instead developed a probabilistic inference task, in which a decision maker is required to learn the information structure of the decision environment through trial-and-error. Then, to identify the most likely decision strategies employed by participants, a set of computational algorithms based on variational Bayesian inference was developed and applied to participants' choice data. Thus, the multi-cue probabilistic classification task and the accompanying strategy analyses procedure provide an innovative way to examine shifts in decision strategies triggered by various external or internal pressures in the presence of uncertainty. Any differences from the outlined procedures are noted for each empirical study under Methods sections of the corresponding chapter. The methods presented in this chapter have been published as a part of a research article in the *Journal of Experimental Psychology: Learning, Memory, and Cognition* (Oh et al., 2016).

2.1 Multi-cue probabilistic classification task

To create an environment where choices have to be made by actively integrating multiple pieces of information (*cues*), a multi-cue probabilistic classification task, inspired by the well-known “weather prediction task” (Knowlton et al., 1996, 1994), was developed (for additional details, see Section 1.4.1). As discussed in Chapter 1, due to limited number of cue card combinations (14 patterns), this task is potentially susceptible to memorization of specific cue-outcome relationship, which can discourage participants from actively combining evidence in each experimental trial. To overcome this limitation, instead of using four cues associated with fixed outcome probabilities, we presented two compound stimuli consisting of combinations of four binary cues on each trial (see below; Figure 1). Participants were asked to predict the stimulus that was likely to “win” as quantified by the combined weights of the cues that comprised the stimuli, which had to be learned via probabilistic feedback throughout the experiment. The large number of possible cue combinations (120 unique trials) prevented participants from memorizing specific patterns of stimuli-outcome combination, instead encouraging them to integrate available information to solve the task, allowing us to track and quantify participants’ use of cue information. Hence, the task reliably models real-life uncertainty, in which participants do not have an explicit access to the true underlying information structure but are required to learn probabilistic cue-outcome associations through trial and error.

2.1.1 Stimuli

The task stimuli consisted of 16 unique compound cues, constructed using four different visual features, color, shape, contour, and line orientation, which we refer to as *cue dimensions* (Figure 1A, Table 1). Each cue dimension was binary, comprising two sub-features or *cue states*. For instance, the cue dimension of color had the two possible cue states of blue and red. Each cue state was associated with a fixed predictive value (or “weight”) for predicting the probability of winning, and these values were complementary and summed to one within each cue dimension (Table 2). For example, in a given participant, the color blue might have a weight of 0.6, which was its probability of winning, with the color red having a weight of 0.4. The “net weight” was the difference between the state weights for a dimension, in this case 0.2. It indicated how important that individual dimension was for selecting the winning stimulus (in this case, relatively unimportant). Table 2 displays the possible cue weights assigned across different cue dimensions. As can be inferred from the net weights, we created a compensatory environment, in which the highest cue, c_1 , can be out-weighted by some combinations of less valid cues. The weights were randomly assigned to the different cue dimensions (24 permutations total) for each participant at the beginning of the experiment but always followed the organization of Table 2. By exhaustively combining all possible cue states, 16 unique compound cue stimuli were constructed (Figure 1B). As explained in detail below, the sum of the weights associated with the four cue states

comprising each stimulus governed the probability of that stimulus being a “winning” stimulus.

Table 1: Compound Stimulus Organization

| Cue dimension | Cue state | |
|------------------|------------|----------|
| | 1 | 2 |
| Color | Blue | Red |
| Shape | Circle | Square |
| Contour | White | Black |
| Line orientation | Horizontal | Vertical |

Table 2: Cue Weight Assignment

| Cue dimension | Cue state 1 | Cue state 2 | Net cue weight |
|---------------|-------------|-------------|---------------------------------|
| c_i | $w_{1,i}$ | $w_{2,i}$ | $w_{net,i} = w_{1,i} - w_{2,i}$ |
| c_1 | 0.9 | 0.1 | 0.8 |
| c_2 | 0.8 | 0.2 | 0.6 |
| c_3 | 0.7 | 0.3 | 0.4 |
| c_4 | 0.6 | 0.4 | 0.2 |

2.1.2 Task

In all three empirical studies, participants performed a probabilistic classification task, in which they were asked to compare two compound stimuli and make a prediction about which stimulus is more likely to win (i.e., to have a higher sum of weights than the other stimulus) by means of a time-restricted button press (Figure 1C). Stimuli were presented on the left and the right side of the screen along the horizontal meridian, equidistance from a central fixation cross. The stimuli were sampled from the

full set of the 16 compound cues (Figure 1B), such that each stimulus was paired with all the other stimuli except for itself, resulting in 120 unique trials. For a given trial, the compound cue stimuli could thus differ in one, two, three, or four cue dimensions, and this difference governed the underlying winning probabilities. That is, the probability that a left (L) or a right (R) stimulus would win was determined based on the cue states comprising the left stimulus, $C_L = \{c_{L,1}, c_{L,2}, c_{L,3}, c_{L,4}\}$, and the right stimulus, $C_R = \{c_{R,1}, c_{R,2}, c_{R,3}, c_{R,4}\}$, and their associated weights, $W_L = \{w_{c_{L,1}}, w_{c_{L,2}}, w_{c_{L,3}}, w_{c_{L,4}}\}$ and $W_R = \{w_{c_{R,1}}, w_{c_{R,2}}, w_{c_{R,3}}, w_{c_{R,4}}\}$:

$$P(L | C_L, C_R) = \frac{10^{\sum_{i=1}^4 (w_{c_{L,i}} - w_{c_{R,i}})}}{1 + 10^{\sum_{i=1}^4 (w_{c_{L,i}} - w_{c_{R,i}})}} \quad \text{Eqn. 1}$$

$$P(R | C_L, C_R) = 1 - P(L | C_L, C_R) \quad \text{Eqn. 2}$$

where i represents cue dimension (Table 2). For example, if there was no difference in overall weights of each stimulus, $\sum_{i=1}^4 (w_{c_{L,i}} - w_{c_{R,i}}) = 0$, then $P(L | C_L, C_R) = P(R | C_L, C_R) = 0.5$. At the other extreme, if the left stimulus consisted of cues having every one of the higher cue states (Table 2), the probability that the left stimulus would win could be calculated as, $P(L | C_L, C_R) = 10^2 / (1 + 10^2) = 0.99$. Based on equations 1 and 2, the winning stimulus was determined probabilistically on a trial-by-trial basis, and the outcome was signaled to the participant by presenting the words “win” or “lose” on the screen as post-decision feedback. Thus, since feedback was probabilistic, it was not providing participants with “correct” feedback in an absolute sense, because

there could be situations where a participant might receive negative feedback for an objectively correct decision.

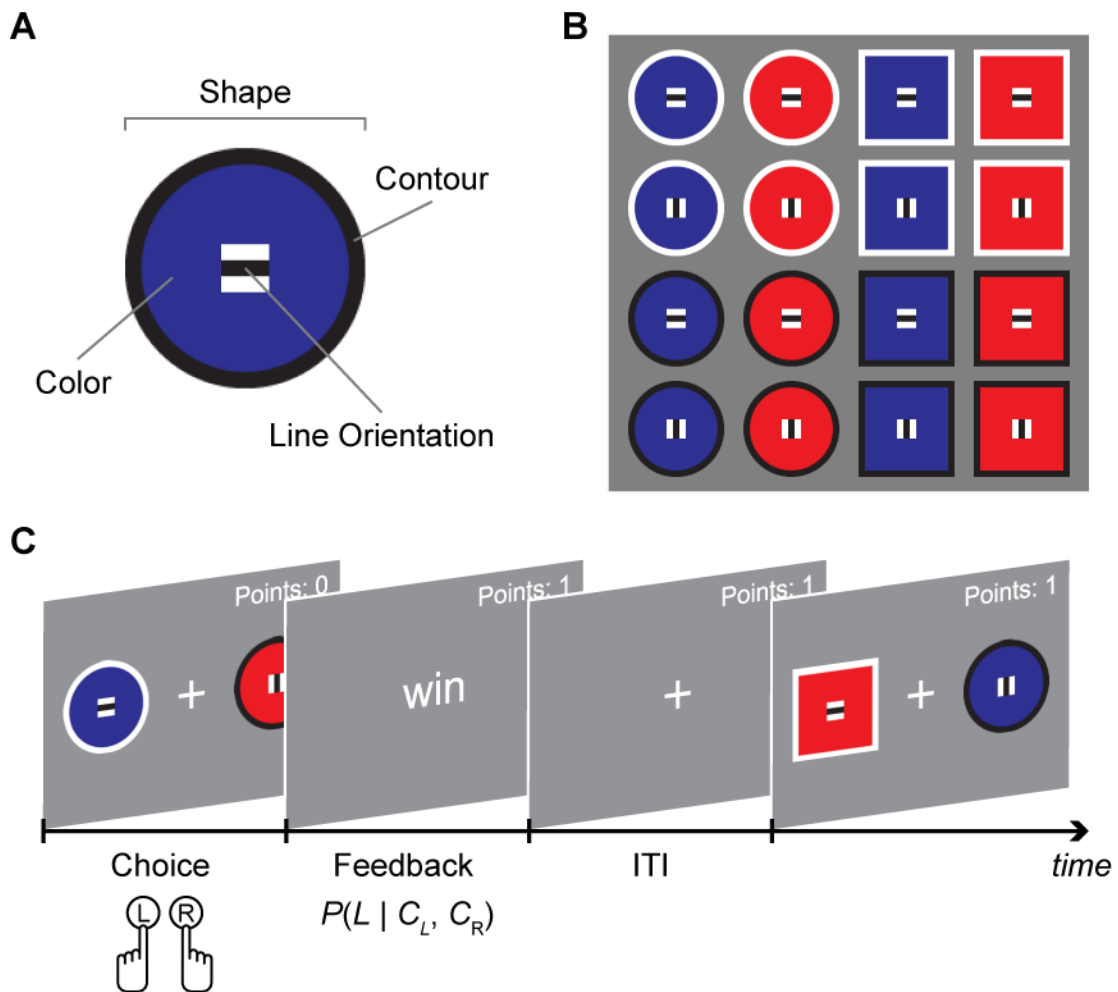


Figure 1: Compound Stimuli and Multi-Cue Probabilistic Classification Task. (A) Stimulus organization with four cue dimensions. (B) Sixteen compound stimuli used in the task. (C) Schematic of the multi-cue probabilistic task; participants were asked to predict a stimulus that is more likely to win, and received a probabilistic outcome (“win” or “lose”), followed by an intertrial interval (ITI).

2.1.3 Post-experimental survey

To gauge participants' knowledge of cue values and decision strategies, at the end of the experiment, participants filled out a survey presented in Appendix A. The post-experimental survey included multiple-choice questions about participants' beliefs concerning the cue rankings, number of cues used in each experimental condition, and cue-outcome probabilities. First, participants were asked to rank cue dimensions in order of their perceived importance in predicting outcomes (Question 1). Then, they were asked to indicate how many cue dimensions they considered in making their decisions during each experimental phase (Questions 2 and 3). Finally, to probe their knowledge on cue states, they were presented with four different compound stimuli pairs varying in only one cue dimension (Question 4). For each pair, participants were asked to choose a stimulus with a higher probability of winning and indicate how likely they will win the given trial along a range from 50% to 100%.

2.2 Decision strategy analysis using variational Bayesian inference

To track and quantify participants' use of cue information, a set of strategy analyses using logistic regression based on variational Bayesian inference (Drugowitsch, 2013) combined with a Bayesian model selection approach (Stephan, Penny, Daunizeau, Moran, & Friston, 2009) was developed. First, standard statistical analyses were used to characterize participants' choice behavior per experimental phase as well as per difficulty of each trial. Then, variational Bayesian inference was employed to identify the

most likely strategy participants may have adopted under various experimental conditions by exploring a large space of plausible decision strategy models.

2.2.1 Behavioral performance analysis

Data analyses were based on optimal choices that were favored by the probability of winning (Eqns. 1-2), independent of the probabilistic outcome provided to participants. In other words, whenever a participant chose the stimulus with the larger sum of weights, the trial was considered “correct” (optimal) even though the probabilistic nature of the feedback could have resulted in negative feedback for that particular trial. For the purpose of evaluating behavioral performance (percent optimal choices), trials with two stimuli that had an equal sum of weights were excluded, because a correct choice cannot be defined in these trials, i.e., $P(L | C_L, C_R) = P(R | C_L, C_R) = 0.5$. Hence, if a participant learned the cue weights optimally, as in case of an ideal observer (see below), and correctly integrated them, that participant could in theory achieve 100% accuracy (which would not be associated with 100% “win” feedback though). However, it has been shown in a number of previous studies that when receiving probabilistic feedback, people tend to match their choice probabilities to the outcome probabilities (see e.g., Vulkan, 2000). From this matching perspective, performance with 79% accuracy represents the ideal percentage of optimal choices in our protocol. In reporting analysis of variance (ANOVA) measures, violations of sphericity assumptions were corrected by Greenhouse-Geisser correction to the degrees

of freedom. Similarly, in reporting t-test results, degrees of freedom were corrected for unequal variance where necessary.

2.2.2 Sum of Evidence (SoE)

To formally define the objective difficulty of reaching the correct decision for a given trial, we computed each trial's SoE, the sum of available evidence, $|\sum_{i=1}^4(w_{c_{L,i}} - w_{c_{R,i}})|$, which is equivalent to the absolute value of log odds that the left (or right) stimulus will be the one yielding a positive outcome. The inner term of the SoE equation corresponds to a decision variable, where, ideally, the subject is assumed to choose left, when $\sum_{i=1}^4(w_{c_{L,i}} - w_{c_{R,i}}) > 0$ and right, when $\sum_{i=1}^4(w_{c_{L,i}} - w_{c_{R,i}}) < 0$. Hence, as the magnitude of SoE decreases, the decision becomes more difficult, which is also associated with an increasing uncertainty of observing a positive outcome. For the SoE analyses, performance was analyzed by sorting trials based solely on the magnitude of SoE and therefore, different combinations of cues that share the same SoE were categorized as the same type of event. This resulted in 11 SoE levels, ranging from 0 (no evidence, $P(L | C_L, C_R) = P(R | C_L, C_R) = 0.5$) to 2 (maximum available evidence, e.g., $P(L | C_L, C_R) = 0.99$).

2.2.3 Ideal observer model

To identify the optimal performance level for the different task phases, an ideal observer model was employed. An ideal observer was exposed to the same trial sequences and feedback as participants and learned the cue weights optimally. This

approach enabled estimation of the “ideal” cue weights that participants could have learned based on the probabilistic feedback that they received.

2.2.4 Subjective cue weights

To characterize the decision strategies participants employed in each experimental phase, the following analyses were focused on quantifying the degree to which each cue dimension affected participants’ choices. As noted above, the net weights (Table 2) correspond to the objective importance of each cue dimension in identifying the “winning” stimulus. We assumed that a decision variable (DV) is computed by cue weights that a given subject learned over time such that the subject chooses left, when $DV > 0$, and right, when $DV < 0$. Ideally, if the subject had learned all the cue weights correctly, as in the case of the ideal observer model, the corresponding psychometric function would be a step function. However, people tend to make mistakes, which can be better modeled using an S-shaped function that accounts for decision errors (see Figure 4A-B). This is equivalent to adding a little bit of noise into the DV or sampling the weights from a distribution of subjective beliefs. Hence, to assess the post-hoc, subjective importance of each cue dimension for each subject, a logistic regression analysis using variational Bayesian inference (Drugowitsch, 2013) was performed. We first defined an optimal strategy model, $Model_{opt}$, which assumes that participants utilized all four cue dimensions to guide their choices. Then, for each participant per phase, an N (total number of trials) \times 4 (number of cue dimensions)

input matrix, X_{opt} , coding for the sign of cue dimension differences on each trial, was constructed. Specifically, each element of matrix X_{opt} on i th cue dimension of n th trial was defined as following:

$$x_{n,i} = \begin{cases} -1 & \text{if } w_{cL,i} < w_{cR,i} \\ 0 & \text{if } w_{cL,i} = w_{cR,i} \\ 1 & \text{if } w_{cL,i} > w_{cR,i} \end{cases} \quad \text{Eqn. 3}$$

Logistic regression was performed based on X_{opt} , which returned parameters of a fitted logit model, w_{opt} , and a lower bound of the marginalized log-likelihood,

$P(D | Model_{opt})$, of the observed choice data, D :

$$P(\text{choice} = L | X_{opt}, w_{opt}) = \frac{1}{1 + e^{-(w_0 + X_{opt} \times w_{opt})}} \quad \text{Eqn. 4}$$

$$P(w_{opt} | \alpha_{opt}) = \mathcal{N}(w_{opt} | 0, \alpha_{opt}^{-1} I) \quad \text{Eqn. 5}$$

$$P(\alpha_{opt}) = \text{Gamma}(\alpha_{opt} | a_0, b_0) \quad \text{Eqn. 6}$$

with w_0 representing a $N \times 1$ vector for estimating the intercept.

For all three studies, two sets of hyper-parameters for the prior over weights were used. First, hyper priors were chosen so that they corresponded to a very weak prior ($a_0 = 0.01$ and $b_0 = 0.0001$). Additionally, consistent with an empirical Bayesian approach, the hyper priors, a_0 and b_0 , were estimated based on the mean (μ) and the variance (σ^2) of the observed distribution of subjective cue weights in the initial learning phase:

$$a_0 = \frac{\mu^2}{\sigma^2}, b_0 = \frac{\mu}{\sigma^2} \quad \text{Eqn. 7}$$

The goal of using informative hyper priors (Eqn. 7) was to rule out a potential source of bias in the Bayesian model selection results (see below) as Bayes factors could be highly sensitive to the prior specification.

Finally, subjective cue weights, w^* , were calculated by transforming the fitted coefficients, w_{opt} , to log base 10:

$$w^* = w_{opt} \times \log_{10} e \quad \text{Eqn. 8}$$

In contrast to the net weights corresponding to the objective importance of each cue dimension (Table 2), subjective cue weights highlight the perceived importance of each cue dimension for each subject. Since the magnitude of subjective cue weights roughly corresponds to decision noise, a subject can be considered relatively optimal when their subjective weights are perfectly correlated with the true net weights.

Similar to the generalized Take-the-Best model (gTTB; Bergert & Nosofsky, 2007), this approach is a probabilistic generalization of the cue usage based on log odds instead of a linear combination of cue weights. Although both models yield roughly similar predictions for the probability of choosing the left stimulus, $P(\textit{choice} = L | X_{opt}, w_{opt})$, the main discrepancy between the two models is that when all cues presented on a given trial (ranging from 1 to 4 cues) are in favor of the left (or right) stimulus, the gTTB will predict $P(\textit{choice} = L | X_{opt}, w_{opt}) = 1$ (or $P(\textit{choice} = L | X_{opt}, w_{opt}) = 0$) whereas the logistic model will predict the probability scaled by the sum of cue weights. For consistency with our probabilistic feedback protocol, which was also based on the sum

of cue weights (see Eqns. 1-2), we chose to adopt a logistic regression analysis instead of the gTTB to achieve better sensitivity in inferring subjective cue weights. Another popular model of decision noise is a “lapse model” that assumes random guessing in some of the trials. However, this model is not appropriate for characterizing the participants’ behavior, since a lapse model predicts psychometric curves that do not saturate at 0 or 100% when the sum of evidence is at the minimum or maximum value. This is not consistent with our empirical observations (again, see Figure 4A-B). We therefore did not attempt to fit lapse model variants to our data (see Chapter 3 for discussion). We also note that, while we have not explicitly considered a strategy which is a mixture of strategies considered, the Bayes factors that we do compute (see below) can be used to approximately determine the mixture proportions on a per subject basis (see Friston et al., 2015 for details).

2.2.5 Decision strategy models

To identify the cue dimensions that were effectively used during each phase, we further expanded the model space and explored a large set of plausible decision strategy models, accounting for every possible combination of cue usage including the optimal model defined above (Figure 2A). For a given model m ($m = 1, \dots, 14, opt$), an input matrix, X_m , was constructed using cue dimensions that were included in the model, which then was used to compute $P(D | Model_m)$ based on Equations 3-7. Then, 15

different decision strategy models¹ were compared to identify the most likely cue usage per experimental condition. The optimal model represents a compensatory strategy model where participants were assumed to integrate all four cue dimensions in making their choices, whereas models 1 through 14 consisted of different variations of sub-optimal cue weight integration. The expected accuracy for each strategy model was estimated (Figure 2B) under the assumption that a decision maker has learned to make optimal choices. As expected, the optimal cue-integration model, $Model_{opt}$, achieves the best performance, which further validates our compensatory cue structure.

Based on the strategy models, the variational Bayesian procedure described in Equations 3-7 was repeated. This approach allows estimation of an approximate posterior distribution over the weights, and it marginalizes out the uncertainty to obtain a marginalized likelihood of the model, m . Hence, Bayes factors, BF_m , were computed for each model in comparison to the optimal cue integration model:

$$BF_m = \frac{P(D|Model_m)}{P(D|Model_{opt})} \quad \text{Eqn. 9}$$

where D denotes observed data and m denotes model number ($m = 1, \dots, 14$). Any model that had BF_m greater than 3 was considered as having greater evidence in its favor than the optimal cue-integration model (Kass & Raftery, 1995). Using this approach, we were able to infer the most likely strategy that each participant may have

¹ Empirical studies included in Chapter 3 also consider the Take-the-Best (TTB) heuristic model (Gigerenzer & Goldstein, 1996).

employed to solve the task. Hence, the advantage of using such analysis is that even if participants did not learn optimally, it is possible to deduce the most likely cue usage based on each individual's choice patterns.

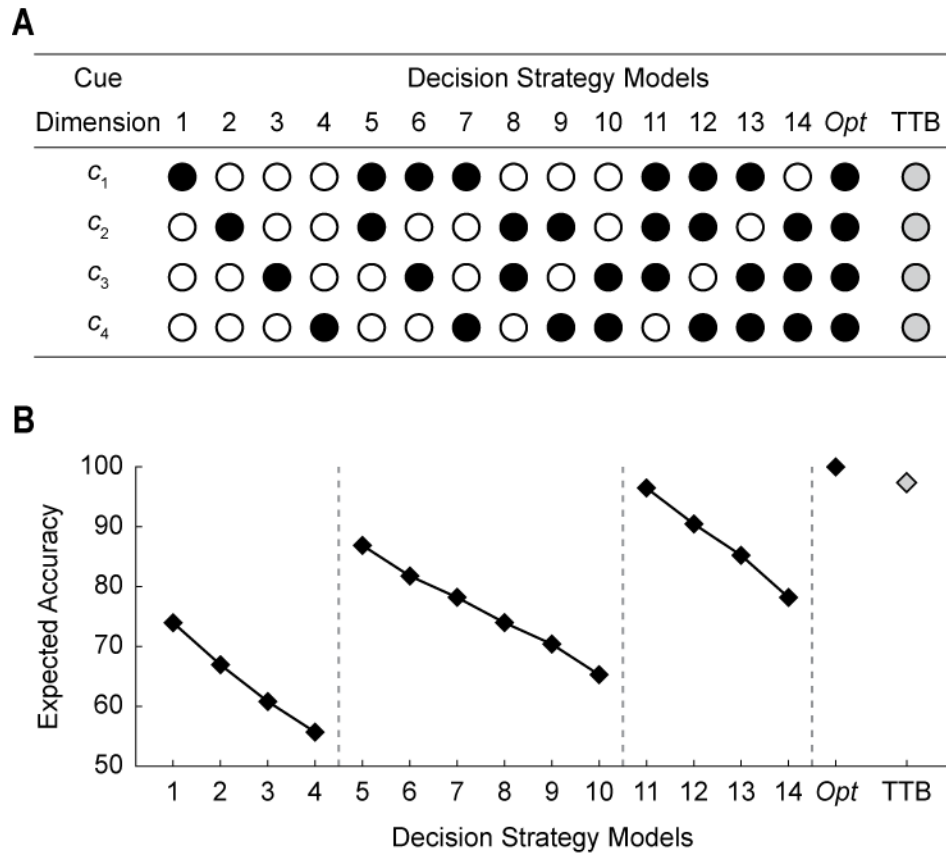


Figure 2: Decision Strategy Models. (A) Sixteen strategy models, including Model *Opt*, an optimal cue-integration model, and Model TTB, the Take-the-Best heuristic model. Filled circles denote the cues that are included in a given model. (B) Expected accuracy per strategy model. Dotted lines separate strategy models according to the number of cues used.

2.2.6 Strategy model selection

To characterize decision strategies at the group level, a Bayesian model selection procedure was employed by submitting the log model evidences obtained from the variational Bayesian inference above (Rigoux, Stephan, Friston, & Daunizeau, 2014; Stephan et al., 2009). This approach fits the hierarchical model by treating models as random effects that could vary across subjects and estimates the parameters of a Dirichlet distribution to obtain the probabilities for all strategy models considered. These probabilities are then used to define a multinomial distribution, which can be used to estimate the probability that model m generated the data of each subject. Finally, the exceedance probabilities estimated from this procedure reflect the belief that a model, m , is more likely than any other model, given the marginalized likelihoods. Unless specified, the Bayesian model selection results reported were calculated using the `spm_BMS` routine of the SPM12 software suite (<http://www.fil.ion.ucl.ac.uk/spm/software/spm12/>).

3. Satisficing Decision Making Under Time Pressure

As discussed in Chapter 1, in the real world, rationality is usually *bounded* by incomplete knowledge of the decision-relevant information and limitations placed upon cognitive resources and available computation time (Simon, 1955, 1956, 1990). Humans are thought to overcome these limitations via *satisficing*, the use of heuristic strategies that simplify decision-making problems by prioritizing some sources of information while ignoring others. In this first empirical chapter, we characterize some of the principles that govern satisficing decision-making due to uncertain information and high time pressure, by combining a novel probabilistic classification task with Bayesian strategy analysis techniques as outlined in Chapter 2.

The presence of time pressure is known to influence the decision process (for an overview, see Svenson & Maule, 1993), fostering the use of heuristics that can be applied quickly within a choice deadline. Several studies have shown that, under time pressure, people engage in a more selective information search (Böckenholt & Kroeger, 1993; Lamberts, 1995; Maule, 1994; Payne et al., 1988; Rieskamp & Hoffrage, 2008) and employ memory-based heuristics when possible (Goldstein & Gigerenzer, 2002; Lamberts, 1995, 2000; Nosofsky & Palmeri, 1997; Pachur & Hertwig, 2006). In fact, under certain circumstances, decisions made under high time constraints or limited cue exposure can even be found to be more accurate compared to those made after a long deliberation period (Ballew & Todorov, 2007; Wilson & Schooler, 1991), suggesting that the use of

fast heuristics can sometimes lead to better choices (see also Gigerenzer & Gaissmaier, 2011). In the perceptual decision-making literature, making choices under time constraints often results in a speed-accuracy trade-off that is characterized by a (conservative) shift of the decision or response threshold (see e.g., drift diffusion model, Ratcliff, 1978). It has been shown that when making such judgments, well-trained human participants can adaptively adjust this threshold to maximize reward rate (Balci et al., 2011; Simen et al., 2009) as well as accuracy (Bogacz, Hu, Holmes, & Cohen, 2010), and optimally integrate multiple pieces of information (Drugowitsch, DeAngelis, Klier, Angelaki, & Pouget, 2014). However, prior studies concerned with comparing cue-based decision strategies under different time pressures have dealt exclusively with rather slow-paced decision scenarios, where time pressure conditions could range from 15-50 seconds (Bergert & Nosofsky, 2007; Payne et al., 1988, 1996; Rieskamp & Hoffrage, 2008). Other studies that required faster decisions (deadlines ranging between 600 to 1600 ms) either extensively trained subjects with exemplars (Lamberts, 1995) or used choice problems that elicited participants' prior knowledge (Pachur & Hertwig, 2006), both of which encourages the use of memory-based heuristics. These circumstances, however, do not always approximate the kind of time pressure often faced in everyday and high-stakes decision-making, like in traffic, medical, or military scenarios, where choices have to be made by actively integrating multiple cues within a fraction of a second.

To gauge decision strategies under severe time constraints, we imposed decision time pressures in the sub-second range. Specifically, we first report two experiments, involving two independent cohorts of participants that each performed the same multi-cue probabilistic classification task using compound (or integrated) cues: following an initial learning period, we assessed and compared decision making between two post-learning task phases, one unpressured phase (2s response window), which was common to both experiments, and a subsequent high time pressure phase, where the degree of time pressure differed between the two groups of participants, ranging from moderate (Experiment 1: 750ms response window) to severe (Experiment 2: 500ms response window). We then employed variational Bayesian inference and Bayesian model selection analyses to infer the subjects' decision strategies under the different time pressure conditions. Participants may feasibly approach this task in a number of different ways, including considering a random subset of cues, using a subset of cues with strong cue weights, using all the available cues, engage in a take-the-best strategy, and so forth. To explore a large space of plausible decision strategies, we developed and contrasted 16 different plausible strategy models, allowing us to systematically track how participants integrate available information under changing time pressure. Then, to test generalizability of our findings, we conducted a third experiment using non-compound cues with a 3s unpressured phase followed by a 750ms time pressure phase. In sum, we used a large set of abstract cue stimuli combined with probabilistic feedback

to infer how people learn to use multiple cues in the presence of uncertainty and how this usage changes as a function of high time pressure. The results establish that, under split-second time pressure, humans satisfice decision making by strategically discounting (or ignoring) the least informative cues. The data presented in this chapter have been published as a research article in the *Journal of Experimental Psychology: Learning, Memory, and Cognition* (Oh et al., 2016).

3.1 Experiment 1: Satisficing under moderate time pressure

3.1.1 Methods

For all experiments, we have reported all measures, conditions, data exclusions, and how we determined our sample sizes.

3.1.1.1 Participants

48 volunteers were recruited online through a human intelligence task (HIT) via Amazon Mechanical Turk (MTurk). Assuming a medium-to-large effect size, we calculated that 38 usable data sets would provide us with a power of 0.8 at a Type I error level of 0.01 (Cohen, 1992). We recruited 10 extra participants in anticipation that about 10 to 20% of respondents would fail to meet performance criteria. All participants provided informed consent in line with Duke University institutional guidelines. To ensure high data quality, we followed MTurk's recommended inclusion criteria of only inviting participants who had previously completed ≥ 500 HITs with a HIT approval rate of $\geq 90\%$ to participate in the experiments. We did not have any restrictions in age,

demographics, or performance in approving the HIT, but assignments with more than 75 invalid trials (> 10% of total number of trials) were rejected. Five participants were excluded from further analysis due to chance-level performance gauged by the percent optimal responses in the initial 480 trials prior to the time pressure phase (one-tailed binomial test, $p > 0.05$; see Method), leaving 43 participants (mean age = 34.4 years, SD = 10.5, 21 – 66 years; 17 female, 26 male). Participants were compensated with \$5.00 upon completion of the experiment, which lasted approximately one hour. In addition, a bonus payment of \$5.00 was given to the participant who earned the highest point.

3.1.1.2 Procedure

Participants performed the multi-cue probabilistic classification task, as described in Chapter 2. Compound stimuli (150 × 150 pixels each) were presented on the left and the right side of the screen (window size: 1000 × 700 pixels) along the horizontal meridian, at an eccentricity of 250 pixels from a central fixation cross. On each trial, stimuli were presented on the screen until a response was made or for the duration of an assigned response window of 2 s or 750 ms depending on an experimental phase. Once a response was made, probabilistic feedback consisting of the words “win” or “lose”, was displayed for 1 s, followed by a 1 s inter-trial interval (Figure 3A). A trial was considered invalid if a response was not made within a given response window or was made faster than 150 ms post-stimulus. At the beginning of the experiment, participants were informed that they needed to learn about the values (weights) associated with the

different cue states by trial-and-error, in order to collect as many points as possible.

Participants earned 1 point for every winning trial and the total score at the end of the experiment was used to select the participant who received a bonus payment.

Our goal was to first allow participants to learn the (uncertain) cue values, and to then compare decision making strategies between conditions of low vs. moderate time pressure. To this end, participants completed three different phases of the probabilistic classification task, (1) an initial learning phase, followed by (2) a low time pressure (LowP) phase, followed by (3) a high time pressure (HighP) phase to create satisficing pressure. In each phase, participants completed 240 trials, consisting of the full set of possible trials presented twice, where stimuli in second set were presented in the opposite locations (i.e., as mirror images) to the first set. Trials were grouped into 12 blocks of 60 trials each, with short breaks in between. In both the initial learning period and LowP phases, participants were given a maximum of 2 seconds to respond. Then, following the LowP phase, participants performed an additional 240 trials of the HighP phase, in which participants were given a 750 ms response window (moderate time pressure). This time pressure was based on pilot work with this task, where, without time pressure, we obtained mean response times of around 700 ms. Hence, trial sets 1 and 2 made up the learning phase, sets 3 and 4 comprised the LowP phase, and sets 5 and 6 formed the HighP phase. Probabilistic feedback was provided throughout all three phases of the experiment. Our analyses focused on assessing and contrasting decision

making strategies in the experimental (post-learning) phases, that is, in the LowP and HighP phases.

3.1.1.3 Survey

At the end of the experiment, participants filled out a post-experimental survey, which is presented in Appendix A. “Self-report strategy models” (as opposed to objectively inferred ones) were constructed based on the answers to Questions 1-3. For example, if a participant ranked cues as color, shape, contour, and line orientation (ordered from most to least informative) and claimed to have used 2 cues during the HighP phase, then his/her self-report HighP model was a model that included only color and shape cues.

3.1.1.4 Data analysis

See Chapter 2 for the complete procedure. Additional data analyses performed in this chapter are detailed below.

Decision strategy models. In all three experiments included in the current chapter, Model TTB was additionally constructed based on the Take-the-Best satisficing algorithm (Gigerenzer & Goldstein, 1996), which searches through the cues in order of descending value until it finds the first cue that differentiates between the two stimuli. For example, although both Model 1 and TTB each use a single cue to arrive at a decision, Model 1 assumes that the subject uses only the highest value (most discriminatory) cue dimension, c_1 , and makes a random guess if this cue does not differ

between the two stimuli, whereas Model TTB searches through the cue dimensions from c_1 to c_4 in descending value order until it finds the highest-value cue that discriminates between the two stimuli.

Bayesian model selection. As outlined in Chapter 2, two sets of hyper priors (Eqn. 6) were used to compute the marginalized log-likelihood. Specifically, we initially employed relatively uninformative hyper priors, $a_0 = 0.01$ and $b_0 = 0.0001$. Then, in a control analysis, we established that changing the hyper priors based on the posterior distributions of the weights computed from the initial learning phase did not yield any significant influence on Bayes factors nor model comparison results (see Figure 5). The informative hyper priors were estimated using Equation 7, which yielded following values: $a_0 = 0.594$ and $b_0 = 0.815$ for Experiment 1, $a_0 = 0.501$ and $b_0 = 0.798$ for Experiment 2, and $a_0 = 0.479$ and $b_0 = 1.804$ for Experiment 3. For the group-level model selection procedure, the exceedance probabilities were calculated by submitting the approximate log model evidences to the `spm_BMS` routine of the SPM8 software suite (<http://www.fil.ion.ucl.ac.uk/spm/software/spm8/>).

3.1.2 Results

3.1.2.1 Task performance

Prior to analyzing how subjects weighted the cues and arrived at their decisions, we briefly summarize general task performance. Mean proportion of optimal choices and response time (RT) for each set and phase are shown in Table 3. Here, an optimal

choice refers to the “correct” decision favored by the sum of cue weights, independent of the actual feedback provided. As can be seen in Figure 3B (black line), participants gradually learned to choose the higher-value stimuli in the learning phase, achieving 72.1% optimal responding by set 2. We observed no significant difference in performance between set 2 and 3 ($t(42) = 0.1, p = 0.92$), suggesting that performance had stabilized by the end of the learning period. However, percent optimal choices increased from set 3 to set 4 ($t(42) = 3.4, p = 0.001$), thus suggesting that some residual cue learning was still taking place during the LowP phase. This residual learning effect is not surprising since probabilistic feedback was provided throughout the experiment. As expected, the percentage of optimal choices and RT reliably scaled with the difficulty of the decisions, as defined by the difference in sum of cue weights between the two stimuli on each trial. More difficult decisions (smaller weight differences) were associated with decreasing percent correct choices (linear trend: $F(1,42) = 322.1, p < 0.001$) and increasing RT (linear trend: $F(1,42) = 45.2, p < 0.001$).

In the subsequent HighP phase (Figure 3B, black line), participants experienced a moderate time pressure of 750 ms, which was close to (but greater than) the mean RT in the LowP phase (see Table 3). Here, the rate of optimal choices did not differ significantly from the LowP phase ($t(42) = 1.6, p = 0.11$), but responses were accelerated ($t(42) = 11.5, p < 0.001$), indicating that participants modified their approach to the task to adapt to the higher time pressure. We found no significant

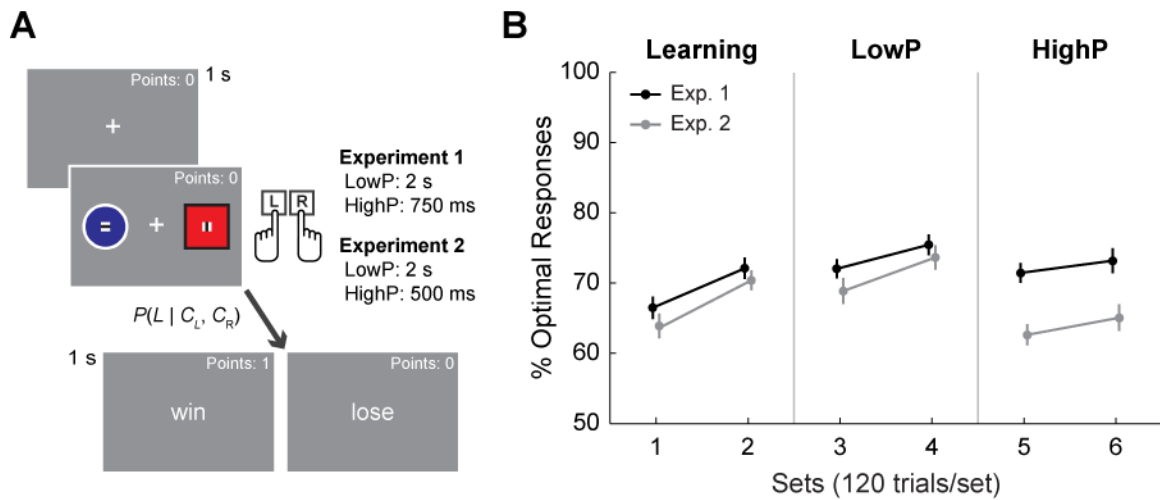


Figure 3: Task and Performance of Experiments 1 and 2. (A) Probabilistic classification task under time pressure; participants performed the task under varying response windows of 2 s (learning and LowP phases), 750 ms (HighP phase in Experiment 1), and 500 ms (HighP phase in Experiment 2). (B) Average percent optimal responses in Experiment 1 (750 ms time pressure) and Experiment 2 (500 ms time pressure) during the learning, LowP, and HighP phases. Error bars are standard errors (SEM).

difference in performance between set 5 and 6 ($t(42) = 1.5, p = 0.15$). Similar to the LowP phase, both percent correct choices (linear trend: $F(1,42) = 137.8, p < 0.001$) and RT (linear trend: $F(1,40) = 117.1, p < 0.001$) were modulated by decision difficulty. Finally, assessing the potential impact of basic individual differences on task performance, we found no effect of gender on choices in either task phase (LowP: $t(41) = 0.9, p = 0.38$; HighP: $t(41) = 1.1, p = 0.29$), and no correlation between age and performance (Pearson's correlation, LowP: $r = -0.2, p = 0.19$; HighP: $r = -0.1, p = 0.51$). In sum, these initial analyses demonstrate basic statistical learning of our task as well as an effect of time pressure on RT. To examine potential shifts in decision making

strategies under time pressure, we next turned to determining participants' cue weighting and cue-integration strategies in performing the task.

Table 3: Summary of Task Performance. Numbers in parenthesis are SEM.

| Phase | Set | Experiment 1 | | Experiment 2 | | Experiment 3 | |
|----------|-----|-------------------|-------------|-------------------|-------------|-------------------|--------------|
| | | % optimal choices | RT (ms) | % optimal choices | RT (ms) | % optimal choices | RT (ms) |
| Learning | 1 | 66.5 (1.6) | 732 (27) | 63.6 (1.8) | 749 (37) | 62.9 (1.5) | 1150 (61) |
| | 2 | 72.1 (1.5) | 702 (24) | 70.4 (1.4) | 729 (33) | 69.4 (1.6) | 1122 (53) |
| LowP | 3 | 72.0 (1.4) | 682 (23) | 68.9 (1.8) | 701 (34) | 72.8 (1.3) | 1126 (52) |
| | 4 | 75.4 (1.5) | 658 (21) | 73.6 (1.7) | 674 (32) | 74.5 (1.5) | 1076 (52) |
| HighP | 5 | 71.4 (1.4) | 457 (8) | 62.6 (1.5) | 359 (7) | 62.7 (1.4) | 510 (13) |
| | 6 | 73.2 (1.8) | 458 (8) | 65.1 (1.9) | 367 (6) | 63.2 (1.6) | 511 (11) |

3.1.2.2 Subjective cue weights

To assess the extent to which each cue dimension (e.g., color) affected participants' decisions, we performed logistic regression (Eqns. 3-6, Figure 4) based on their performance in the LowP and HighP phases. The output of the logistic regression analysis corresponds to log odds of choosing the left stimulus given the presence of differing cue states, and thus, provides a way to measure how much subjective net weight was assigned to each of the four cue dimensions (as contrasted with the

objective, *a priori* established net weights; see Table 2). Logistic regression curves for an example participant are displayed in Figure 4A and B. At the population level (Figure 4C), an ANOVA with assigned cue weights ($w_{net.i}$) and experimental phase revealed a significant main effect of cue weights ($F(2.2,91.98) = 20.3, p < 0.001$), which was characterized by a significant linear trend ($F(1,42) = 33.7, p < 0.001$), suggesting that participants had correctly learned the relative ranks of the cues. There was no effect of LowP/HighP phase ($F(1,42) = 2.0, p = 0.17$) and no cue weight \times phase interaction ($F(3,126) = 0.02, p = 0.99$). Subjective cue weights of each participant were further compared to weights produced by an ideal observer, by computing Spearman's rank correlation coefficients. An ideal observer was exposed to the same stimulus and feedback sequences and hence, ideal weights reflect weights that participants would have acquired if they had learned optimally. Table 4 shows a summary measure of correlation coefficients averaged across participants (tested for significance using standard t-tests). All the mean correlation coefficients were significantly greater than zero, confirming positive correlations between participants' subjective cue weights and the ideal weights. In sum, the subjective cue weight analysis indicates that participants learned to accurately rank the informational values of the four cues. Next, we determined how participants combined the cue weights in reaching their decisions.

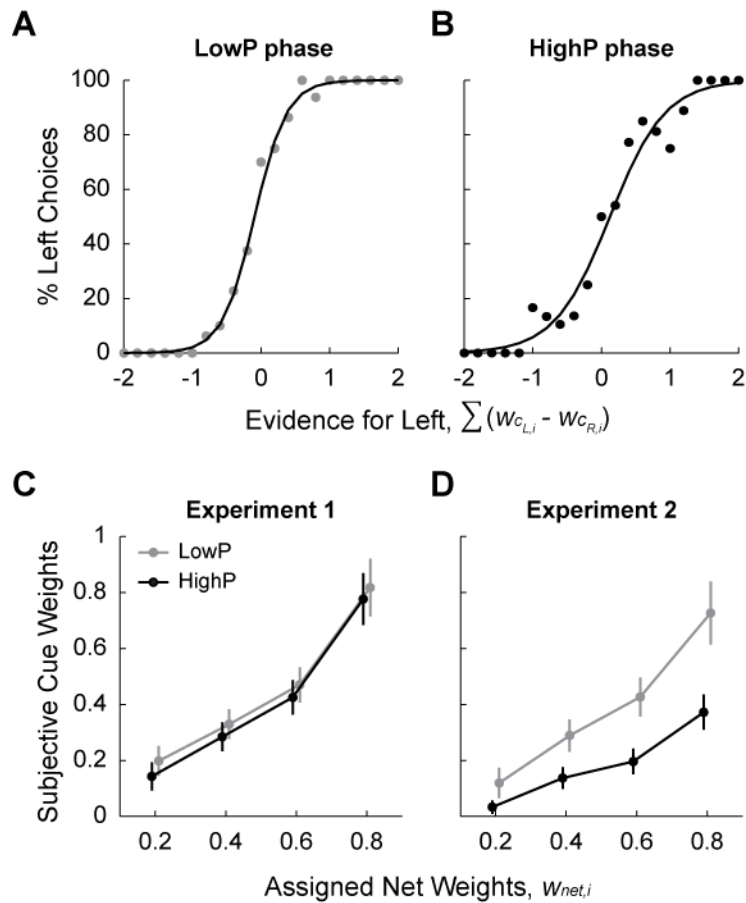


Figure 4: Logistic Regression Fit of Individual Performance and Estimation of Subjective Cue Weights. (A) Performance of a single example participant in the LowP phase (2 s response window). (B) Performance of the same participant in the HighP phase (750 ms response window). An individual data point reflects percentage of left choices made given the sum of evidence. Curves are logistic regression fits to the data. (C) Average subjective cue weights for Experiment 1 (750 ms HighP). (D) Average subjective cue weights for Experiment 2 (500 ms HighP). Error bars are SEM.

3.1.2.3 Decision strategy model comparison and selection

To quantify participants' decision-making strategies, we used variational Bayesian inference (Drugowitsch, 2013) to gauge evidence for different decision models

in reference to the optimal cue-integration model, Model *Opt* (Figure 2A; see Methods and Chapter 2). A marginalized likelihood of each model per participant was then used to characterize overall decision strategies at the group level through Bayesian model comparison (Stephan et al., 2009). In the LowP phase, the optimal cue-integration model, presented as Model 15, was the winning model with exceedance probabilities of 0.78 (Figure 5A). Under 750 ms time pressure (HighP phase), however, a model that used only the three most predictive cues (Model 11, using c_1 , c_2 and c_3) was the most likely model with an exceedance probability of 0.67 (Figure 5A). The additional control analysis using informative priors estimated based on the posterior distributions of the weights from the initial learning phase yielded highly similar results (Figure 5B). These results suggest that participants shifted from using all four cues when having no time pressure to using only the three highest-value cues under moderate time pressure of 750 ms.

Taken together, the subjective cue weight and decision model analyses suggest that (a) participants learned to correctly rank the values of the cue dimensions, and (b) under moderate time pressure, they disregarded the least valuable cue dimension. We reasoned that the latter “dropping” of the worst cue dimension from the decision process under moderate time pressure could reflect one of two processes: it could either reflect a strategic shift in processing, whereby participants categorically ignore the worst cue dimension in their decision making, or it could simply be due to the fact that

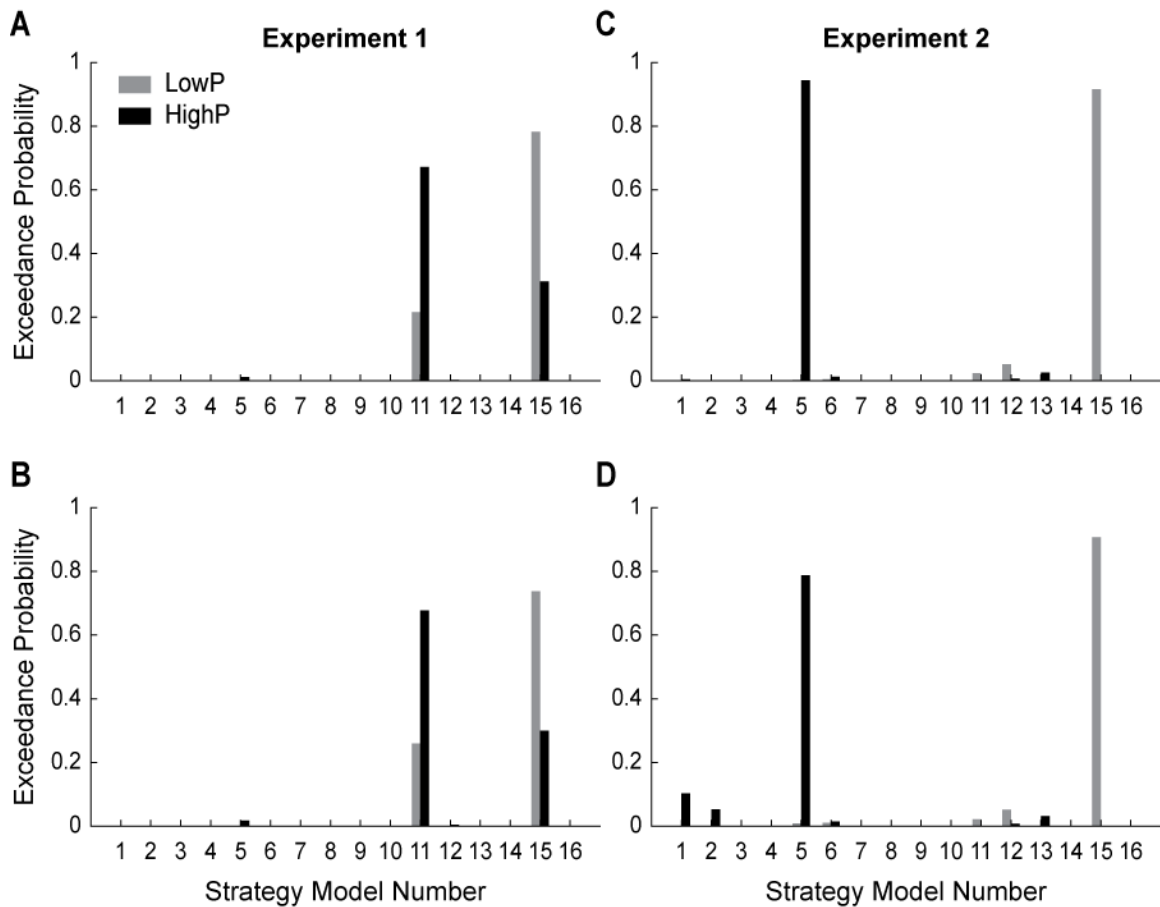


Figure 5: Bayesian Strategy Model Selection. (A) Experiment 1 (750 ms HighP) using uninformative priors. (B) Experiment 1 using informative priors. (C) Experiment 2 (500 ms HighP) using uninformative priors. (D) Experiment 2 using informative priors. Model 15 represents the optimal cue integration model (*Opt*) and Model 16 represents the Take-the-Best (TTB) model.

participants are running out of time in a serial, value-ranked cue integration process. In the latter scenario, participants would still attempt to use the worst cue, but on most trials they do not have sufficient time to process it, because they first attend to higher valued cues. The fact that the average RT for these trials in the HighP phase was well below the response deadline of 750 ms ($mean = 481$ ms; one-sample t-test: $t(42) =$

29.2, $p < 0.001$) speaks against the hypothesis that participants simply ran out of time in considering this cue.

However, to adjudicate between these two possibilities in greater detail, we analyzed choices in trials where the worst cue was the *only* distinguishing dimension between the two stimuli (16 trials/phase). In these trials, participants should be able to evaluate that cue dimension, since they do not have to spend time on integrating differential values over the other cue dimensions, and the single differentiating visual feature should be quite salient perceptually. Thus, if subjects dropped the worst cue in the HighP phase due to a lack of time in serial cue integration, they should nevertheless perform above chance on these trials. By contrast, if subjects strategically disregarded the worst cue dimension under time pressure, they would simply guess on these trials. This analysis revealed that 30 out of 43 participants (70%) performed no different from chance level (binomial test, $p > 0.05$) when only the least important cue differentiated the stimuli, indicating that majority of participants guessed in these trials. To further corroborate this result, we computed Bayes Factors (BF) comparing two models, with H_0 : the probability of correct choice, $p = 0.5$ and H_1 : all possible values of p in $[0, 1]$ is equally probable. This analysis revealed that 28 out of 43 participants (65%) had $BF > 1$ favoring H_0 although only 7 participants showed significant effect ($BF > 3$; Kass & Raftery, 1995). At the group level, we estimated BFs for use in a paired t-test (adopted from Rouder, Speckman, Sun, Morey, & Iverson, 2009) based on average performance

between the LowP ($mean = 58.9\%$, $SEM = 3.5$) and HighP ($mean = 59.6\%$, $SEM = 3.1$) phases. This t-test did not show any significant difference ($BF = 5.88$, favoring the null hypothesis). In sum, we obtained equivocal results in this selective analysis of low-value cue decisions, which do not allow us to draw strong conclusions about whether participants truly performed at chance on these trials. Note though that power in these analyses is limited, as we only considered a small number of trials (16 trials per phase) per participant.

Comparison of the HighP phase results between the estimated subjective cue weights (Figure 4C) and the strategy model selection (Figure 5A, B), at a glance, may seem odd since the population average of the weight of the least important cue (c_4) does not drop to zero, while the modelling results suggest that these cues were dropped from the decision making process. It should be noted though that although the majority of participants learned the correct cue rankings, there was a lot of variability in estimated cue weights across subjects. Given that the magnitude of cue weights also roughly corresponds to the amount of decision noise in each participant, fitted subjective cue weights were also highly dependent on individual task performance. The combination of an increase in decision noise and a decrease in number of observations in the HighP phase, therefore, may have contributed to the increased variability in the subjective cue weights. To rule out the possibility that this increase in noise affecting the estimation subjective cue weights under time pressure is what is driving our group results of the

Bayesian model selection, we re-ran the analysis using only the least noisy half of our data set, that is “high performing” subjects, based on a median split of performance accuracy in the initial learning phase (average performance $\geq 65\%$). The results of this analysis remained the same as for the whole sample (data not shown). Furthermore, we observed a significant down-weighting of the least important cue in the HighP phase compared to the LowP phase ($t(20) = 2.5, p < 0.05$), whereas no significant difference was observed for the rest of the cue weights.

Overall, our decision strategy analysis provides evidence that participants engaged in optimal, exhaustive cue value integration when time pressure was low, at least when dealing with compound cues. When faced with moderate time pressure, however, our analyses provided some evidence indicating that the participants satisfied strategically through cue discounting, compensating for the lack of time via a “drop-the-worst” cue strategy, that is, they ignored the least predictive sub-set of the cue information. As shown in Figure 2B, switching from the optimal cue integration model, Model *Opt*, to Model 11 leads to only a negligible difference in expected accuracy (97%), which further supports the adaptive nature of strategy selection under moderate time pressure.

3.1.2.4 Survey results

To gauge the degree to which decision-making was driven by explicit knowledge concerning the different cue states, we analyzed a range of survey questions. When

participants were asked to define a cue state with higher weight within a cue dimension (e.g., blue vs. red), the mean number of correct response (out of a possible maximum of 4) was 3.4, indicating that majority of participants were able to explicitly understand the relative importance of the cue states within each dimension. Then, to identify a “self-report strategy model” for each participant (see Methods), we surveyed relative ranks of cue dimensions and number of cue dimensions considered in each phase. As can be seen in Table 4, participants’ ratings of relative rank of cue dimensions, from the most to the least informative, showed a significant positive correlation with (1) the subjective cue weights estimated using logistic regression analysis, and (2) the ideal weights estimated from the ideal observer model.

In terms of number of cue dimensions used in each phase, participants answered that they used significantly fewer cues ($t(42) = 7.8, p < 0.001$) during the HighP phase ($mean = 1.8, SD = 0.7$) compared to the LowP phase ($mean = 2.6, SD = 0.7$). These self-report models, however, did not match well with our findings based on the objectively inferred strategy models. That is, only five participants out of 43 were able to correctly identify their own decision strategies for both the LowP and HighP phases. In sum, the survey results indicate that participants gained some explicit knowledge of the relative cue weights; however, they did not have much insight into the strategies they employed in translating these cue weights into decisions.

Table 4: Mean Spearman’s Rank Correlation Coefficients Between Ideal Weights, Subjective Weights, and Explicit Survey Ratings¹.

| | LowP | HighP |
|---------------------|-------|-------|
| Experiment 1 | | |
| Ideal - Subjective | .48** | .45** |
| Ideal - Survey | .45** | .44** |
| Subjective - Survey | .54** | .48** |
| Experiment 2 | | |
| Ideal - Subjective | .55** | .45** |
| Ideal - Survey | .31* | .27* |
| Subjective - Survey | .37** | .41** |
| Experiment 3 | | |
| Ideal - Subjective | .55** | .50** |
| Ideal - Survey | .61** | .54** |
| Subjective - Survey | .67** | .71** |

3.1.3 Discussion

In order to characterize satisficing decisions under uncertainty and split-second time pressure, we developed a new multi-cue statistical learning protocol and applied a range of analyses to infer participants’ learning of cue values and their use in the decision process. The results of Experiment 1 document that participants reliably acquire knowledge of cue values from trial and error learning, and Bayesian model selection suggests that they employ a decision strategy of (optimal) exhaustive cue information

¹ Ideal = ideal weights produced by an ideal observer; Subjective = subjective weights computed based on participants’ behavioral data; Survey = participants’ self-reported cue ranks.

* $p < 0.05$ and ** $p < 0.001$ from the one-sample t -test.

integration when they are not under time pressure (2s response window); however, when put under moderate pressure (750ms response window), participants appear to strategically adapt their decision process by ignoring the least valuable cue dimension. Thus, Experiment 1 seems to have revealed a novel satisficing strategy of “drop-the-worst” cue under time pressure. However, a confirmatory analysis focusing on those trials where only the lowest-value cue differentiated between the two stimuli produced equivocal results. The latter might be attributable to a small trial count ($n = 16$ trials/phase) and/or a subset of participants who were able to perform optimally even under moderate time pressure (see Figure 5A). Therefore, we followed up these initial findings with a second experiment that pursued two main goals. First, we aimed to replicate the decision making pattern observed in Experiment 1. Second, we sought to characterize decision making when time pressure was increased even further. If “drop-the-worst” is a reliable general strategy in split-second satisficing, then we would expect participants to further trim their usage of lower value cues as the time pressure increases. To test this hypothesis, we ran an exact replication of Experiment 1 in a new cohort of subjects, with the only difference being that the HighP phase in Experiment 2 was reduced, from 750 ms to 500 ms.

3.2 Experiment 2: Satisficing under severe time pressure

3.2.1 Methods

3.2.1.1 Participants

A new, non-overlapping cohort of 40 volunteers was recruited through a separate HIT on Amazon MTurk with the expectation of collecting 38 usable data sets as in Experiment 1. The same HIT approval criteria used in Experiment 1 were applied to ensure high data quality. Four participants were further excluded due to chance-level performance (one-tailed binomial test on percent optimal responses in the initial 480 trials without time pressure; $p > 0.05$), leaving a total of 36 participants (mean age = 32.0 years, $SD = 10.4$, 21 – 64 years; 15 female, 20 male, 1 unknown). Participants were paid \$5.00 upon completion of the experiment and an additional \$5.00 bonus payment was given to the participant who achieved the highest score. The age ($t(76) = 1.0, p = 0.31$) and gender ($\chi(1) = 0.1, p = 0.77$) distribution of participants in Experiment 2 was equivalent to those of Experiment 1.

3.2.1.2 Procedure

Experiment 2 was identical to Experiment 1, except that the response time window in the HighP phase was reduced from 750 ms to 500 ms (less than the mean reaction time of the LowP phase in Experiment 1) to observe how decision strategies change under more severe time pressure.

3.2.2 Results

3.2.2.1 Task performance

Similar to Experiment 1, participants were able to gradually improve their performance throughout the first 240 trials in the learning phase. We observed no significant difference in performance between set 2 and 3 ($t(35) = 1.0, p = 0.31$), suggesting that performance had stabilized by the end of the learning period. However, as in Experiment 1, optimal choices increased from set 3 to set 4 ($t(35) = 3.8, p = 0.001$), thus suggesting that some residual cue learning was still taking place during the LowP phase as participants continued to observe probabilistic feedback provided upon each choice (see Figure 3B, gray line, and Table 3). As expected, in the LowP phase, the percentage of optimal choices and RT were affected by decision difficulty with more difficult decisions resulting in a significant decrease in performance (linear trend: $F(1,35) = 267.4, p < 0.001$) and increase in RT (linear trend: $F(1,35) = 29.3, p < 0.001$). Also as expected, given the equivalence between the learning period and LowP phases between Experiments 1 and 2, no difference in performance of set 1 through 4 between experiments was observed (main effect of experiment: $F(1,77) = 1.5, p = 0.23$; experimental group \times set interaction: $F(3,231) = 0.3, p = 0.81$), indicating comparable cue learning and LowP performance across the two cohorts.

In the HighP phase (Figure 3B, gray line), however, both optimal responses ($t(35) = 5.5, p < 0.001$) and RT ($t(35) = 11.0, p < 0.001$) decreased significantly compared to the LowP phase. Accordingly, the percentage of optimal choices in the

HighP phase was significantly lower in the severe pressure group than in the moderate pressure group (between-subject effect of experiment: $F(1,77) = 15.3, p < 0.001$). There was no difference in performance between set 5 and 6 ($t(35) = 1.7, p = 0.10$). Despite a significant change in performance, the effect of decision difficulty on percent optimal choices (linear trend: $F(1,35) = 32.1, p < 0.001$) and RT (linear trend: $F(1,29) = 68.6, p < 0.001$) was still present in this phase. Similar to Experiment 1, there was no significant effect of gender on task performance (LowP: $t(33) = 0.2, p = 0.83$; HighP: $t(33) = 0.7, p = 0.48$) as well as no significant correlation between age and performance (Pearson's correlation, LowP: $r = 0.02, p = 0.90$; HighP: $r = -0.1, p = 0.42$). In sum, throughout the learning period and LowP phases, which were equivalent between Experiment 1 and 2, participants successfully acquired the (uncertain) values of the different cue dimension via trial-and-error learning. However, there was clear evidence that the 500 ms time pressure in the HighP phase of Experiment 2 had a detrimental effect on performance, both compared to the LowP phase of Experiment 2, as well as to the 750ms HighP phase in Experiment 1.

3.2.2.2 Subjective cue weights

Separate sets of subjective cue weights for the LowP and HighP phases were obtained using logistic regression (Figure 4D). An ANOVA revealed a significant effect of cue weights ($F(2.2, 76.9) = 14.2, p < 0.001$), with a significant linear trend of cue weights ($F(1,35) = 27.7, p < 0.001$) indicating that participants learned the correct

relative ranks of cue weights. The main effect of LowP/HighP phase ($F(1,35) = 47.2, p < 0.001$), and the cue weight \times phase interaction ($F(2.4, 84.1) = 4.1, p = 0.01$) were also significant. Post hoc analyses revealed significant main effects for both LowP ($F(2.2,76.2) = 12.3, p < 0.001$) and HighP ($F(2.5,86.0) = 11.0, p < 0.001$) phases on cue weights. In addition, there was a significant effect of time pressure phase on subjective cue weights ($F(1,143) = 47.9, p < 0.001$), which was further confirmed by significant pairwise cue weight differences for cues 0.4 ($t(35) = 3.5, p = 0.001$), 0.6 ($t(35) = 4.0, p < 0.001$), and 0.8 ($t(35) = 4.5, p < 0.001$), demonstrating an overall down-weighting of cue values under severe time pressure. The main effect of phase on subjective cue weights as well as cue weight \times phase interaction observed in Experiment 2 (which was not found in Experiment 1) is indicative of robust changes in decision strategy as well as increased decision noise under more severe time pressure, which may account for the significant performance difference between the LowP and HighP phases reported above. In addition, we observed significant positive correlations between participants' subjective cue weights and ideal weights (Table 4).

3.2.2.3 Decision strategy model comparison and selection

In accordance with Experiment 1, in Experiment 2, the optimal cue-integration model (presented as Model 15) had the highest exceedance probability (0.92) when participants were under no time pressure (LowP phase). Under 500 ms time pressure (HighP phase), however, Model 5, using only the two most predictive cues, c_1 and c_2 ,

was the most likely strategy model, with an exceedance probability of 0.94 (Figure 5C). The control analysis using informative priors also revealed the same pattern of strategy shift (Figure 5D). Note that the Bayesian model comparison enables us to distinguish between whether the least valuable cue was assigned a small weight (perhaps as a consequence of generally lower cue weights in the speeded conditions) versus it being simply ignored altogether, and our results support the latter strategy. Again, switching from the optimal model to Model 5 results in reasonably high (“good-enough”) expected accuracy (87%) but with only half of the number of cues to consider (Figure 2B), which demonstrates the adaptive nature of cue usage under severe time pressure.

As in Experiment 1, we next sought to test whether the dropping of the two least valuable cue dimensions was strategic in nature, or whether participants were still trying to employ these cues but ran out of time. Accordingly, we analyzed performance in the HighP phase for trials where the two most important cues were identical between the two stimuli (32 trials/phase). Average RT for these trials was 369 ms, which was well below the response deadline of 500 ms (one-sample t -test: $t(35) = 61.5, p < 0.001$), and thus argues against the running-out-of-time hypothesis. Next, as for Experiment 1, we performed a set of analyses on the choice data to determine whether participants performed at chance-level on these trials in the HighP phase. Here, 33 out of 36 participants (92%) performed at chance level (binomial test, $p > 0.05$). Additional analyses based on BFs revealed that 26 out of 36 participants (72%) had a $BF > 1$

(favoring the hypothesis of chance-level performance), with 17 participants showing a significant effect ($BF > 3$). Next, we used BF analysis for testing for significant differences in performance between LowP and HighP phases using a paired t-test (Rouder et al., 2009). This analysis yielded a significant difference between the LowP ($mean = 59.5\%$, $SEM = 2.4$) and HighP ($mean = 49.3\%$, $SEM = 2.0$) phase performance ($BF = 67.4$, favoring alternative), and importantly, a one-sample t-test of the HighP phase performance showed a $BF = 5.2$, favoring the null hypothesis, that is, significant evidence for group performance being at chance-level in this condition. Thus, unlike in Experiment 1, in Experiment 2 we obtained unequivocal support for the hypothesis that the dropping of the weakest cues was strategic and categorical, rather than a function of participants running out of time to employ these cues. The reason for this significant finding in Experiment 2 compared to equivocal support in Experiment 1 might of course be the more severe degree of time pressure, but it might also be partly attributable to the fact that this analysis entailed a higher number of trials in Experiment 2 (32 per phase per participant, as compared to 16 in Experiment 1).

To address the apparent discrepancy between the results of subjective cue weights and strategy model selection in the HighP phase, as in Experiment 1 we here repeated the analysis using only the high performing participants (average performance of the learning phase $\geq 65\%$). Given that the magnitude of cue weights are highly dependent on individual task performance, a significant decrease in the percentage of

optimal responses along with the increase in number of no-response trials in Experiment 2 may have contributed to the overall down-weighting of estimated subjective cue weights, which was not observed in Experiment 1. The re-run of the Bayesian model selection analysis using only the least noisy half of our data set revealed the identical results as we obtained above (data not shown). This suggests that the overall decrease of subjective cue weight estimates in the HighP phase reflects an increased decision noise and a lower number of observations. The results of Bayesian model selection analysis, indicating a dropping of the weakest cues from the decision process, however, does not appear to be driven by this increase in noise.

In sum, the decision strategy analysis in Experiment 2 documented again that participants engaged in optimal cue integration when time pressure was low, but when faced with time pressure, they satisfied through cue discounting; dropping some information sources from the decision-making process. Importantly, the severe time pressure (500 ms) applied in Experiment 2 led participants to disregard the *two* least predictive cues, compared to the dropping of only the single least predictive cue under moderate pressure (750 ms) in Experiment 1.

3.2.2.4 Survey results

When asked about the relative importance of cue states within each cue dimension, the majority of participants were able to correctly identify the higher cue states with a mean correct response of 3.3 ($SD = 0.7$) out of 4. There was no significant

difference in these results between Experiment 1 and 2 ($t(77) = 0.7, p = 0.49$), indicating that regardless of experimental group, participants were able to explicitly identify the more predictive cue states in each cue dimension to a similar extent. Similar to Experiment 1, participants' ratings of relative rankings of cue dimensions showed a significant positive correlation with the subjective cue weights and the ideal weights, although this association was more modest for the ideal weights (Table 4).

When asked about the number of cue dimensions used during each phase, participants indicated that they used significantly fewer cues ($t(35) = 6.2, p < 0.001$) during the HighP phase ($mean = 1.6, SD = 0.7$) compared to the LowP phase ($mean = 2.5, SD = 0.7$). These results, however, were not significantly different from those of Experiment 1 for both the LowP phase ($t(77) = 0.37, p = 0.71$) and the HighP phase ($t(77) = 1.28, p = 0.21$). In addition, there were only a total of four (out of 36) participants, whose indicated subjective strategy models matched the objectively identified strategy models for both experimental phases. Thus, as in Experiment 1, we observed some evidence for explicit knowledge of cue values, but little evidence that participants had insight into their decision-making strategies.

3.2.3 Discussion

In Experiment 2, we successfully replicated and extended the results of Experiment 1. First, participants again displayed reliable statistical learning of the probabilistic cue values, and engaged in exhaustive cue integration when solving the

task in the absence of time pressure (LowP phase). Second, we obtained stronger evidence for the use of the “drop-the-worst” satisficing strategy, in that (a) increasing time pressure from moderate (Experiment 1) to severe (Experiment 2) led to a further trimming of cue usage (from the single to the two least valuable cues), and (b) in selectively analyzing trials where the two highest valued cues were identical between stimuli, we observed clear evidence for intentional ignoring of the weakest cues, and against the notion that participants simply ran out of time in trying to integrate all available cues. In sum, participants strategically disregarded the two least predictive cues to adapt to severe time pressure, although there was little evidence of explicit knowledge about this shift in cue usage in the survey data.

3.3 Experiment 3: Satisficing using non-compound cues under time pressure

Many of the previous studies that documented the use of lexicographic heuristics under satisficing pressures have used non-compound stimuli (e.g., Bröder, 2000; Dieckmann & Rieskamp, 2007; Payne, Bettman, & Johnson, 1988; Rieskamp & Otto, 2006), presenting cues independently, for example in the form of an information matrix. In Experiments 1 and 2, we took a departure from this approach and used compound cue stimuli, which integrated all cue dimensions to construct a single object. This raises the question whether our findings of a drop-the-worst satisficing strategy is for some reason unique to the case of integrated, compound cues. To investigate the generalizability of our “drop-the-worst” findings to non-compound cues, we therefore

conducted a third experiment, where we adopted the same type of probabilistic classification task that was used in the previous experiments but using non-compound cue presentation. Given that it is difficult to create unambiguous segregated cues out of some of the cue dimensions we employed in Experiments 1 and 2 (e.g., we cannot show a contour stimulus that does not also have a shape), we here opted to use a new set of cue symbols that lends themselves well to being presented in a cue matrix.

3.3.1 Methods

3.3.1.1 Participants

A new group of 48 volunteers was recruited through a separate HIT on Amazon MTurk with the goal of collecting 38 usable data sets. The same HIT approval criteria used in Experiments 1 and 2 were applied to ensure high data quality. Data from nine participants were excluded, one because of data loss, and eight due to chance-level performance (one-tailed binomial test on percent optimal responses in the initial 480 trials without time pressure; $p > 0.05$), leaving a total of 39 participants (mean age = 33.9 years, $SD = 8.2$, 21 – 55 years; 23 female, 16 male). Participants were compensated with \$4.00 upon completion of the experiment and an additional bonus ranging from \$0.50 to \$2.00 based on their performance.

3.3.1.2 Stimuli

The task stimuli consisted of four pairs of unique cues that comprised each of four cue dimensions (Figure 6A). Similar to the previous experiments, each cue

dimension was binary and consisted of two cue states that belong to the same category. For instance, the cue dimension of “weather” had two possible states, sunny and rainy. In addition, each cue dimension was assigned to a fixed location in a stimulus composed of a 2×2 matrix, so that the weather cue always appeared in the top left (*location 1*), the transportation cue in the top right (*location 2*), the activity cue in the bottom right (*location 3*), and the building cue in the bottom left (*location 4*). The same weights as in Experiments 1 and 2 (see Table 2), were randomly assigned to the different cue dimensions for each participant at the beginning of the experiment. Hence, except for the use of non-compound cues, the cue weight assignment and probabilistic nature of the stimuli remained identical to Experiments 1 and 2.

3.3.1.3 Procedure

Participants performed a probabilistic classification task, in which they compared two stimuli each composed of four cues and chose the one that is most likely to win (Figure 6B). Stimuli (194×194 pixels each) were presented on the left and the right side of the screen (window size: 1000×700 pixels) along the horizontal meridian, at an eccentricity of 250 pixels from a central fixation cross. Similar to the previous experiments, by exhaustively combining all possible cue states, we constructed 120 unique trials, in which the stimuli could differ in one to four cue dimensions. Stimuli were presented on the screen until a response was made or for the duration of an assigned response window of 3 s or 750 ms. Upon each valid response, probabilistic

feedback (determined by Eqns. 1-2), indicating “win” or “lose”, was displayed for 1 s, followed by a 1 s inter-trial interval. Participants were provided with identical task instructions as Experiment 1 (adapted to the new cue symbols), and earned 1 point for every winning trial. The total cumulative score at the end of the experiment was used to determine the amount of a bonus payment for each participant.

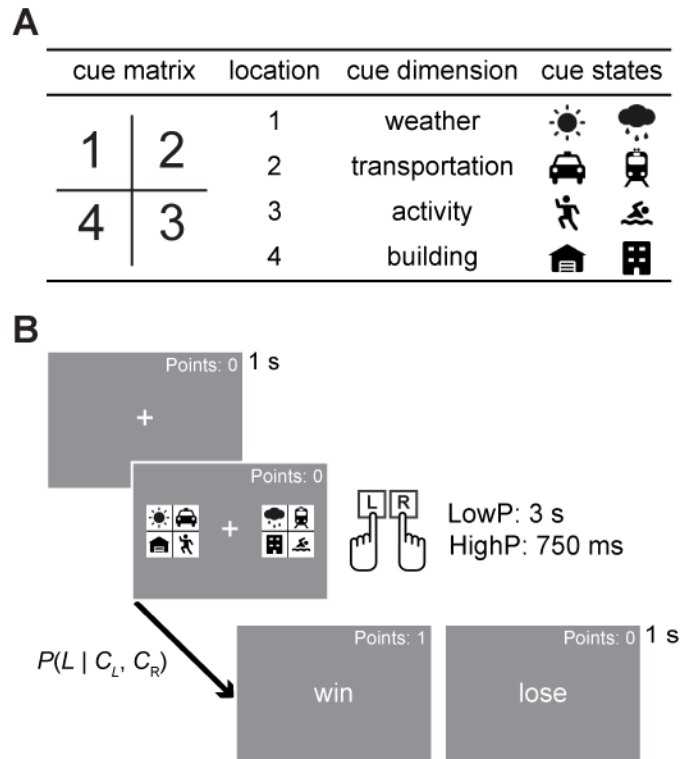


Figure 6: Cue Organization and Task of Experiment 3. (A) Non-compound stimulus organization using four independent cue dimensions (weather, transportation, activity, and building), each with binary states of varying weights. (B) Schematic of the non-compound multi-cue probabilistic classification task.

Again, to compare decision making strategies between conditions with and without time pressure, participants completed three phases sequentially, (1) an initial learning phase (sets 1 and 2), (2) an LowP phase (sets 3 and 4), and (3) a HighP phase (sets 5 and 6), each of which consisted of 240 trials grouped into 4 blocks of 60 trials. In both the initial learning period and the LowP phase, participants were given a maximum of 3 s to make a decision. This response window was chosen based on a pilot study, which used a 2 s response window for phases 1 and 2. In this study, participants achieved an average of 70.6% optimal responses at the end of the LowP phase (set 4), which was significantly lower than the LowP phase performance of Experiment 1 ($t(79) = 2.4, p = 0.02$). Therefore, to provide participants with sufficient time to observe the cues and thereby to facilitate the learning process, we extended the response window to 3 s. Following the LowP phase, participants performed an additional 240 trials of the HighP phase with 750 ms time pressure. Identical to the previous experiments, probabilistic feedback was provided throughout the entire experiment.

3.3.1.4 Survey

At the end of the experiment, participants completed a brief survey, which tested participants' explicit knowledge of cue values and their decision strategies in each phase. All the questions were identical to the previous version used in Experiments 1 and 2, except that the choices of cue dimensions and states were modified to match non-

compound stimulus conditions. Based on the answers to the survey questions, self-report strategy models were constructed.

3.3.1.5 Data analysis

Data analyses were carried out identical to Experiments 1 and 2.

3.3.2 Results

3.3.2.1 Task performance

Similar to Experiments 1 and 2, participants were able to gradually improve their performance throughout the initial learning phase (Figure 7A, Table 3). Even after the completion of the first 240 trials, participants continued to learn, indicated by significant difference in the percentage of optimal choices between set 2 and 3 ($t(38) = 3.2, p < 0.01$). We found no difference in performance between set 3 and 4 ($t(38) = 1.4, p = 0.16$), suggesting that performance had stabilized during the LowP phase. The percentage of optimal choices and RT were influenced by decision difficulty, which is characterized by the difference in sum of cue weights between the two stimuli, with more difficult decisions associated with decreasing performance (linear trend: $F(1,38) = 269.5, p < 0.001$) and increasing RT (linear trend: $F(1,38) = 43.4, p < 0.001$).

In the subsequent HighP phase (Figure 7A), participants experienced a time pressure of 750 ms, which was substantially below the mean RT of the LowP phase (see Table 3). Accordingly, both the rate of optimal choices ($t(38) = 8.9, p < 0.001$) and RT ($t(38) = 12.4, p < 0.001$) decreased significantly compared to the LowP phase. We

found no significant difference in performance between set 5 and 6 ($t(38) = 0.3, p = 0.74$). Regardless, the effect of decision difficulty on the percentage of optimal choices (linear trend: $F(1,38) = 134.9, p < 0.001$) as well as RT (linear trend: $F(1,38) = 4.2, p < 0.05$) was still present in the HighP phase. Similar to the previous experiments, there was no significant effect of gender on performance (LowP: $t(37) = 1.8, p = 0.08$; HighP: $t(37) = 1.7, p = 0.10$), and no significant correlation between age and performance (Pearson's correlation, LowP: $r = 0.1, p = 0.67$; HighP: $r = -0.1, p = 0.77$). Taken together, using non-compound stimuli, we obtained a comparable statistical learning profile to that observed in Experiments 1 and 2. That is, participants successfully learned to solve the task through trial-and-error learning over the course of the initial learning and the LowP phases. In the HighP phase, however, time pressure had a detrimental effect on performance, thus creating conditions for possible changes in decision strategy.

3.3.2.2 Subjective cue weights

To examine the effect of each cue dimension on participants' decisions, we again carried out logistic regression based on choice data in the LowP and HighP phases (Eqns. 3-6, Figure 7B). In accordance with our previous findings based on compound stimuli, an ANOVA with assigned cue weights ($w_{net,i}$) and experimental phase showed a significant main effect of cue weights ($F(2.0,75.5) = 26.8, p < 0.001$), which was characterized by a significant linear trend ($F(1,38) = 73.8, p < 0.001$). This suggests that participants had correctly learned the relative importance of the different cues even with

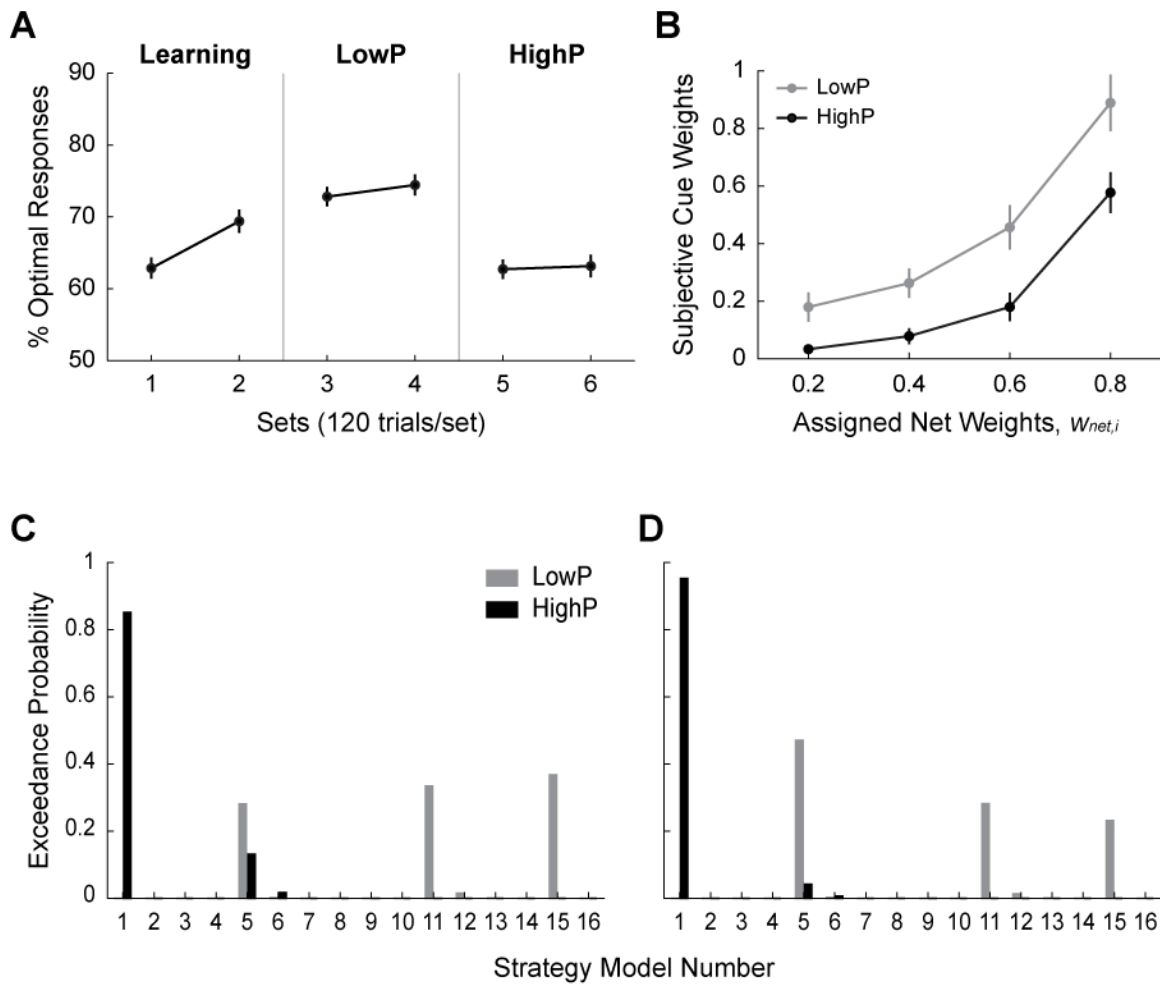


Figure 7: Task Performance, Subjective Cue Weights, and Bayesian Strategy Model Selection of Experiment 3. (A) Average percent optimal responses in the learning, LowP (3s response window) and HighP (750 ms response window) phases. (B) Average subjective cue weights as a function of LowP and HighP phases. Error bars are SEM. (C) Strategy model selection group results using informative priors. (D) Strategy model selection group results using informative priors.

non-compound stimuli. In addition, there was a significant effect of LowP/HighP phase ($F(1,38) = 48.0, p < 0.001$) with marginal cue weight \times phase interaction ($F(3,114) = 2.6, p = 0.06$). Post hoc test revealed that there was a significant overall down-weighting of cues in the HighP phase compared to the LowP phase (mean difference = 0.23, $SE =$

0.03, $p < 0.001$). This main effect of experimental phase, as in the case of Experiment 2, is indicative of a shift in decision strategy along with increased decision noise under time pressure, which also accounts for a significant decrease in performance in the HighP phase compared to the LowP phase. Moreover, comparison between subjective cue weights and weights produced by an ideal observer revealed a significant positive correlation (Table 4), which further confirms that, on average, participants learned to accurately rank the cue dimensions.

3.3.2.3 Decision strategy model comparison and selection

We again used variational Bayesian inference (Drugowitsch, 2013) and Bayesian model comparison (Stephan et al., 2009) to quantify participants' decision making strategies at an individual as well as at the group level (see Methods in Experiment 1 and Chapter 2). In the LowP phase, unlike Experiments 1 and 2, the model selection resulted in no clear winning model, with exceedance probability evenly shared among Model *Opt* (0.37 for uninformative and 0.23 for uninformative priors), Model 11 (0.33 for uninformative and 0.28 for uninformative priors), and Model 5 (0.28 for uninformative and 0.47 for informative priors) (Figure 7C, D; see also Figure 2A). This suggests that unlike the comparison between compound stimuli, when subjects are comparing non-compound stimuli, they are more inclined to consider a lower number of cues even when the time pressure is fairly low. It is, however, worth noting that this reduction of cue usage appears to be strategic and adaptive, and conforms to the drop-the-worst

strategy: compared to the optimal model (Model *Opt*), Model 11 ignores the least important cue with 97% expected accuracy, and Model 5 ignores the two least important cues with 87% expected accuracy (Figure 2B), both of which result in only small drop-offs in performance. By contrast to the LowP phase, performance in the HighP phase, with a 750 ms response window, was best captured by strategy Model 1. This model, which only uses the most valuable cue dimension and thus ignores the three weakest cues, was the most likely strategy model with an exceedance probability of 0.85 (uninformative priors) and 0.95 (informative priors) (Figure 8B, C). These results conceptually replicate those of Experiments 1 and 2, in that they suggest that under higher time pressure participants reduced the search space following a drop-the-worst principle, here using only the single highest cue, c_1 , to arrive at their decisions.

Next, as in Experiment 1 and 2, we aimed to tackle the question whether the observed shift in strategy from the LowP to HighP phase is strategic in nature or participants simply didn't have enough time to process additional cues. Accordingly, we analyzed performance in trials when the most important cue was identical between the two stimuli (112 trials/phase). Average RT for these trials in the HighP phase was significantly below the response deadline of 750 ms ($mean = 519$ ms; one-sample t-test: $t(38) = 18.6, p < 0.001$), which is in contradiction with the running-out-of-time hypothesis. To further examine the validity of the model selection result, we conducted a set of additional analyses to test whether performance on these trials is at chance in the

HighP phase. This analysis revealed that 35 out of 39 participants (90%) performed no different from chance level (binomial test, $p > 0.05$), indicating that the great majority of participants guessed when the most important cue did not discriminate between the two stimuli. Additional analysis based on BFs revealed that 35 out of 39 participants (90%) had a $BF > 1$ (favoring the hypothesis of chance-level performance), with 28 participants showing a significant effect ($BF > 3$). At the group level, we performed a Bayesian paired t-test (Rouder et al., 2009) to compare average performance between the LowP ($mean = 64.7\%$, $SEM = 1.7$) and HighP ($mean = 51.5\%$, $SEM = 1.1$) phases. This t-test yielded a highly significant difference ($BF = 4.4 \times 10^7$, favoring alternative), and an additional one-sample t-test of the HighP phase performance revealed $BF = 2.5$ (standard t-test: $t(38) = 1.3$, $p = 0.19$), thus favoring the null hypothesis that group performance is at chance level on these trials under time pressure. In sum, using non-compound cues, we again obtained evidence supporting that dropping of the weakest cues is a strategic decision to cope with time pressure.

In summary, the decision strategy analysis in Experiment 3 revealed that when multiple cues were presented in a spatially segregated manner, participants were more prone to consider a lower number of cues and to adopt a noncompensatory strategy even without high time pressure. Under 750 ms time pressure, participants satisfied by dropping the weakest three cues from their decision making process, considering only the highest cue and resorting to guessing when this cue did not discriminate between

the two cue matrices. In addition, Bayesian model comparison showed that only three participants out of 39 (one in the LowP and two in the HighP phase) preferred the Take-the-Best model (Model TTB) over the optimal cue integration model, suggesting that the majority of participants did not adopt a sequential, lexicographic strategy. Rather, participants seem to strategically disregard the least predictive cues and focus only on a small subset of cues to guide their decisions.

3.3.2.4 Survey results

Similar to Experiments 1 and 2, most participants were able to correctly identify the cue states with a higher predictive value within each cue dimension with a mean correct response of 3.7 ($SD = 0.6$) out of 4. Participants' explicit ratings of relative importance of cue dimensions showed a significant positive correlation that matched very closely to their inferred subjective cue weights as well as the ideal weights (Table 4).

When asked about the number of cue dimension usage in each phase, participants indicated that they used significantly fewer cues ($t(38) = 12.5, p < 0.001$) in the HighP phase ($mean = 1.5, SD = 0.5$) compared to the LowP phase ($mean = 2.8, SD = 0.7$). Comparison between subjective strategy models, which were defined by each participant's response to the survey questionnaire, and objectively identified strategy models revealed that only a total of five out of 39 participants were able to correctly identify their strategies for both experimental phases. Thus, in accordance with

our previous findings, we obtained some evidence for explicit knowledge of relative importance of cue dimensions and states, but only ambiguous support for any insight into the decision-making strategies employed in the LowP and HighP phases.

3.3.3 Discussion

In Experiment 3, we sought to examine whether our findings of a “drop-the-worst” satisficing strategy based on compound cues can be generalized to a probabilistic classification task using non-compound cues. We found that similar to the previous experiments, participants were able to learn the probabilistic cue values through trial-and-error. In the LowP phase, however, only a small portion of participants appear to engage in exhaustive cue integration and many participants adopted sub-optimal strategies. Importantly, though, those sub-optimal strategies conformed to a drop-the-worst pattern, with participants using only the two or three highest value cues to make decisions. In the 750 ms HighP phase, participants further reduced the search space and considered only the single most important cue. Our model selection results provided no support for widespread use of the Take-the-Best strategy. In sum, using non-compound cues, we successfully replicated our main results from Experiments 1 and 2, confirming the reliable use of the “drop-the-worst” strategy.

The difference in the results of the LowP phase from our previous findings may have been caused by participants’ strategic decision to decrease the number of cues to consider to reduce their effort in the presence of low, but not negligible, time pressure.

As reported in Table 3, participants took longer to arrive at decisions in the learning and LowP phases in Experiment 3 compared to the experiments using compound cues, suggesting that comparing non-compound multi-cue stimuli may be more effortful than comparing compound stimuli. In addition, previous studies that reported the use of the weighted additive strategy using non-compound cues have provided participants with ample time to observe and compare the cues before making a choice (e.g., Glöckner & Betsch, 2008; Pachur & Olsson, 2012). However, our task imposed a fixed response deadline of 3 s during the learning and LowP phases, which may have already encouraged the use of sub-optimal strategies. Regardless, rather than adopting the lexicographic, Take-the-Best heuristic that relies on a single discriminating cue through sequential search (Gigerenzer & Goldstein, 1996; e.g., Payne, Bettman, & Johnson, 1988; Rieskamp & Hoffrage, 2008), or dropping random cues from their decision process, participants systematically integrated the two or three most valuable cues to arrive at decisions in the LowP phase.

In the HighP phase, participants strongly relied on the single highest cue to solve the task. Through a set of control analyses, we were able to confirm that participants performed at chance when faced with trials when the highest cue did not discriminate between two stimuli. It is, however, difficult to clearly distinguish whether participants were trying to employ the Take-the-Best strategy but did not have enough time to evaluate the second highest cue or they strategically disregarded the rest of the cues

altogether. Given the saliency of the cues as well as RT data, it is unlikely that participants did not have enough time to process the second cue visually but it is difficult to completely rule out the possibility of the use of lexicographic strategy under high time pressure.

3.4 General discussion

Decision-making in everyday life is beset by uncertainty due to noisy and incomplete information, and limited available decision time and cognitive resources. It has long been held that humans adaptively use satisficing strategies that can simplify the decision-making process to save time and cognitive effort to arrive at good-enough solutions (Gigerenzer & Goldstein, 1996; Shah & Oppenheimer, 2008; Simon, 1955, 1956, 1990), and some previous studies have begun to investigate the nature of shifts in decision-making under time pressure. The goal of the present study was to expand this literature by creating a protocol that emphasized two key aspects of certain real-life, high-stakes decision making, namely uncertainty of cue values combined with severe time constraints, i.e., in split-second decision making. Specifically, we examined how people solve a probabilistic classification task under varying time pressure, thereby connecting the often disparate literatures on learning and decision making. Using compound cues, we found that, under low time pressure (LowP phase), participants were able to correctly weight and integrate all available cues to arrive at near-optimal decisions. With increasing time pressure (HighP phase), however, participants shifted

their decision strategies by dropping cues from the information-integration process. Importantly, this selective discounting of a subset of cue information was clearly strategic and adaptive, in that participants specifically dropped the one (Experiment 1, moderate time pressure group with 750 ms response window) or two (Experiment 2, severe time pressure group with 500 ms response window) least informative cues from the decision-making process. Moreover, control analyses confirmed that disregarding of the least valuable cue(s) was not an expression of simply running out of time during an attempt at integrating those cue values. Rather, the weakest cues seem to have been categorically excluded from the decision process under high time pressure. We replicated these results using non-compound cues (Experiment 3), which demonstrated that under 750 ms time pressure, participants dropped the three least informative cues, utilizing only the best cue to make decisions. Post-experiment survey results suggested that participants had at least some explicit knowledge concerning the informative value of the cues but lacked insight into the decision strategies they adopted. Our results thus document, and quantify, adaptive shifts in decision strategies under uncertainty to compensate for limited decision time. Specifically, we showed that participants engaged in adaptive cue discounting, ignoring the least valuable information sources, a satisficing variant we here call “Drop-the-Worst”.

Our discovery of the decision strategy that drives satisficing under uncertainty and high time-pressure is significant, because this knowledge, in principle, renders split-

second human choices (e.g., in traffic or combat) predictable, which in turn can inform the optimal design of safety measures and/or autonomous agents that interact with humans. The fact that the nature of satisficing in split-second decision making cannot be anticipated on the basis of strategies uncovered in more slow-paced environments becomes clear when we contrast the present findings with those of several prior investigations using cue-based search paradigms. These studies examined the effect of time pressure on strategy selection in slow-paced (> 15-50s) scenarios and have found that under low time pressure, similar to the present study, a strategy model that integrates all available information (e.g., weighted additive strategy) performed best at predicting participants' choices (Rieskamp & Hoffrage, 2008). However, under higher time pressure, people adapted by using a simple lexicographic heuristic (i.e., a "Take-the-Best" heuristic for a binary choice problem; Gigerenzer & Goldstein, 1996) (Payne et al., 1988, 1996; Rieskamp & Hoffrage, 2008), which looks for "one clever cue" to base decisions on. These conclusions are partly drawn from the information search structure of tasks that allow subjects to inspect only one piece of information at a time (Payne et al., 1988; see Glöckner & Betsch, 2008 for discussion). Specifically, Take-the-Best requires searching through cues in descending order of cue validity until the decision maker finds the first (and the "best") cue that discriminates between the alternatives and such sequential search paradigm can strongly encourage subjects to adopt one-cue heuristic under time pressure. Indeed, when cues were presented all at once, people were able to

integrate cue information using compensatory strategies relatively quickly (Glöckner & Betsch, 2008; Pachur & Olsson, 2012), which is in accordance with our findings under no time pressure. Our present findings further extend the previous work on bounded rationality by documenting, for the first time, that under much higher time pressure that enforces split-second decision making, subjects instead adopted a “Drop-the-Worst” heuristic, whereby the least valuable cues are simply ignored (or discounted) altogether.

This difference in satisficing strategy likely represents an adaptive shift in information integration to accommodate the less certain nature of cue information and much higher time pressure in the present task. First, in order to implement the Take-the-Best approach, it is crucial for the decision maker to be confident in their knowledge of the exact rankings of cue validities; in fact, the majority of previous work on this heuristic has provided participants with explicit cue validities on each trial (e.g., Payne et al., 1988, 1996; Bröder, 2000, 2003; Rieskamp and Otto, 2006; Dieckmann and Rieskamp, 2007; Rieskamp and Hoffrage, 2008). Moreover, it has been shown that experts who have better knowledge of cue validities than novices are more likely to adopt the Take-the-Best heuristic (Garcia-Retamero & Dhami, 2009; Pachur & Marinello, 2013). By contrast, participants in our study were not provided with explicit cue weights and had to infer the correct cue ranking through trial-and-error learning, likely rendering them less confident about the exact rankings. In addition, due to a large number of possible cue combinations, each associated with probabilistic feedback, our

experimental design likely prevented subjects from making memory-guided decisions (see Juslin, Olsson, & Olsson, 2003). Given this uncertainty about cue information and the complexity of the task structure, it may be more adaptive to first set a satisfactory cutoff level, integrate cues in an order that is most likely to reach this cutoff, and then make a decision without evaluating all available cues (Shah & Oppenheimer, 2008), especially under high time pressure. This strategy reduces cognitive effort by integrating less information and choosing an alternative that is simply “good-enough” whereas Take-the-Best reduces effort by examining only one cue at a time, but it can require searching for multiple cues if the first cue does not discriminate between alternatives (Shah & Oppenheimer, 2008).

Second, given that Take-the-Best requires serially searching through the cues in order of their validity until one finds the first discriminating cue, in some cases this strategy can require more time than the Drop-the-Worst strategy, in particular when the least informative cue (Experiment 1) or cues (Experiments 2 and 3) are the only cues differentiating between the two stimuli. Therefore, if there is high pressure to make a quick decision with some uncertainties in cue validities, lowering the satisfactory cutoff level and hence, focusing only on a fixed subset of cues in the decision process may be a more efficient way to save cognitive effort and achieve good-enough accuracy.

Accordingly, our findings document that participants did not employ the Take-the-Best strategy (presented as Model 16), even when the cues were presented separately, but

instead used a Drop-the-Worst approach, where only a subset of cues with the highest validities (weights) was considered, and random guesses were made if those cues did not differ between the two stimuli. Although Bergert & Nosofsky (2007)'s response-time approach demonstrated more use of Take-the-Best strategy compared to the weighted additive strategy using compound cues, the use of deterministic feedback during the training phase as well as a task structure that yielded comparable results between the use of Take-the-Best and the optimal strategies may have influenced the participants to rely on a single cue heuristics (see Juslin et al., 2003). It is also important to note that the observed strategy shift to noncompensatory heuristics can be induced by the effect of learning and experience over the course of the experiment (Garcia-Retamero & Dhimi, 2009; Johnson & Payne, 1986; Pachur & Marinello, 2013; Rieskamp & Otto, 2006). Our results are no exception to this observation, and because participants were provided with probabilistic feedback throughout all phases of the experiment, and the HighP phase always followed the LowP phase, it is possible that continued learning may also have played a role in inducing the switch to the Drop-the-Worst strategy.

One intriguing question arising from the present protocol is whether multiple cues are processed in a serial manner (i.e., integrating cues serially from the most important to the least using an additive rule) or in a parallel fashion. For instance, one could imagine evidence for all cue dimensions being accumulated simultaneously in a drift-diffusion type model (Smith & Ratcliff, 2004), where the drift rate of evidence may

increase as a function of cue weight, corresponding to an effect of attention. With our current paradigm, however, it is difficult to distinguish whether behavior stems from a serial or parallel cue processing strategy, since equivalent predictions seem to follow from both models. For instance, in the LowP phase of both experiments, we found significant linear trends of RT as a function of cue weights when only a single cue was different between two stimuli (Exp. 1: $F(1,42) = 22.1, p < 0.001$; Exp. 2: $F(1,35) = 4.1, p < 0.05$). In other words, participants took longer time to make a choice when a less valuable cue was the only distinguishing feature between the two stimuli, in spite of the fact that this single difference should be perceptually quite salient. This might be interpreted as support for the serial processing model, as participants may have evaluated cues serially in order of importance, thus leading to slower response times for less valuable cues even when only that cue distinguishes the two stimuli. However, this result could equally be driven by decision difficulty, with more difficult trials (i.e., with less cue value difference between stimuli) requiring more time to make a choice. Taken together, the precise manner in which multiple cues are integrated and how the shift in decision strategy is instantiated computationally and neurally under time pressure represents an exciting challenge for future research.

Several previous studies have explored ways of identifying strategies that participants use in solving the classic weather prediction task (Gluck, Shohamy, & Myers, 2002; Lagnado, Newell, Kahan, & Shanks, 2006; Meeter, Myers, Shohamy,

Hopkins, & Gluck, 2006; Meeter, Radics, Myers, Gluck, & Hopkins, 2008; Speekenbrink, Lagnado, Wilkinson, Jahanshahi, & Shanks, 2010), which has close parallels with the present protocol. Gluck et al. (2002) introduced a method using a least-mean-squared-error measure that compares each participant's data to the ideal response profiles constructed from three different strategy models, the multi-cue (optimal model), one cue, and singleton strategies. Hence, a strategy that resulted in the lowest error was defined to be the best-fit model. This model-based approach was later extended by using Monte Carlo simulations that can be harnessed to infer switches from one strategy to another over the course of the experiment (Meeter et al., 2006). Another approach in identifying an individual's strategy is to use "rolling regression" methods that estimate subjective weights of each cue through a moving window of consecutive trials, which can be applied to characterize how the learning occurs during the course of the task (Kelley & Friedman, 2002; Lagnado et al., 2006). Meeter et al. (2008) demonstrated that both strategy analysis and rolling regression analysis results in a more or less equivalent ability to predict responses.

We here applied a new analytical approach that captures these key aspects of prior strategy analyses and enabled us to infer both subjective cue values and the manner in which these were combined to reach decisions. Specifically, we inferred how the cues were weighted in each phase of the experiment using logistic regression, and identified the most likely strategy employed by using Bayesian model comparison at an

individual as well as the group level. Hence, instead of simply categorizing strategies into the number of cues used (cf. Gluck et al., 2002; Lagnado et al., 2006; Meeter et al., 2008), we explored an exhaustive set of plausible strategy models to identify the exact cues used and their relative importance in making decisions. For instance, even if a given participant adopted a sub-optimal strategy that resulted in poor performance, we were nevertheless able to infer the most likely underlying cue structure that the participant may have developed throughout the task. The present analysis approach therefore may have great potential for enhancing the inferences drawn from future studies of statistical learning.

Moreover, the present results, in accordance with previous studies using the classic weather prediction task (Gluck et al., 2002; Lagnado et al., 2006; Meeter et al., 2006, 2008; Speekenbrink et al., 2010), highlight the fact that there is likely considerable variability in participants' cognitive strategies in probabilistic decision making scenarios. Previous studies reported that people tend to start with a simple strategy using a single cue (i.e. singleton strategy) but switch to an optimal multi-cue strategy toward the end of the task. Although we did not assess changes in strategies across multiple time points within each experimental phase (which would require a yet higher trial count), our performance data suggest that participants learned to use the cues in a near-optimal manner as they learned the cue characteristics throughout the initial learning and the

LowP phase. In addition, our data on the subjective cue weights indicate that the majority of participants learned the correct relative importance of cue weights.

Another intriguing model for solving a multi-cue probabilistic classification task, especially under high time pressure, is a mixture strategy model where participants some of the time integrate cues in an optimal manner and some of the time decide randomly. We had considered the possibility of this type of mixture model, but rejected it based on the fact that this mixture model predicts psychometric curves that do not saturate to 100% accuracy when the sum of evidence is at the maximum value, unless when the guessing parameter is set to zero. By contrast, the sub-optimal cue usage models such as Models 11 and 5, predict that accuracy will suffer more when the evidence is low but the performance will eventually saturate toward 100%. Note that this pattern also holds for predictions at very low guess rates ($g \leq 0.2$), which are in a reasonable range if we consider decision noise and imperfect knowledge of cue weights. Hence, even though participants are making more errors (possibly random guesses) under time pressure, the performance of a majority of participants, especially the high performers, is scaled by the sum of evidence, reaching the peak when the evidence is at maximum. The shape of these performance curves corresponds to the mixture model with $g \leq 0.2$, indicating that performance in the HighP phase does not reflect a random increase in the guess rate, but is reflective of a systematic lapse in decision making dependent on the sum of evidence. In addition, in case of severe time pressure

(Experiments 2 for compound cues and 3 for non-compound cues), we showed that performance is at chance when the low-value cues are the only informative cues. Therefore, although the mixture model can account for some proportion of our observations, since our goal was to uncover patterns of cue usage under varying degrees of time pressure, we did not include the model in our strategy model comparison.

Finally, there have been mixed findings regarding whether participants have explicit knowledge about how they solve the classic weather prediction task (Gluck et al., 2002; Knowlton et al., 1994; Lagnado et al., 2006; Newell, Lagnado, & Shanks, 2007). This knowledge can be divided into two different insights that might not necessarily coincide: (1) insight into the cue structure of the task, and (2) insight into strategy use (Lagnado et al., 2006). Gluck et al. (2002) reported that there was little or no evidence that participants had explicit insight about the task structure or their strategy use based on their post-experiment questionnaire. On the other hand, Lagnado et al. (2006), using more detailed and frequent measures, found that having accurate knowledge of the task structure and self-insight is necessary for achieving optimal performance. Our post-experiment survey results seem to indicate that the majority of subjects in the present study had some explicit knowledge about the cue structure, but that this knowledge did not necessarily lead to having accurate insight into strategy use.

In summary, the present study characterized the nature of adaptive shifts in decision strategies in an uncertain and fast-paced environment. By combining a

probabilistic classification task under varying time pressure with new analytical approaches to quantifying decision models, we showed for the first time that, when forced to make split-second decisions in an uncertain environment, participants strategically discount the least valuable piece(s) of information, providing novel evidence for a “Drop-the-Worst” cue satisficing decision-making strategy.

4. Satisficing Decision Making Under Anticipatory Anxiety

In Chapter 3, we characterized satisficing decision making under externally-induced pressure. Specifically, in response to increasing time pressure, we showed that participants adaptively shift their choice strategies by employing the Drop-the-Worst heuristic. To examine the effect of internally-driven pressure on decision strategies, in the current chapter, we introduced a new “threat-of-shock” condition to induce anticipatory anxiety (for reviews, see Grupe & Nitschke, 2013; Robinson, Vytal, Cornwell, & Grillon, 2013). While an ability to make effective decisions under uncertainty and threat, such as in military combat or other high-stakes scenarios, is particularly critical for survival, surprisingly little is known about how people utilize satisficing decision strategies under anxiety. The present study, therefore, aims to characterize the manner in which complex decisions involving multiple probabilistic pieces of information, are influenced by anticipatory anxiety.

Anticipating a negative event induces a state of anxiety that can influence a large array of cognitive abilities (for a review, see Robinson et al., 2013). The threat-of-shock paradigm is a robust, well-established technique, which uses infrequent, randomly occurring electrical shocks to induce sustained anxiety in healthy participants (Schmitz & Grillon, 2012). Here, we specifically focused on task non-contingent, unpredictable, and uncontrollable threats, which trigger an anticipatory anxiety response characterized

by an increase in skin conductance level (SCL) (Rhudy & Meagher, 2000). Such task-unrelated anxiety can result in both detrimental and facilitating effects on a broad range of executive functions such as attentional control (e.g., Eysenck, Derakshan, Santos, & Calvo, 2007; Hu, Bauer, Padmala, & Pessoa, 2012), working memory (e.g., Shackman et al., 2006; Vytal, Cornwell, Letkiewicz, Arkin, & Grillon, 2013), and decision-making (e.g., Clark et al., 2012; Keinan, 1987). Importantly, in line with predictions from the satisficing decision theory (Simon, 1955, 1990), Easterbrook (1959) proposed that behavioral changes observed under anxiety are driven by “attentional narrowing”, which reduces the range of cue utilization and thereby enhances focus on the most important task-relevant cues. This attentional narrowing effect has often been observed through enhanced early perceptual processing of salient cues (e.g., Cain, Dunsmoor, LaBar, & Mitroff, 2011; Cornwell et al., 2007), although the extension of this effect to subsequent decision-making processes remains unclear (cf., Shackman, Maxwell, McMenemy, Greischar, & Davidson, 2011). Additionally, whether selective cue processing under anxiety is executed in an *adaptive* manner – a core characteristic of satisficing – via facilitating cue evaluation in order of their importance during decision-making is currently unknown.

The influence of anxiety on probabilistic decision-making has been rarely examined, but prior reports based on acute stress induced by the cold pressor test offer some useful insights. During a two-step reinforcement learning task, stress attenuated

the contribution of model-based, goal-directed learning but not model-free, habitual decision-making (Otto, Raio, Chiang, Phelps, & Daw, 2013). Similarly, during probabilistic classification learning, a pressor-induced stress group was more likely to engage in a striatum-dependent habit learning strategy whereas the control group tended to adopt a hippocampus-dependent declarative strategy (Lars Schwabe & Wolf, 2012). Importantly, this difference was characterized behaviorally by a significant reduction in explicit task knowledge under stress without any observable difference in overall classification accuracy, suggesting that traditional decision performance measures (i.e., accuracy or response time) may not be strongly reflective of the underlying changes in decision strategy. Taken together, consistent with the multiple memory systems theory (Knowlton et al., 1996; Poldrack et al., 2001), anxiety may shift decision-making to the habitual system to facilitate decision processes without sacrificing overall performance, which is often reflected by reduction of explicit task knowledge (for a review, see Schwabe & Wolf, 2013).

To investigate whether this adaptive shift to satisficing decision-making generalizes to decisions under anticipatory anxiety, or whether other satisficing strategies emerge under these conditions, we combined the multi-cue probabilistic classification task with the threat-of-shock paradigm, while recording participants' SCL throughout the experiment. Specifically, we primarily based our analyses on variational Bayesian inference, which yields better sensitivity in characterizing participants' implicit

decision strategies than traditional performance metrics. This allowed us to detect and compare strategies utilized during neutral versus threat-of-shock conditions, both at the within- and between-subjects level, as a function of individual differences in the effectiveness of the threat manipulation as quantified by SCLs. We predicted that anticipatory anxiety would (1) trigger shifts in strategy, promoting the use of satisficing heuristics under high arousal, and (2) induce reduction in declarative knowledge of the task-related information.

4.1 Methods

4.1.1 Participants

Sixty right-handed individuals from the Duke University community participated in return for course credits. Since we planned to investigate individual differences based on sub-grouping of participants by SCL response, we targeted the sample size of ~20 participants/group based on previous studies on the effect of threat-of-shock on decision-making (e.g., Hu et al., 2012; Robinson, Overstreet, Charney, Vytal, & Grillon, 2013). Three participants were excluded from the analysis, two due to a lack of measurable skin conductance responses to shock and one due to an equipment malfunction, leaving 57 participants (38 females; mean age = 19.1 years; range = 18-21 years). All participants provided informed consent in accordance with Duke University Institutional Review Board guidelines and reported no history of neurological or psychiatric illness, and no current psychoactive medication use.

4.1.2 Stimuli

The task stimuli consisted of 16 compound cues as outlined in Chapter 2 (Figure 1; Tables 1-2). Additionally, in the threat-of-shock phase (see below), five shock stimuli were administered to the right wrist for 6 ms. Shock intensity was calibrated for each participant using an ascending staircase procedure to be perceived as “highly annoying but not painful” (mean = 32.8 V, SD = 8.4; Dunsmoor, Mitroff, & LaBar, 2009).

Throughout all phases of the experiment, SCL was measured with Ag-AgCl electrodes placed on the middle phalanx of the second and third digits of the left hand. Shock delivery and SCL were recorded using a MP-150 system connected to AcqKnowledge software (BIOPAC Systems, Goleta, CA).

4.1.3 Procedure

Prior to the experiment, participants completed both State and Trait parts of the State-Trait Anxiety Inventory (STAI; Spielberger, 1983). The multi-cue probabilistic classification task (see Chapter 2; Figure 8A) was divided into three phases: an initial learning (IL) phase, followed by a neutral (no-shock, NS) phase, and a threat-of-shock (TS) phase. Each phase was comprised of 240 trials, which were randomized and grouped into 4 blocks of 60 trials per phase, with short breaks in between. Specifically, blocks 1 and 2 were comprised of the full set of all 120 possible stimulus pairs, followed by blocks 3 and 4, which consisted of the mirror images of the first set. On each trial, participants were presented with a stimulus pair and asked to indicate their choice via

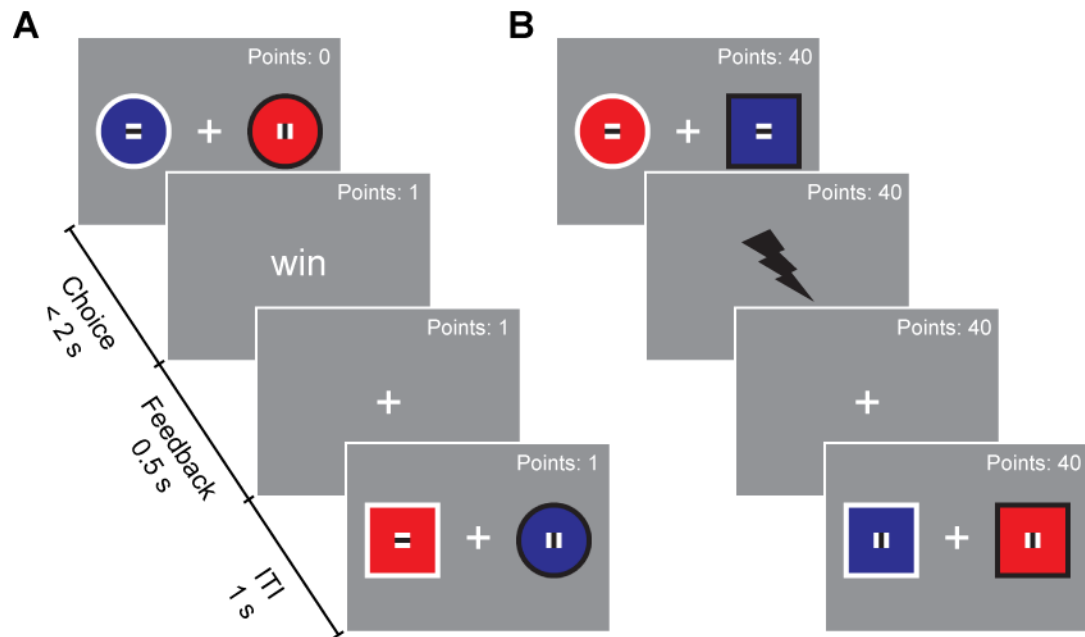


Figure 8: Schematic of the Experimental Procedure. (A) Schematic of the standard multi-cue probabilistic classification task in the initial learning and neutral phases. (B) Schematic of the threat-of-shock phase; participants additionally experienced randomly occurring shocks at the time of feedback along with a shock symbol.

keyboard button press using their right index (left stimulus) and middle (right stimulus) fingers, within a 2 s response window. Once a valid response was made, feedback was presented for 500 ms, which was followed by a 1 s inter-trial-interval.

Enhanced startle responses can be triggered by placement of shock electrodes on participants' wrist alone (Grillon & Ameli, 1998), indicating that this procedure, even in the absence of threat, could induce a general increase in baseline anxiety. To avoid this kind of spillover effect of the shock condition contaminating decision behavior in neutral blocks, the treat-of-shock manipulation was only introduced to participants upon completion of the IL and NS phases, when we performed the shock intensity calibration.

Then, subjects were informed that a few shocks would be delivered throughout the phase, independent of their performance. During the TS phase, five shock trials were included at random time points, in addition to 240 valid trials, with a maximum of two shocks per block. Shocks were administered at the time of feedback along with a visual shock symbol (Figure 8B), and these trials were excluded from data analyses.

Additionally, upon finishing the TS phase, participants completed a post-experimental survey assessing their explicit task knowledge (Appendix A).

4.1.4 Skin conductance level (SCL) analysis

SCLs were calculated as the mean response over each block using the Autonomate toolbox (Green, Kragel, Fecteau, & LaBar, 2014). To minimize the influence of shock-triggered responses on SCLs, time epochs containing shock delivery and three trials following the shock (~ 8 s) were removed from the analysis. For each participant, SCLs were range-corrected based on subject-specific minimum and maximum SCLs observed during the experiment to account for individual differences (Lykken & Venables, 1971), and averaged across blocks per phase. Then, changes in SCL between the NS and TS phases were used ($\Delta SCL = SCL(TS) - SCL(NS)$) to quantify the effect of threat-of-shock in inducing anticipatory anxiety.

4.1.5 Participant clustering

To characterize how anticipatory anxiety differentially influenced decision strategies employed during the TS phase, we partitioned participants based on ΔSCL .

Specifically, participants were clustered into low ($n = 14$), mid ($n = 20$), and high ($n = 23$) responders using the Ckmeans.1d.dp algorithm (Wang & Song, 2011). This algorithm yields optimal one-dimensional clustering by minimizing within-cluster distances from each data point to its corresponding cluster mean, and therefore, provides an objective means of reliable and natural grouping of participants. Note, however, that using a tercile split does not change the overall results we report.

4.1.6 Behavioral performance analysis

Performance data analyses were based on correct choices favored by the probability of winning, as described in Chapter 2.

4.1.7 Decision strategy model selection

To identify decision strategies that individual participants may have applied to solve the task, we explored a large set of plausible decision models, accounting for every possible combination of cue dimension usage, as outlined in Chapter 2 (Figure 2A). Since relatively uninformative cue dimensions, such as c_3 and c_4 , do not strongly influence final outcomes, the expected accuracy of some sub-optimal strategy models is close to optimal (Figure 2B). Given the probabilistic nature of the task, human participants are prone to have higher decision noise and to make more errors than the ideal observer, which could make the performance outcomes of these sub-optimal and the optimal models indistinguishable. Therefore, to achieve better sensitivity, we performed model comparison using variational Bayesian inference (Drugowitsch, 2013)

with $a_0 = 0.55$ and $b_0 = 0.79$ as hyper priors (Eqns. 3 – 8; see Section 2.2.4). Changing these values to correspond to relatively uninformative hyper priors ($a_0 = 0.01$ and $b_0 = 0.0001$) did not change the overall results of model comparison.

To summarize the model comparison results at the group level, we employed a Bayesian model selection procedure as described in Chapter 2 (Rigoux et al., 2014; Stephan et al., 2009). Specifically, to rule out the possibility that differences in observed model frequencies could be driven by chance, we calculated *protected* exceedance probabilities and associated Bayesian omnibus risks (BORs). The protected exceedance probability reflects the belief that a model, m , provides a more likely explanation of the data than any other model, beyond chance, given the marginalized likelihoods obtained from the variational Bayesian inference procedure above. The BOR evaluates the probability that the observed difference in model frequencies occurred by chance, i.e., the null-hypothesis whereby all model frequencies are equal. Hence, BOR can be considered as analogous to a classical p-value.

4.1.8 Explicit knowledge of the cue structure

Post-experimental survey responses were analyzed to examine participants' explicit knowledge about the cue structure and strategies employed in each phase. Spearman's rank-order correlation was used to quantify explicit knowledge by comparing each participant's survey rating with assigned weights, $w_{net.i}$ ($i = 1, \dots, 4$; Table 2), as well as inferred subjective cue weights, w_{opt} (Eqns. 3-7). The correlation

between the survey rating and assigned weights, r_w , describes how well participants learned the true underlying cue structure. More importantly, the correlations between the survey rating and subjective cue weights, $r_{sw,NS}$ and $r_{sw,TS}$, indicate the degree of correspondence between participants' explicit knowledge of the cue structure versus the actual cue dimension weighting employed in the NS and TS phases. For further analyses, all correlation coefficients were Fisher z-transformed to achieve normal distribution.

To further examine the relationship between participants' cue usage and their explicit knowledge across conditions, we compared z-transformed correlation coefficients between the NS and TS phases. Specifically, the concordance between actual strategy and declarative knowledge was measured by changes in subjective cue weights from the NS to TS phase in relation to participants' explicit survey ratings ($\Delta r_{sw} = z(r_{sw,TS}) - z(r_{sw,NS})$). Hence, a positive Δr_{sw} indicates a greater concordance between participants' applied cue weighting and their expressed explicit knowledge in the TS compared to the NS phase. Conversely, a negative Δr_{sw} suggests an increased mismatch between the cue usage and survey rating in the TS phase. Based on previous studies that demonstrated a reduction in the use of declarative task knowledge under stress (e.g., Schwabe & Wolf, 2012), we hypothesized that the participants with a greater increase in arousal (ΔSCL) would show a greater dissociation between declared and applied task knowledge in the TS relative to the NS phase. To investigate this relationship, we

performed a quadratic regression with ΔSCL as the independent variable and Δr_{sw} as the dependent variable:

$$\Delta r_{sw} = \beta_0 + \beta_1 \times \Delta SCL + \beta_2 \times \Delta SCL^2 + \varepsilon \quad \text{Eqn. 10}$$

where β_1 and β_2 represent regression coefficients, β_0 is a constant, and ε is a random error. In addition to reporting adjusted R^2 (R_{adj}^2), significance of the quadratic regression was tested via ANOVA.

4.2 Results

4.2.1 Physiological response to threat-of-shock

At the population level, analysis of SCLs revealed a main effect of phase, $F(2,112) = 29.43, \eta_p^2 = 0.35, p < 0.001$. Post-hoc pairwise comparisons using Bonferroni correction showed a significant increase of SCLs in the TS phase compared to the IL and NS phases ($ps < 0.001$), with no difference between the IL and NS phases ($p = 0.213$). To confirm that this selective increase in SCLs in the TS phase was not driven by any general trend for increased SCLs over time, we performed a 3 (phases) \times 2 (sets) repeated-measures ANOVA (Figure 9A). Each set was comprised of 120 unique stimulus pairs and therefore, this approach additionally controlled for any variance in SCL caused by task difficulty. This analysis revealed a significant effect of set, $F(1,56) = 5.29, \eta_p^2 = 0.09, p = 0.03$, which was due to an overall decrease rather than increase in amplitude with time. Taken together, these results confirm that the threat-of-shock manipulation successfully enhanced autonomic arousal.

To examine how autonomic arousal responses to the threat-of-shock manipulation varied across individuals, and how this responsivity may modulate decision strategies, we divided participants into low ($M = -0.13, SD = 0.08$), mid ($M = 0.10, SD = 0.07$), and high ($M = 0.39, SD = 0.09$) responders based on differences in SCLs between the TS and NS phases (ΔSCL) via the Ckmeans.1d.dp clustering algorithm (Wang & Song, 2011) (Figure 9B). As expected, one-way ANOVA revealed a significant difference in $\Delta SCLs$ between the three groups, $F(2,54) = 198.01, \eta_p^2 = 0.88, p < 0.001$, suggesting a reliable clustering of participants based on their level of autonomic arousal under threat-of-shock.

Lastly, we sought to examine the relationship between participants' STAI scores and changes in SCLs throughout the experiment. Neither state ($M = 34.28, SD = 8.37$) nor trait ($M = 38.35, SD = 10.04$) anxiety scores differed across the three groups (one-way ANOVA, $ps > 0.5$). Interestingly, trait anxiety scores were significantly correlated with SCLs in the IL ($r = 0.26, p = 0.049$) and TS ($r = 0.29, p = 0.027$) phases but not in the NS ($r = 0.12, p = 0.390$) phase, suggesting that participants with high trait anxiety displayed enhanced arousal at the beginning of the experiment as well as under the threat-of-shock condition. No significant correlations were found with state anxiety scores, nor between STAI scores and $\Delta SCLs$.

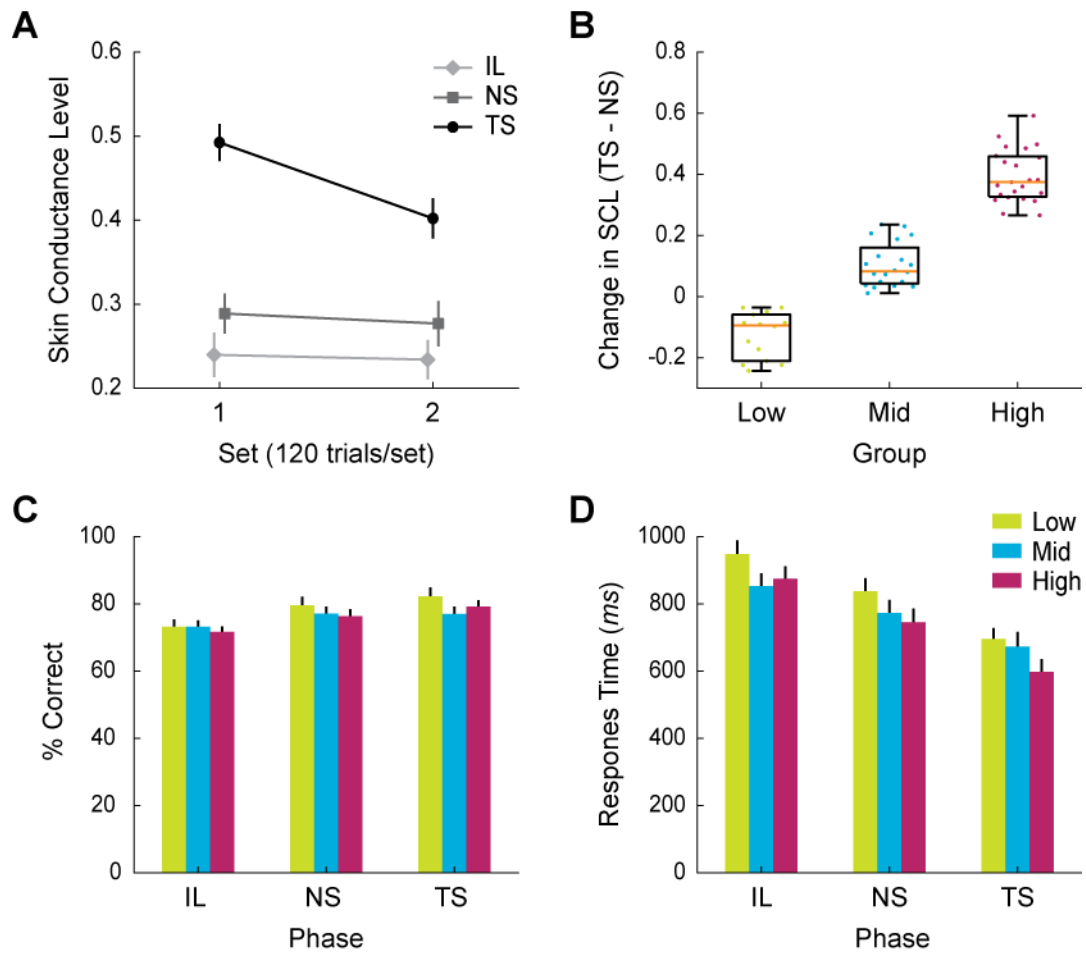


Figure 9: Participant Clustering and Decision Performance. (A) Average skin conductance level (range-corrected) per set in each experimental phase. (B) Participant clustering based on changes in SCL (Δ SCL). Dots represent individual participants. (C) Percent correct choices per group. (D) Average response time (ms) per group. Error bars indicate SEM.

4.2.2 Task performance

Task performance was assessed based on the correct decision defined by the sum of cue weights, independent of the feedback provided (see Section 2.2.1). All three SCL groups learned the task reliably, gradually improving and stabilizing performance over the course of the experiment, which was demonstrated by a significant main effect of

phase, $F(2,108) = 30.85, \eta_p^2 = 0.36, p < 0.001$ (Figure 9C). We observed no between-groups effect, $F(2,54) = 0.52, \eta_p^2 = 0.02, p = 0.597$, nor a phase \times group interaction, $F(4,108) = 0.18, \eta_p^2 = 0.06, p = 0.178$, indicating that all groups performed the task equally well. This observation was in line with the changes in RT throughout the experiment (Figure 9D). RT decreased as participants become more familiar with the task, which was shown by a significant main effect of phase, $F(2,108) = 128.00, \eta_p^2 = 0.70, p < 0.001$. Again, we observed no between-groups effect, $F(2,54) = 1.30, \eta_p^2 = 0.05, p = 0.282$, nor a phase \times group interaction, $F(4,108) = 2.25, \eta_p^2 = 0.08, p = 0.069$.

In sum, summary performance measures of mean RT and accuracy were not affected by responsiveness to threat-of-shock. However, these data cannot convey whether the underlying decision strategies producing these compound results may have been altered by the threat-of-shock manipulation, especially since the expected accuracy of some sub-optimal strategy models is near-optimal (Figure 2B; see Methods and Chapter 2). Hence, to achieve better sensitivity, we employed variational Bayesian inference to examine how threat-of-shock may affect decision strategy.

4.2.3 Decision strategy model comparison

We employed variational Bayesian inference to gauge participants' decision making strategies (Drugowitsch, 2013) and to quantify the most likely strategy model at the group level (Rigoux et al., 2014; Stephan et al., 2009). We explored 15 different plausible decision strategy models covering every possible case of cue usage, with

Model *Opt* representing the optimal cue-integration model (Figure 2A). The protected exceedance probabilities characterize how much more frequent a particular strategy model was employed than any other model, above and beyond chance, given the group model evidences, which are also reflected in BORs (Rigoux et al., 2014). In the NS phase, the optimal cue integration model, Model *Opt*, was the winning model for all three groups, with protected exceedance probabilities of 0.93 ($BOR = 0.013$; Low, Figure 10A), 0.96 ($BOR < 0.001$; Mid, Figure 10B), and 0.93 ($BOR < 0.001$; High, Figure 10C), respectively. The threat-of-shock manipulation did not have any effect on cue usage for the low and mid responder groups, for whom Model *Opt* remained the winning model with an exceedance probability of 0.96 ($BOR < 0.001$; Figure 10A) and 0.98 ($BOR < 0.001$; Figure 10B). The most likely decision model for the high responder group in the TS phase, however, switched to Model 11, with a protected exceedance probability of 0.84 ($BOR < 0.001$; Figure 10C). This result indicates that participants who experienced the greatest increase in arousal under the threat-of-shock manipulation stopped using the weakest cue in their decision making process, instead relying on only the three most informative cues.

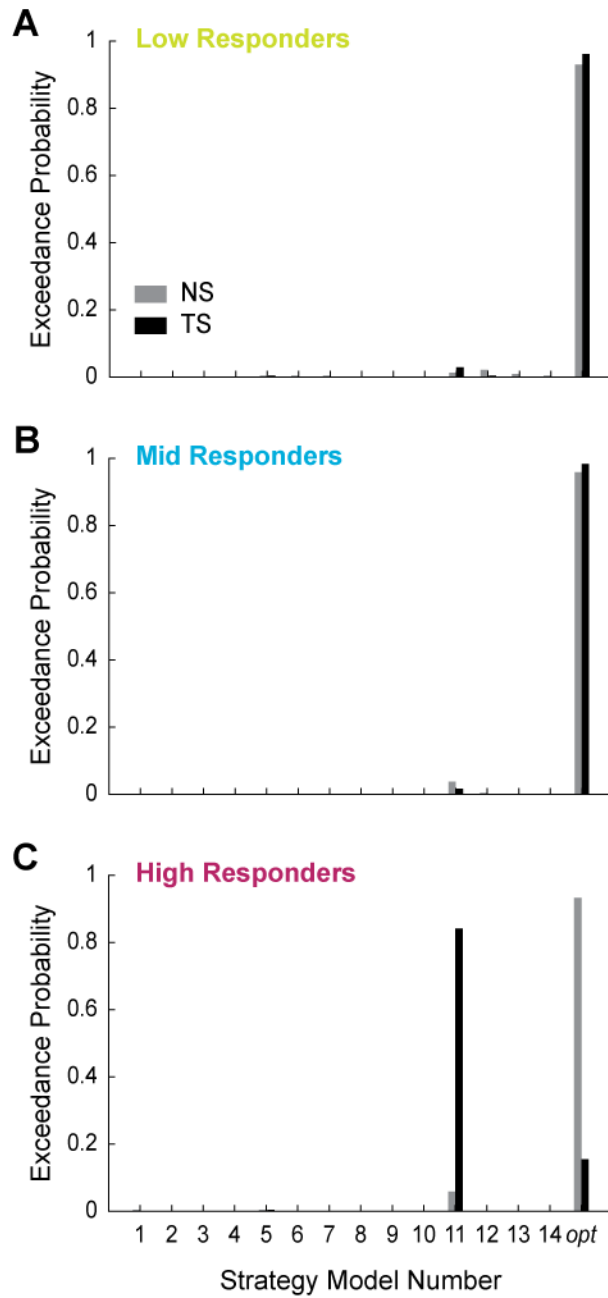


Figure 10: Strategy Model Selection per Responder Group. (A) Protected exceedance probabilities for the low responder group. (B) Protected exceedance probabilities for the mid responder group. (C) Protected exceedance probabilities for the high responder group.

4.2.4 Post-experimental survey

In the post-experimental survey, participants were first asked to rank cue dimensions from the most to the least informative. Participants' ratings of cue dimensions showed a significant positive correlation with the assigned weights, $w_{net.i}$ (mean $r_w = 0.65$, $SEM = 0.06$; one-sample $t(56) = 8.28$, $p < 0.001$), without any between-group difference, $F(2,54) = 0.53$, $p = 0.592$, suggesting that all three groups of participants were able to explicitly describe the relative importance of the true underlying cue weights correctly to some extent. Similarly, their survey cue ratings and the subjective weights, $w_{opt'}$, showed significant positive correlations in both the NS (mean $r_{sw,NS} = 0.57$, $SEM = 0.06$; $t(56) = 7.35$, $p < 0.001$) and the TS (mean $r_{sw,TS} = 0.63$, $SEM = 0.06$; $t(56) = 8.81$, $p < 0.001$) phases. No difference was observed across the responder groups, $F(2,54) = 1.60$, $\eta_p^2 = 0.06$, $p = 0.211$, along with no main effect of phase, $F(1,54) = 0.58$, $\eta_p^2 = 0.01$, $p = 0.451$, nor phase \times group interaction, $F(2,54) = 0.53$, $\eta_p^2 = 0.02$, $p = 0.591$.

To further investigate the relationship between participants' explicit knowledge and cue usage across phases, we assessed strategy concordance with declarative knowledge (Δr_{sw} ; see Section 4.1.8 above), which reflects changes in subjective weights from the NS to TS phase relative to survey cue ratings. If participants formed more accurate insights to their cue usage or refined strategies according to their declarative knowledge of the cue structure as the experiment progressed, Δr_{sw} would be positive.

Conversely, if there was an increased discordance between participants' explicit knowledge and cue usage in the TS phase, Δr_{sw} would be negative, indicating that subjects failed to employ their declarative knowledge of cue weights in executing decisions under the threat-of-shock condition. To examine the effect of arousal (ΔSCL) on strategy concordance (Δr_{sw}), we fit a quadratic model (Eqn. 10), which revealed a significant inverted-U-shaped relationship, $R_{adj}^2 = 0.12, F(2,54) = 4.92, p = 0.011$ (Figure 11). A simple linear regression did not show any meaningful relationship, $R_{adj}^2 = 0.04, F(1,55) = 1.97, p = 0.166$. These results suggest that participants who experienced moderate levels of arousal during the TS phase displayed enhanced concordance between their declarative knowledge and applied strategy, whereas those who experienced low or high levels of arousal showed increased discordance between their subjective cue weighting and explicit knowledge under the threat-of-shock. Additionally, the high responder group demonstrated a significantly negative correlation between ΔSCL and Δr_{sw} (Pearson's $r = -0.60, p = 0.003$), which is in line with the possibility that this group deviated from using their internal model consistently when moving from the NS to the TS phase. In other words, even though high SCL responders were able to articulate the relative utility of all four cues to a comparable extent as other participants after the experiment, during the TS phase they did not employ this knowledge consistently to guide their decisions.




Figure 11: The effect of anticipatory anxiety on task knowledge. Solid line represents a significant quadratic fit ($p = 0.011$) between changes in SCL (ΔSCL) and the implicit-explicit strategy concordance (Δr_{sw}). Colored dots indicate individual participant per group. Dashed lines represent the 95% CI for the regression line.

Next, participants were asked the number of cue dimensions used in each phase.

On average, subjects reported that they used 2.09 cues ($SEM = 0.09$) in the NS phase and 2.32 cues ($SEM = 0.09$) in the TS phase, which demonstrated a significant main effect of phase, $F(1,54) = 4.86, \eta_p^2 = 0.08, p = 0.032$, but without a phase \times group interaction, $F(2,54) = 0.10, \eta_p^2 = 0.004, p = 0.906$, nor a between-groups effect, $F(2,54) = 1.56, \eta_p^2 = 0.06, p = 0.220$. The discrepancy between results from the Bayesian strategy model selection and survey responses seems to suggest that a majority of participants did not have precise insight into the strategies (i.e., cue utilization) they employed to solve the task, which replicated previous experiments using this protocol (Oh et al., 2016). Lastly, subjects were asked to identify the cue state with the higher weight within each cue dimension (e.g., choice of blue vs. red). The mean number of correct responses

was 3.63 ($SEM = 0.08$) out of 4 and there was no difference across responder groups, $F(2,54) = 0.77, p = 0.466$, indicating that all responder groups were able to identify the better cue state to a similar extent.

4.3 Discussion

To characterize the impact of anticipatory anxiety on complex decision-making, we tested healthy participants on a multi-cue probabilistic classification task under neutral and threat-of-shock conditions. Specifically, we focused on examining changes in post-learning decision performance, once participants had formed sufficient understanding of the cue structure of the task. To account for individual differences in susceptibility to the threat-of-shock manipulation, participants were clustered into three responder groups based on changes in SCLs from the NS to TS phase. In the neutral phase, all three groups of participants utilized information near-optimally, appropriately weighting and integrating all available cues to make choices. Under threat of shock, however, the high responder group adopted the drop-the-worst satisficing strategy by considering only the three most informative cues and ignoring the least important cue, whereas the low and mid responder groups did not show any change in strategy. Although suboptimal, this shift to the drop-the-worst heuristic was adaptive, since it yielded a comparable performance outcome while using less information. In addition, strategy concordance with participants' declarative knowledge exhibited an inverted-U relationship with ΔSCL , suggesting that a moderate level of arousal can foster

the formation and use of accurate, explicit strategy knowledge. This concordance, however, decreased with increasing arousal, indicating that participants who experienced greater anxiety demonstrated higher divergence of implicit strategy from their declarative cue knowledge. These findings suggest that high levels of anticipatory anxiety may trigger a shift from near-optimal to satisficing decision making, which is also associated with decreased implicit-explicit strategy concordance.

Faced with real-world problems, where obtaining and combining all decision-relevant information for normatively optimal decision-making is often not feasible, people exhibit *bounded rationality* by resorting to good-enough solutions (Gigerenzer & Goldstein, 1996; Simon, 1955). Such adaptive satisficing heuristics depend heavily on exploiting the structure of the environment, thereby simplifying the decision problem without any significant impairment in performance (Gigerenzer & Gaissmaier, 2011; Gigerenzer & Goldstein, 1996; Simon, 1990). Although affect and emotion are thought to influence this type of satisficing decision behavior (Simon, 1990), only a handful of studies have systematically examined this relationship based on risky decision scenarios (e.g., Pachur, Hertwig, & Wolkewitz, 2014), where the probabilities of possible outcomes were explicitly available to participants. While prominent theories have conjectured that anxiety reduces available cognitive resources, and therefore, restricts utilization of threat-irrelevant information (Easterbrook, 1959; Eysenck et al., 2007), the exact nature of this change during decision-making, especially under uncertainty, has remained

unclear. Here, in accordance with this broad prediction, we demonstrated a reduction in low-information cue utilization under heightened anxiety. Importantly, this change in choice behavior was adaptive: to cope with anxiety, participants experiencing high arousal during the TS phase employed a satisficing heuristic that used less information without sacrificing overall performance. This observation is also consistent with our previous findings on satisficing under time pressure, where participants demonstrated adaptive cue discounting by systematically ignoring the least informative cue(s) with increasing time pressure (Oh et al., 2016). We here expanded these results by demonstrating that internally-induced pressure (i.e., anxiety) similarly fosters the switch to the drop-the-worst satisficing decision-making strategy.

Much work in probabilistic decision-making suggests that acute stress fosters habitual learning and attenuates contributions of the goal-directed system (Otto et al., 2013; Schwabe & Wolf, 2012; for reviews, see Phelps, Lempert, & Sokol-Hessner, 2014; Schwabe & Wolf, 2013). Although adopting either strategy often does not lead to a significant difference in probabilistic classification accuracy, the use of the habitual system is often linked to diminished explicit task knowledge (Foerde, Knowlton, & Poldrack, 2006; Lars Schwabe & Wolf, 2012). A similar trend was observed for participants under emotional arousal (Thomas & LaBar, 2008). That is, compared to a control group, participants in the emotional condition, where outcomes were paired with phobic stimuli (e.g., snake/spider), were more prone to using a simple, suboptimal

strategy, with less declarative insights about probabilistic cue-outcome associations. Interestingly, a recent functional neuroimaging (fMRI) study reported an inverted-U-shaped relationship between arousal and saliency-executive network cohesion (Young et al., 2017), which is indicative of a suboptimal engagement of executive control network under high arousal. Congruent with this finding, we here report that participants who experienced high arousal under threat-of-shock demonstrated significantly increased discordance between implicit-explicit cue knowledge, indicating a possible shift to habitual decision-making during the TS phase. Hence, the inverted-U relationship between arousal and concordance measure, may suggest that moderate level of arousal or anxiety can aid participants to form more concrete, declarative insights into their implicit strategy over time.

Shifting from goal-directed to habitual decision-making can serve as an adaptive coping mechanism in response to the sustained stressor. Compared to goal-directed behavior, habitual processes require less cognitive resources and thus, lead to faster and more efficient decisions, all of which could help maintain performance (Lars Schwabe & Wolf, 2013). Additionally, in some cases, anxiety can facilitate information processing by reducing interference from task-irrelevant cues (e.g., Easterbrook, 1959; Hu, Bauer, Padmala, & Pessoa, 2012). In line with this notion, accuracies in a dual-target visual search task, performed under neutral and threat-of-shock conditions, were comparable for high-salience targets, whereas detection of a second, low-saliency target was

impaired under threat-of-shock (Cain et al., 2011). Hence, to cope with anxiety, healthy participants might employ heuristics by selectively focusing on high-value information while ignoring low-valued cues. Many studies have demonstrated that, if employed appropriately, such heuristic models could perform equivalent or sometimes even better than more complex statistical models (e.g., Czerlinski, Gigerenzer, & Goldstein, 1999; Martignon, Katsikopoulos, & Woike, 2008). That is, successful use of satisficing heuristics is heavily dependent on adaptive exploitation of task structure by accurately identifying and utilizing the most important set of information. In our task, we examined post-learning decision behavior once participants had formed a good understanding of the task structure, which might have contributed to the “adaptive” reduction of cue space under threat-of-shock.

Although the threat-of-shock anxiety-induction technique provides an effective way of directly manipulating state anxiety in healthy participants, clinical studies have reported somewhat diverging behavioral effects of pathological anxiety from induced-anxiety (Robinson, Vytal, et al., 2013). That is, patients with clinical anxiety disorders, compared to healthy controls, typically demonstrate elevated baseline arousal (Grillon, Morgan, Davis, & Southwick, 1998) and exhibit heightened sensitivity to threat-related information (Bar-Haim, Lamy, Pergamin, Bakermans-Kranenburg, & van IJzendoorn, 2007), which in turn can impair various executive functions such as attentional control (e.g., Martin, Williams, & Clark, 1991), working memory (e.g., Kizilbash, Vanderploeg, &

Curtiss, 2002), and decision making (e.g., Sailer et al., 2008). Therefore, it seems likely that pathological anxiety may interfere with feedback-based learning (Petzold, Plessow, Goschke, & Kirschbaum, 2010), which could result in inaccurate information usage and inappropriate application of heuristics. Applying our task and methods to clinical populations in future studies may help better understand underlying causes of deficits in learning and decision-making observed in pathological anxiety disorders.

In conclusion, the present study characterized how anticipatory anxiety induced by threat-of-shock leads to satisficing decision behavior. Throughout the learning and neutral phases, all three groups of participants performed near-optimally, appropriately utilizing all available information to arrive at choices. Under threat-of-shock, however, the high SCL responder group shifted to using the drop-the-worst satisficing strategy by focusing only on a subset of three most informative cues. Additionally, an inverted-U relationship was observed between arousal and implicit-explicit strategy concordance, suggesting that high levels of anxiety increased divergence of participants' actual strategy use from their explicit knowledge.

5. Neural Mechanisms of Satisficing Decision Making Under Time Pressure

In Chapters 3 and 4, we have established that people adaptively employ the Drop-the-Worst heuristic to cope with externally (e.g., time) and internally (e.g., anticipatory anxiety) induced satisficing pressures. Specifically, in Chapter 3, we demonstrated that, under low time pressure, information was integrated near-optimally across all available cues. By contrast, under high pressure, participants dropped the weaker, less predictive cues from the decision-making process. What remains to be addressed are computational principles and neural mechanisms underlying such satisficing decision making. Hence, to elucidate the neural dynamics underlying this shift in decision modes from optimal to satisficing, in the present chapter, we combined the multi-cue probabilistic classification task, performed under low (1500 ms) and high (500 ms) time pressure, with functional magnetic resonance imaging (fMRI). Using variational Bayesian inference, we quantified participants' cue usage and related it to changes in regional blood-oxygen-level dependent (BOLD) signals.

While we are not aware of any previous study assessing the neural mediators of probabilistic inference under time pressure, prior reports on statistical learning under stress, and studies of the speed-accuracy tradeoff in perceptual decision-making, offer grounds for tentative hypotheses. Probabilistic inference has been studied extensively through variants of the weather prediction task (Gluck & Bower, 1988; Knowlton et al.,

1994), where acquiring probabilistic cue-outcome relationship through feedback has been shown to be associated with activity in the striatum, hippocampus (Knowlton et al., 1996; Poldrack et al., 2001; Shohamy, Myers, Grossman, et al., 2004) and parietal cortex (Yang & Shadlen, 2007). In addition, other recent studies of probabilistic decision-making suggest an important role for the frontoparietal attentional control network in mediating learning in a multidimensional decision environment (Niv et al., 2015), and the orbital/ventromedial prefrontal cortex (vmPFC) in encoding expected reward, subjective value, outcome predictions, and credit assignment (Akaishi et al., 2016; Daw, O'Doherty, Dayan, Seymour, & Dolan, 2006; D. J. Levy & Glimcher, 2012; John P. O'Doherty, Kringelbach, Rolls, Hornak, & Andrews, 2001).

Stress has shown to bias decision-making strategies by reducing contributions of the prefrontal cortex (PFC) and encouraging habitual stimulus-response processes (Dias-Ferreira et al., 2009; L. Schwabe & Wolf, 2009). Specifically, learning the weather prediction task under stress induced by the cold pressor test has been associated with increased use of implicit, striatum-mediated strategies (Lars Schwabe & Wolf, 2012). Similarly, time pressure on perceptual decision making has been associated with a deterioration in information processing in early sensory areas (Ho et al., 2012) and increased activity in the striatum (Bogacz, Wagenmakers, Forstmann, & Nieuwenhuis, 2010; Forstmann et al., 2008), indicating that the striatum may promote faster but possibly premature or sub-optimal decisions.

Here, we characterized how the brain encodes probabilistic cue information as participants shift from employing optimal to satisficing decision strategies with increasing time pressure. Based on the above studies, we predicted that probabilistic decisions will be mediated by both subcortical (striatum, hippocampus) as well as prefrontal (lateral and medial PFC) and parietal regions under low time pressure, with a preferential involvement of the striatum under high time pressure. The data supported this hypothesis and revealed details of the networks involved in this cortical-to-subcortical shift of activity.

5.1 Methods

5.1.1 Participants

Thirty-two healthy volunteers participated in this experiment. Seven participants were excluded from further analysis: Five participants due to chance-level performance and two participants due to excessive head movement ($> 20\text{ mm}$). The final sample consisted of twenty-five subjects (13 females, mean age = 27 years, range = 18 – 40 years). All participants provided informed consent in line with Duke Medical Center institutional guidelines and were compensated with \$40 for their time (2 h).

5.1.2 Procedure

Participants performed the multi-cue probabilistic classification task based on 16 unique compound stimuli (Figure 1B; see Chapter 2 for details). Importantly, the weights outlined in Table 2 were randomly assigned to the different cue dimensions for

each participant at the beginning of the experiment, and every possible weight permutation (24 total) was used at least once. Thus, our neuroimaging analyses focused on encoding of “cue feature-invariant” informational value, by dissociating cue weights from specific cue dimensions and visual features across the subject population.

Similar to previous studies, on each trial, participants were presented with two compound stimuli and asked to indicate their choice via keypress within a specified choice window (Figure 1C). Upon the choice deadline, stimuli disappeared from the screen, and the outcome of the choice (“win”, “lose”) or a no-response warning (“miss”) was displayed for 500 ms. After a variable intertrial interval (ITI) of 3 - 5 s drawn from a pseudo-exponential distribution (mean ITI = 3.5 s), the next trial began with a new pair of stimuli. Prior to the scan, all participants completed a 240 trial learning phase, comprised of two successive sets of all unique trials presented in random order. During the learning phase, participants were given a 1.5 s choice window to register their responses. The goal of this phase was to allow participants to explore and learn the cue weights by trial and error and familiarize themselves with the probabilistic classification task. Once the initial learning phase was complete, participants performed two more task phases (240 trials/phase) inside the scanner: (1) a low time pressure (LowP) phase with 1.5 s choice window (identical to the practice phase), and (2) a high time pressure (HighP) phase with a 0.5 s choice window. Each phase was separated into four runs (60 trials/run).

5.1.3 Behavioral performance analysis

See Chapter 2 for the complete procedure. Additional data analyses performed in this chapter are detailed below.

5.1.3.1 Sum of Evidence (SoE)

In addition to the objectively defined amount of evidence (see Section 2.2.2), we quantified the *subjective* SoE on each trial using the sum of inferred, subjective cue weights, w^* (Eqn. 8), under low and high time pressure conditions. Similar to the objective SoE, subjective SoE is closely related to perceived decision difficulty of a given trial. For the behavioral data analyses, we employed SoE as a basic manipulation check, by testing whether % correct choices and response times scale with SoE. In the neuroimaging analyses, we employed both subjective and objective SoE to probe which brain regions tracked the amount of evidence or outcome uncertainty, tailored to each participant's decision strategy under low and high time pressure (detailed below).

5.1.3.2 Subjective cue weights

Subjective cue weights, w^* , were estimated based on the optimal strategy model, $Model_{opt}$, using hyper priors, $a_0 = 0.345$ and $b_0 = 0.584$ (Eqn. 3 – 8; see Section 2.2.4).

5.1.3.3 Comparison between decision strategies adopted under low and high time pressure

To identify the cue dimensions that were effectively used during each time pressure phase, we explored 15 different decision strategy models, accounting for every possible combinations of cue usage (Figure 2). Log model evidences were estimated via

the variational Bayesian logistic regression method (Eqn. 3 – 6) and were submitted to the group-level Bayesian model selection procedure. Note that using relatively uninformative hyper priors ($a_0 = 0.01$ and $b_0 = 0.0001$) in Equation 6 does not change the overall results of model comparison.

5.1.4 fMRI data acquisition

Images were acquired on a 3 T GE MR750 scanner. A T1-weighted structural images were scanned parallel to the AC-PC plane (146 slices, slice thickness = 1 mm, TR = 8.124 ms, FoV = 256 mm × 256 mm, in-plane resolution = 1 mm × 1 mm). Functional images were scanned using a T2*-weighted single-shot gradient EPI sequence (42 slices, slice thickness = 3 mm, TR = 2 s, TE = 28 ms, flip angle = 90°, FoV = 192 mm × 192 mm, in-plane resolution = 3 mm × 3 mm). 169 functional images per run were acquired for the first four runs of the LowP phase and 139 images per run were acquired for the last four runs of the HighP phase.

5.1.5 Image preprocessing

Preprocessing and univariate statistical analyses were performed using SPM12 (<http://www.fil.ion.ucl.ac.uk/spm/software/spm12/>). After discarding the first four scans of each run, functional images were realigned to their mean image and corrected for slice timing. Each participant's structural image was co-registered to the mean functional image and segmented into gray matter, white matter, cerebro-spinal fluid, bone, soft tissue, and air/background. The deformation field map obtained through

segmentation was applied to normalize the subject's functional images to the Montreal Neurological Institute (MNI) template space. Functional images were resampled into $3 \times 3 \times 3$ mm voxel size and spatially smoothed with a Gaussian kernel of 5 mm full-width at half maximum (FWHM).

5.1.6 fMRI data analyses

Assuming that participants learned the cue structure and associated weights, the decision task can be approached in two distinct but complementary (and not mutually exclusive) ways: (1) analyzing each cue dimension difference between the two compound stimuli present on a given trial based on *subjective cue weights*, and adding them up, and/or (2) summing the cue states comprising each compound stimulus and estimating the difference between the sums of the two stimuli, which is also equivalent to the *subjective SoE*. Both approaches would lead to the same solution, with the first approach having more emphasis on parsing the compound stimuli out into individual cue dimension differences prior to integrating the cues to evaluate the total amount of evidence. The preferential use of either of these two approaches might depend on a participant's subjective cue weight distribution (e.g., weights are evenly distributed or one cue weight dominates all others) as well as on trial type (i.e., which cue dimensions differed between stimuli on a given trial). We assumed that participants likely used a mixture of these approaches, and we therefore focused our fMRI analyses on characterizing neural substrates of both subjective cue weight and SoE representation as

a function of time pressure. For the following fMRI analyses, we used objectively correct trials (as favored by the probability defined in Eqn. 1 – 2) as the events of interest.

5.1.6.1 Neural representation of subjective cue weights

To identify brain regions that represent information about subjective cue weights of individual cue dimensions, we employed a multivariate decoding technique based on support vector machine regression (SVR). Recall that these analyses solely concern the informational value of the cues, rather than specific visual cue features, as our task design dissociated cue weights from specific cue dimensions/features across the subject population. To achieve maximal sensitivity, we used realigned and slice time corrected functional images in each participant's native space without spatial normalization and smoothing. Functional data were first analyzed using the standard general linear model (GLM) approach (Friston et al., 1994) to regress the BOLD signal against task models and to estimate parameter β s for conditions of interest. Four regressors, each representing a unique cue dimension, tracked the presence of the corresponding cue dimension difference between a compound stimulus pair on each trial. For example, if two stimuli were different in c_1 and c_2 dimensions but not in c_3 and c_4 on trial t , boxcar functions were created at trial t only in the regressors for c_1 and c_2 . Following this rule, all correct trials were modeled as boxcar functions of durations 2 s (LowP) or 1 s (HighP) aligned to trial onsets, capturing both the stimulus presentation and the subsequent feedback in each phase. As regressors of no interest, we included a categorical regressor

of win (1) and lose (-1) probabilistic feedback. To account for differences in response time (RT) across trials and experimental phases, a parametric regressor of RT aligned to the trial onsets was also included. In addition, incorrect trials and trials with no response were modeled separately, along with six head-motion parameters and grand means of each run. All regressors were convolved with the canonical hemodynamic response function. The resulting parameter estimates of the four cue dimension regressors, β_{C_i} ($i = 1, 2, 3, 4$), were used to search for brain areas that contain information about the subjective cue weights under low and high time pressure conditions.

Specifically, we conducted multivoxel pattern analysis (MVPA) with a whole-brain searchlight approach (Haynes et al., 2007; Kriegeskorte, Goebel, & Bandettini, 2006) that scanned through spheres of gray matter voxels (searchlight radius = 4 voxels) identified using each participant's gray matter mask produced from T1 segmentation. The MVPA was performed by using a linear SVR with a constant regularization parameter of $C = 1$ (Kahnt, Heinzle, Park, & Haynes, 2011), implemented in MATLAB. We iteratively used 3 of the 4 runs to train an SVR model and then used the remaining run as the test data set to predict the value of the subjective cue weights (i.e., a leave-one-run-out procedure). Prediction accuracy was determined by the Fisher's Z-transformed correlation coefficients (Kahnt et al., 2011) between the predicted values and the subjective cue weights, w^* , of the test data set. An accuracy map was constructed by averaging the prediction accuracy from 4-fold cross-validation for each

center voxel of the searchlight spheres. We conducted the decoding analyses separately for the LowP and HighP phases, using the corresponding subjective cue weights.

The accuracy maps of each participant were normalized into the template MNI space and smoothed with a Gaussian kernel (FWHM = 5 mm) to account for differences in activation localization across subjects (Kahnt et al., 2011). The group analysis was performed by entering the accuracy maps into one-sample *t*-tests, separately for LowP and HighP phases. The *p*-value maps were corrected for multiple comparison using the CorrClusTh.m function developed for use with SPM (<http://www2.warwick.ac.uk/fac/sci/statistics/staff/academic-research/nichols/scripts/spm/spm8/corrclusth.m>), which determined that an uncorrected voxelwise threshold of $p < 0.005$ combined with a cluster size of 46 voxels corresponded to cluster level corrected threshold of $p < 0.05$.

5.1.6.2 Modulation of neural activity by SoE

Subjective SoE reflects the integrated sum of subjective cue weights of a given trial, which should be highly correlated with the objective amount of evidence and decision difficulty. Therefore, brain regions modulated by subjective SoE likely encode not only the integrated sum of cue weights but also decision processes involved in producing a final choice output. Since subjective SoE can vary over up to 40 distinct levels, it would not yield a sufficiently reliable estimation of parameter β s (~5 trials/level) to utilize the multivariate analysis approach carried out on the subjective cue

weights above (see *Neural representation of subjective cue weights*). Therefore, we instead employed a parametric modulation analysis within the standard mass-univariate GLM approach to detect areas whose activation was modulated by SoE. Correct trials were modeled as boxcar functions of 2 s (LowP) or 1 s (HighP) duration, aligned to trial onsets. A parametric regressor of subjective SoE, estimated separately for the LowP and HighP phases using the associated subjective weights, was attached to trial onsets. As the regressors of no interest, a categorical feedback regressor and a parametric regressor of RT were aligned to trial onsets. In addition, incorrect trials, trials with no response, six head-motion parameters, and grand means of each run were also included as regressors of no interest. All regressors were convolved with the canonical hemodynamic response function and regressed against the BOLD signal in each voxel. Within-subject effects of modulation of neural activity by subjective SoE were assessed for the LowP and HighP phases separately. The resulting single-subject contrast maps were entered into group-level analyses, which treated subjects as random effects.

To probe for differences between experimental conditions, we further contrasted effects of SoE modulation between the LowP and HighP phases. For all analyses, the p -value maps were corrected for multiple comparison using the function `CorrClusTh.m`, which determined that an uncorrected voxelwise threshold of $p < 0.005$ combined with a cluster size 96 to 105 voxels ensured a false discovery rate < 0.05 . In addition, although SPM orthogonalizes parametric modulators to compute the GLM, results could

potentially be influenced by the order of the modulators. To ensure that the observed effects are independent from this influence, we additionally ran GLMs by varying the order of parametric regressors (subjective SoE, feedback, and RT), which did not yield any significant difference in overall neuroimaging results.

5.1.6.3 Region-of-Interest analysis of objective SoE

To further investigate the relationship between neural encoding of objective and subjective SoE, we performed a region-of-interest (ROI) analysis. First, we identified brain regions modulated by subjective SoE in both low *and* high time pressure phases by applying a “logical AND” conjunction analysis (Nichols, Brett, Andersson, Wager, & Poline, 2005). Voxels that passed the multiple comparison correction in both the LowP and HighP phases were included in the conjunction map. For clusters that span multiple regions, we applied anatomical masks to include only the voxels within a specified area. Anatomical ROIs were defined using the WFU PickAtlas toolbox (Maldjian, Laurienti, & Burdette, 2004; Maldjian, Laurienti, Kraft, & Burdette, 2003) and clusters with less than 10 voxels were excluded from the analysis. Then, we employed the same GLM approach as above but treated each objective SoE level as a separate condition. To roughly match the number of trials included in each regressor, we merged the three highest SoE levels (1.6, 1.8, 2), which resulted in nine parameter estimates β_{SoE_i} with i ranging from 0 (no evidence) to 1.6* (high evidence). Similar to the previous GLMs, correct trials corresponding to each SoE level were modelled with boxcar functions along with

parametric modulators of feedback and RT. Additionally, incorrect trials, trials with no response, six head-motion parameters, and grand means of each run were included. To investigate the effect of objective SoE on modulating activities in regions sensitive to subjective SoE, the parameter estimates obtained from the GLM were extracted and averaged within each ROI, and tested for linear trend using a repeated-measure ANOVA.

5.1.6.4 Practice effect analysis

Since participants completed the low and high time pressure phases in a block-wise, sequential manner, it is important to rule out that changes in BOLD signal across experimental phases are not merely driven by task practice. Therefore, we conducted an additional GLM analysis to examine the presence of the neural practice effects. Specifically, we hypothesized that we would observe gradual BOLD signal changes over time if practice effects were present. To characterize overall changes in trial-induced BOLD activation over the course of the experiment, we divided our task into four different sets (120 trials/set, 2 sets/phase), each consisting of a complete set of unique stimulus combinations. Then, correct trials were modeled as boxcar functions of 2 s (LowP) or 1 s (HighP) duration, aligned to trial onsets, with parametric regressors of feedback and RT. Additionally, incorrect trials, trials with no response, six head-motion parameters, and grand means of each run were included as regressors of no interest. All regressors were convolved with the canonical hemodynamic response function. Then,

the percent signal changes (MarsBaR; <http://marsbar.sourceforge.net/>) and the parameter estimates, β_{Trial} , were extracted and averaged within the predefined ROIs and entered into 2 (sets) \times 2 (LowP/HighP phases) repeated-measures ANOVAs. To further investigate changes in subjective SoE modulation per set, we repeated the same GLM analysis using subjective SoE (see above), but splitting the data into four sets instead of two phases.

5.2 Results

5.2.1 Behavioral data

5.2.1.1 Task performance

Participants performed a multi-cue probabilistic classification task under low (1500 ms) and high (500 ms) time pressure inside the fMRI scanner. Their task was to compare two compound stimuli comprised of four different visual features (color, shape, contour, and line orientation) and make a prediction on a stimulus that is more likely to win (Figure 1). Upon each choice deadline, the outcome (“win” or “lose”), determined probabilistically based on the cue weights, was displayed (Table 2). Performance was evaluated based on the number of correct choices favored by the cue weights, independent of the probabilistic outcome feedback participants experienced. Throughout the learning phase prior to the scan, participants were able to gradually improve their decision performance as characterized by a significant main effect of block (accuracy: $F_{(3,72)} = 8.36, p < 0.001$; RT: $F_{(3,72)} = 7.54, p = 0.001$) and a linear trend

(accuracy: $F_{(1,24)} = 14.52, p = 0.001$; RT: $F_{(1,24)} = 9.80, p = 0.005$) (Figure 12A left). In the LowP phase (1500 ms response window), participants achieved better mean decision accuracy, $t_{24} = 5.62, p < 0.001$, and expedited decision speed, $t_{24} = 4.12, p < 0.001$, compared to the learning phase as they became yet more accustomed to the task, but performance remained stable throughout the LowP phase (main effect of block, accuracy: $F_{(3,72)} = 0.70, p = 0.56$; RT: $F_{(3,72)} = 2.66, p = 0.06$) (Figure 12A middle). In addition, performance scaled with objective SoE, defined by the sum of cue weight differences between the stimulus pair. As SoE increased, percent correct choices increased (linear trend, $F_{(1,24)} = 153.03, p < 0.001$) and decision time decreased (linear trend, $F_{(1,24)} = 30.54, p < 0.001$), indicating that participants learned to integrate the predictive value of the four cue dimensions to base their decisions on (Figure 12B-C).

As anticipated, increased time pressure in the HighP phase (500 ms response window) speeded up decision time, $t_{24} = 20.86, p < 0.001$, compared to the LowP phase, but had a detrimental effect on decision making, revealed by a large decrease in accuracy, $t_{24} = 8.08, p < 0.001$, which remained constant throughout the HighP phase (main effect of block, accuracy: $F_{(3,72)} = 1.53, p = 0.21$; RT: $F_{(3,72)} = 0.62, p = 0.56$) (Figure 12A right). Despite this significant decrement in overall performance, decision accuracy nevertheless scaled with objective SoE in the HighP phase (linear trend, $F_{(1,24)} = 30.54, p < 0.001$), confirming that participants were able to use available cue information to guide their choices (Figure 12B). By contrast, RT was no longer related to

SoE (linear trend, $F_{(1,24)} = 0.96, p = 0.34$), presumably due to the severe time pressure enforced in this condition (Figure 12C).

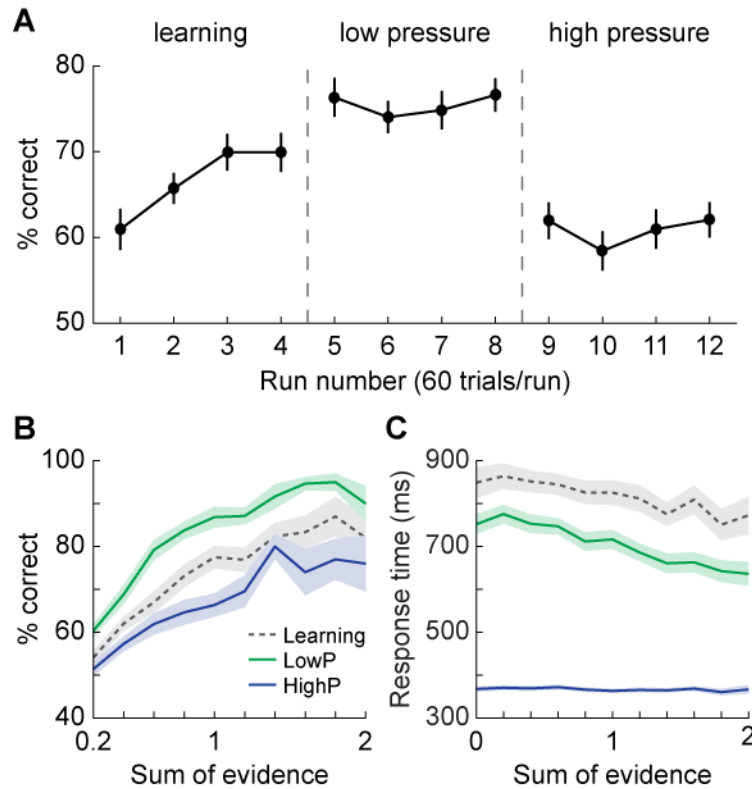


Figure 12: Behavioral Results. (A) Performance throughout the task runs. Plotted is the percentage of correct choices favored by the sum of cue weights, regardless of outcome feedback. (B) Percentage of correct choices as a function of objective SoE. (C) Response time (ms) as a function of SoE. Error bars and shaded area represent SEM.

5.2.1.2 Subjective cue weights

To examine the relative importance of each cue dimension in guiding participants' choices, separate sets of subjective cue weights for the LowP and HighP

phases were obtained using logistic regression (Figure 13A, Eqns. 3 – 8). A repeated-measures ANOVA revealed a main effect of cue weights, $F_{(3,72)} = 21.29, p < 0.001$, which was characterized by a significant linear trend, $F_{(1,24)} = 39.02, p < 0.001$, indicating that participants were able to correctly rank the cues according to their objective order of importance, which was true for both the LowP, $F_{(3,72)} = 23.55, p < 0.001$, and the HighP, $F_{(3,72)} = 8.86, p < 0.001$, phases. Nevertheless, the main effect of phase was significant, $F_{(1,24)} = 38.71, p < 0.001$, as subjective cue weights were overall smaller in the HighP than in the LowP phase, reflecting a general down-weighting of cue weights under time pressure. Since the magnitude of fitted weights also corresponds to the decision noise, this overall decrease of cue weights reflects the deterioration of decision making performance in the HighP phase. Finally, the phase \times cue weight interaction, $F_{(3,72)} = 11.23, p < 0.001$, was also significant, as the relative difference between cue weights between the two phases increased as a function of the assigned cue weights.

5.2.1.3 Decision strategy model selection

To characterize the difference in decision strategy under the conditions of low and high time pressure, we explored 15 different plausible strategy models covering every possible combinations of cue usage (Figure 13D). We estimated marginalized log-likelihood by fitting a logistic function using variational Bayesian inference (Eqns. 3 – 6; Drugowitsch, 2013) for each model per participant. These log model evidences were

then used to fit the hierarchical model (Rigoux et al., 2014; Stephan et al., 2009) to estimate the most likely strategy model employed in each experimental phase at the group level (Figure 13C). In the LowP phase, the optimal cue integration model was the most likely model with an exceedance probability of 0.72. In the HighP phase, however, Model 11, which utilizes only the three most informative cues, was the winning model, with an exceedance probability of 0.79. This shift in strategy from the optimal model to Model 11 suggests that participants dropped the worst cue in their decision-making process under increasing time pressure, replicating our previous results (Oh et al., 2016). Additionally, to address the possibility that participants may have engaged in a stimulus-based decision process by learning and memorizing weights of each of the 16 compound stimuli rather than the four cue dimensions, we also included a strategy model based on summed cue weights in the model comparison with the other 15 aforementioned cue-based models per phase. The group-level model comparison results yielded very weak support for the stimulus-based model in all three phases (exceedance probabilities < 0.005), suggesting that participants engaged in a cue dimension-based decision process, regardless of time pressure.

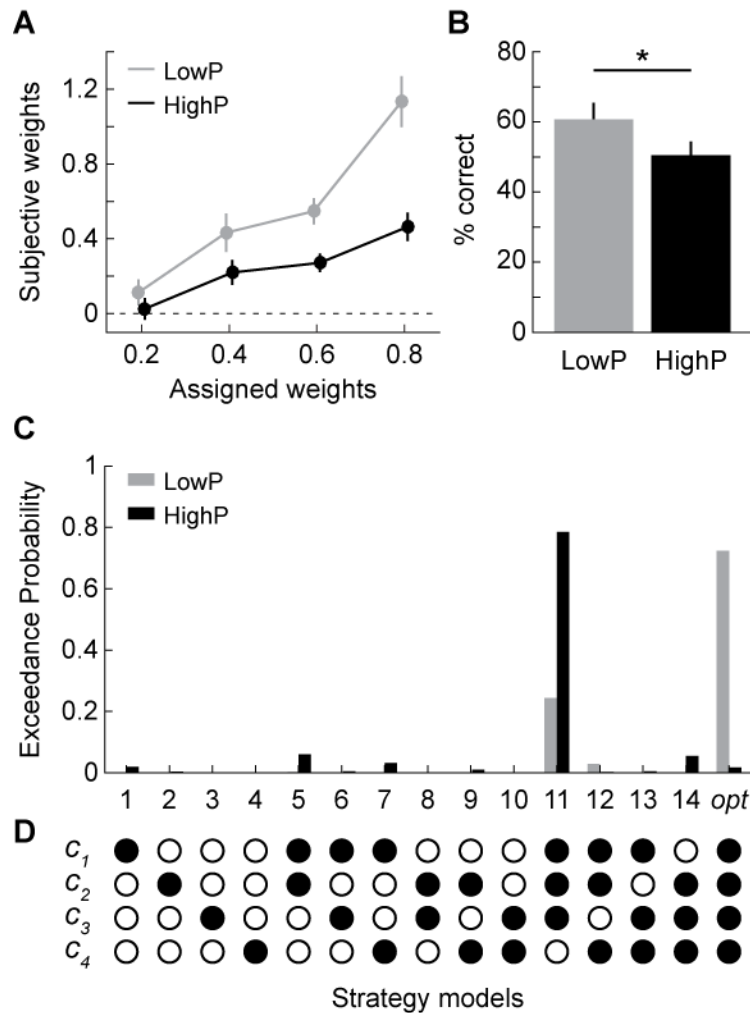


Figure 13: Subjective Cue Weights and Strategy Model Selection. (A) Average subjective cue weights as a function of objective, pre-assigned cue weights. **(B)** Percentage of correct choices when only the least important cue is different between the stimulus pair. Error bars indicate SEM; * $p < 0.01$. **(C)** Bayesian model selection group results presented in exceedance probabilities under low and high time pressure conditions. **(D)** Decision strategy models used in model comparison. Filled circles denote the cue dimensions that are included in a given model.

To further investigate whether this shift in strategy was simply due to running out of time to integrate all cues under severe time pressure, rather than due to a strategic neglect of the weakest cue, we examined performance on trials when only the least

important cue, c_4 , was different between the stimulus pair. Given the visual saliency of only one cue differing between the two stimuli, and the fact that only that one cue (c_4) had to be evaluated in these trials, we would not expect a participant to run out of time in this condition. However, as shown in Figure 13B, the difference in choice accuracy between the LowP and HighP phases was significant, $t_{24} = 2.86, p = 0.009$, with HighP phase performance not differing significantly from chance, $t_{24} = 0.13, p = 0.90$. These results further support the conclusion that the shift in cue usage was a strategic choice of participants to restrict their search space and therefore to integrate less cues to arrive at good-enough decisions under severe time pressure. Having established that time pressure produced noisier decisions and a dropping of the weakest cue, we turned to ask how these changes in decision making are reflected in brain activity.

5.2.2 Neuroimaging data

5.2.2.1 Neural representation of subjective cue weights

Once participants have formed an understanding of the cue structure, one way to effectively solve the classification task is to compare a given compound stimulus pair and extrapolate the differences in cue dimensions using subjective cue weights. Then, one can use this cue information to arrive at a final choice directly (e.g., when one cue out-weighs the rest of available cue(s)) or sum the weights to arrive at the (subjective) SoE. Hence, this process requires breaking down compound stimuli into individual cue dimensions and weighting them appropriately, which serves as a precursor to the

estimation of subjective SoE and/or a final choice output. Therefore, we first sought to examine which brain regions contained information about the subjective importance of each cue dimension in solving the classification task. We employed a whole-brain searchlight SVR to find BOLD patterns that are significantly predictive of the magnitude of subjective cue weights in each experimental phase. As our behavioral data suggest, participants used varying decision strategies under low and high time pressure, with increased decision noise and a reduced information search space under high pressure. This change is represented by relative over-weighting of the cues that each participant deemed important and down-weighting the cues that are considered less informative (Figure 13A). Therefore, to account for large individual differences in cue usage, we focused our neuroimaging analyses on using subject- and phase-specific subjective cue weights (or the sum of weights represented by subjective SoE; see below), rather than fixed objective cue weights, to achieve higher sensitivity in detecting brain regions that encode decision evidence. Hence, our decoding approach using a linear SVR aimed to delineate the brain regions that are associated with the subjective evaluation of individual cue dimensions under each experimental condition.

The results are summarized in Figure 14 and Table 5 ($p < 0.05$, corrected). Under low time pressure, subjective cue weights could be successfully decoded from neural signals in the dorsal striatum, dorsolateral prefrontal cortex (dlPFC), posterior cingulate cortex (PCC), and precuneus (Figure 14A). Under increased time pressure, however,

above-chance decoding of cue weight information was found in the postcentral gyrus and cerebellum (Figure 14B). Furthermore, there were significant differences in decoding accuracies between task phases in the dorsal striatum ($t_{24} = 3.26, p = 0.003$), dlPFC ($t_{24} = 2.45, p = 0.02$), and PCC ($t_{24} = 4.45, p < 0.001$), indicating deterioration in individual cue weight information in these clusters under increased time pressure (Figure 14C). Interestingly, we found significantly enhanced encoding of cue information in the cerebellum in the HighP phase ($t_{24} = 3.84, p < 0.001$). Thus, in line with deterioration in behavioral decision making and down-weighting of subjective cue weights, encoding of individual cue weight information in fronto-parietal and striatal regions under low pressure was dampened under high pressure. Note that the results remain consistent over variations of the regularization parameter, C , in SVR. We next turned to interrogating how this loss of fidelity of subjective cue weight representations for specific cue dimensions may be accompanied by changes in the neural encoding of the SoE or decision variable.

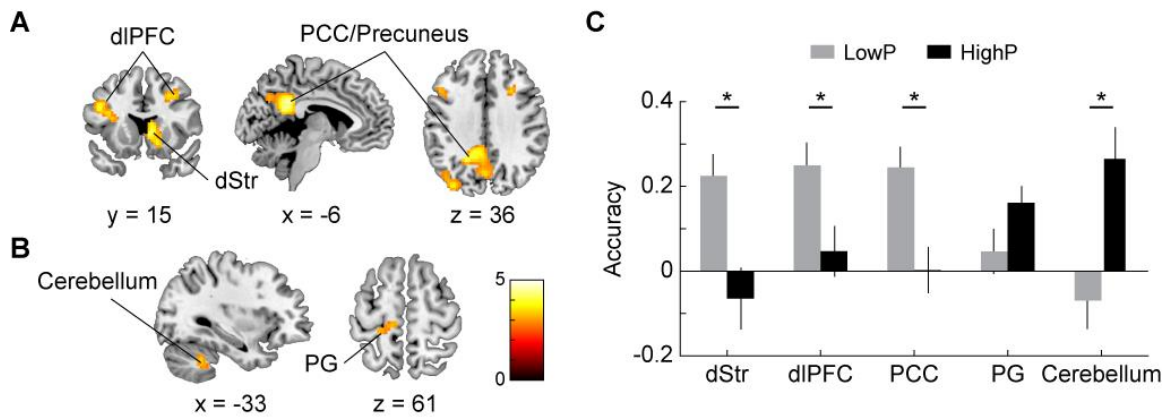


Figure 14: Neural Representation of Subjective Cue Weights. (A) Brain regions significantly predictive of subjective cue weights under low time pressure. (B) Brain regions significantly predictive of subjective cue weights under high time pressure. All maps are $p < 0.05$ corrected. Image display according to neurological convention (left is left). (C) Average prediction accuracy (Fisher's Z-transformed correlation coefficients) of each cluster. Accuracy of zero represents at chance decoding performance. dStr = dorsal striatum; PG = postcentral gyrus; Error bars indicate SEM; * $p < 0.05$.

Table 5: Clusters Showing Significant Decoding Accuracy of Subjective Cue Weights.

| Region | Peak MNI (x, y, z) | Peak t | Cluster size (searchlights) |
|-------------------------------------|---------------------------|-------------|--------------------------------|
| Low time pressure | | | |
| L/R. Precuneus/posterior cingulate | (-6, -46, 29) | 4.73 | 1219 |
| R. Caudate/putamen | (9, 14, 2) | 4.29 | 160 |
| L. Middle/inferior frontal gyrus | (-45, 17, 29) | 4.07 | 108 |
| R. Middle frontal gyrus | (27, 17, 38) | 4.00 | 86 |
| High time pressure | | | |
| L. Postcentral/medial frontal gyrus | (-12, -25, 59) | 3.58 | 65 |
| L. Cerebellum | (-33, -49, -37) | 3.45 | 66 |

5.2.2.2 Modulation of neural activity by SoE

The magnitude of SoE represents the log odds of observing a positive outcome from choosing either a left or right compound stimulus. Therefore, it represents the total amount of evidence available to a participant, which in turn, determines the objective difficulty of a given trial. That is, uncertainty of outcome increases with decreasing SoE, which makes decision-making more difficult. Similarly, subjective SoE reflects participants' perceived difficulty of a given trial. Thus, we hypothesized that brain regions modulated by SoE should encode both the integrated sum of cue weights and associated uncertainty/difficulty that produce the final choice output. Specifically, to account for individual variability in cue learning and decision strategies employed under low and high time pressure conditions, we sought to examine the areas modulated by *subjective* SOE, which was estimated using the sum of subjective weights of the cue dimension present on a given trial (see Methods). To this end, we employed the standard GLM approach using trial-by-trial subjective SoE as a parametric modulator.

The results are summarized in Figure 15 and Tables 6 and 7 ($p < 0.05$, corrected). We found distinct sets of brain regions displaying either a positive modulation effect (Table 6), reflecting increasing activity with greater decision evidence and certainty, or a negative modulation effect (Table 7), where activity increased with greater decision difficulty and uncertainty (i.e., decreasing magnitude of SoE). In the LowP phase,

positive scaling with SoE magnitude was found in the vmPFC, extending to the rostral anterior cingulate and superior frontal gyrus; in superior temporal gyrus, extending to the insula, putamen and hippocampus; and in precuneus and PCC (Figure 15A, Table 6). By contrast, negative SoE modulation was found in the dlPFC, dorsal medial PFC, and lateral inferior parietal lobule (IPL) (Figure 15A, Table 7), a set of regions that are commonly activated under high attentional demand and are often referred to as the fronto-parietal cognitive control network (Dosenbach et al., 2006; Duncan & Owen, 2000; Niendam et al., 2012; Wager, Jonides, & Reading, 2004). Under increased time pressure, similar areas positively modulated by subjective SoE in the LowP condition were found, with more confined activation in the putamen, vmPFC, and PCC. In addition, areas related to motor control such as the supplementary motor area (SMA), cerebellum, precentral and postcentral gyrus also showed enhanced modulation (Figure 15B, Table 6). Again, similar to the LowP phase, significant negative modulation effects were observed in the right dlPFC and IPL (Figure 15B, Table 7), suggesting reduced engagement of the fronto-parietal control network when choices have to be made quickly under severe time pressure.

When testing for regions with similar response profiles across low and high time pressure by means of a conjunction analysis, positive modulation of SoE was commonly found in the putamen, vmPFC, PCC, precuneus, parahippocampal gyrus, and insula. Additionally, negative modulation of SoE was commonly observed in the dlPFC and

IPL, suggesting a critical role of these regions associated with an accumulation of evidence or subjective confidence regardless of time pressure. Finally, and most importantly, to determine the way in which time pressure alters decision making in terms of the modulation of neural activity by decision evidence, we directly contrasted the modulation effect of subjective SoE between the LowP and HighP phases (Figure 15C, Table 6). The results revealed a significantly *greater* positive modulation of activity by SoE in the putamen, thalamus, dopaminergic midbrain, and cerebellum during the HighP compared to the LowP phase. The midbrain cluster consisted of substantia nigra (SN, 83 voxels) and ventral tegmental area (VTA, 38 voxels), which were identified based on a probabilistic atlas (Murty et al., 2014). In sum, the results highlight enhanced sensitivity in the basal ganglia, thalamus, and cerebellum to decision-relevant evidence under high time pressure.

Since participants' performance scaled linearly with objective SoE (see Figure 12B), we sought to examine whether the regions modulated by subjective SoE are similarly sensitive to objective SoE. That is, since objective and subjective SoEs are highly correlated, especially for high performers, we expected that the similar regions would also be modulated by objective SoE. We conducted an additional ROI analysis based on the clusters sensitive to subjective SoE modulation identified through the conjunction analysis. We applied anatomical masks to separate clusters that span multiple areas, and clusters with more than 10 voxels were used for ROI analysis. This

included the putamen (70 voxels), vmPFC (57 voxels), and PCC (59 voxels) from the positive SoE modulation effect, and the dlPFC (81 voxels) and IPL (84 voxels) from the negative SoE modulation effect. For both the low and high time pressure phases, parameter estimates, β_{SoE_i} ($i = 0, \dots, 1.6^*$; $SoE = 1.6^*$ includes trials from the three highest SoE levels, 1.6, 1.8, and 2), from all five regions showed highly significant linear trend as a function of objective SoE. Specifically, the putamen (LowP: $F_{(1,24)} = 21.28, p < 0.001$; HighP: $F_{(1,24)} = 19.73, p < 0.001$) (Figure 15D), vmPFC (LowP: $F_{(1,24)} = 20.63, p < 0.001$; HighP: $F_{(1,24)} = 13.56, p = 0.001$) (Figure 15E), and PCC (LowP: $F_{(1,24)} = 21.47, p < 0.001$; HighP: $F_{(1,24)} = 9.10, p = 0.006$) showed significantly positive linear relationship with objective SoE. Similarly, both the dlPFC (LowP: $F_{(1,24)} = 13.97, p = 0.001$; HighP: $F_{(1,24)} = 5.88, p = 0.02$) and IPL (LowP: $F_{(1,24)} = 17.39, p < 0.001$; HighP: $F_{(1,24)} = 9.03, p = 0.006$) demonstrated significantly negative linear relationship with objective SoE. The results suggest that these regions closely track the integrated sum of decision-relevant information on each trial at both subjective and objective level, regardless of time pressure.

Lastly, to ensure that our results were not simply driven by differences in RT profiles in the LowP and HighP phases, we conducted additional analyses. For all GLMs reported above, trial-by-trial RTs were modeled explicitly as a parametric modulator, and we found no significant differences in RT modulation across phases. Additionally, evidence accumulation models such as the drift diffusion model (e.g., Ratcliff, 1978)

suggest that time pressure induces lowering of the decision bound, which would in turn result in a change in the magnitude of the subjective cue weights but not their ratios. To account for this prediction, we analyzed fMRI data using normalized subjective cue weights, and found no difference in results, both in the subjective cue weights (SVR) and the subjective SoE (parametric GLM) analyses. Therefore, it is highly unlikely that the phase differences in neural cue weight and SoE encoding reported above are merely reflective of RT differences between conditions.

Taken together, results from the SoE analysis showed that activity in the putamen, dopaminergic midbrain (SN/VTA), thalamus, and cerebellum is positively correlated with trial-by-trial variations in subjective SoE, and display a significantly greater sensitivity to the decision variable under severe time pressure than under low time pressure. Furthermore, regardless of time pressure, the putamen, vmPFC, and PCC were positively modulated by subjective SoE, whereas activity in the fronto-parietal cognitive control network was consistently sensitive to decision uncertainty, characterized by negative SoE modulation. Additionally, we found that these areas also closely tracked objective SoE, confirming their role in encoding integrated decision-relevant information.

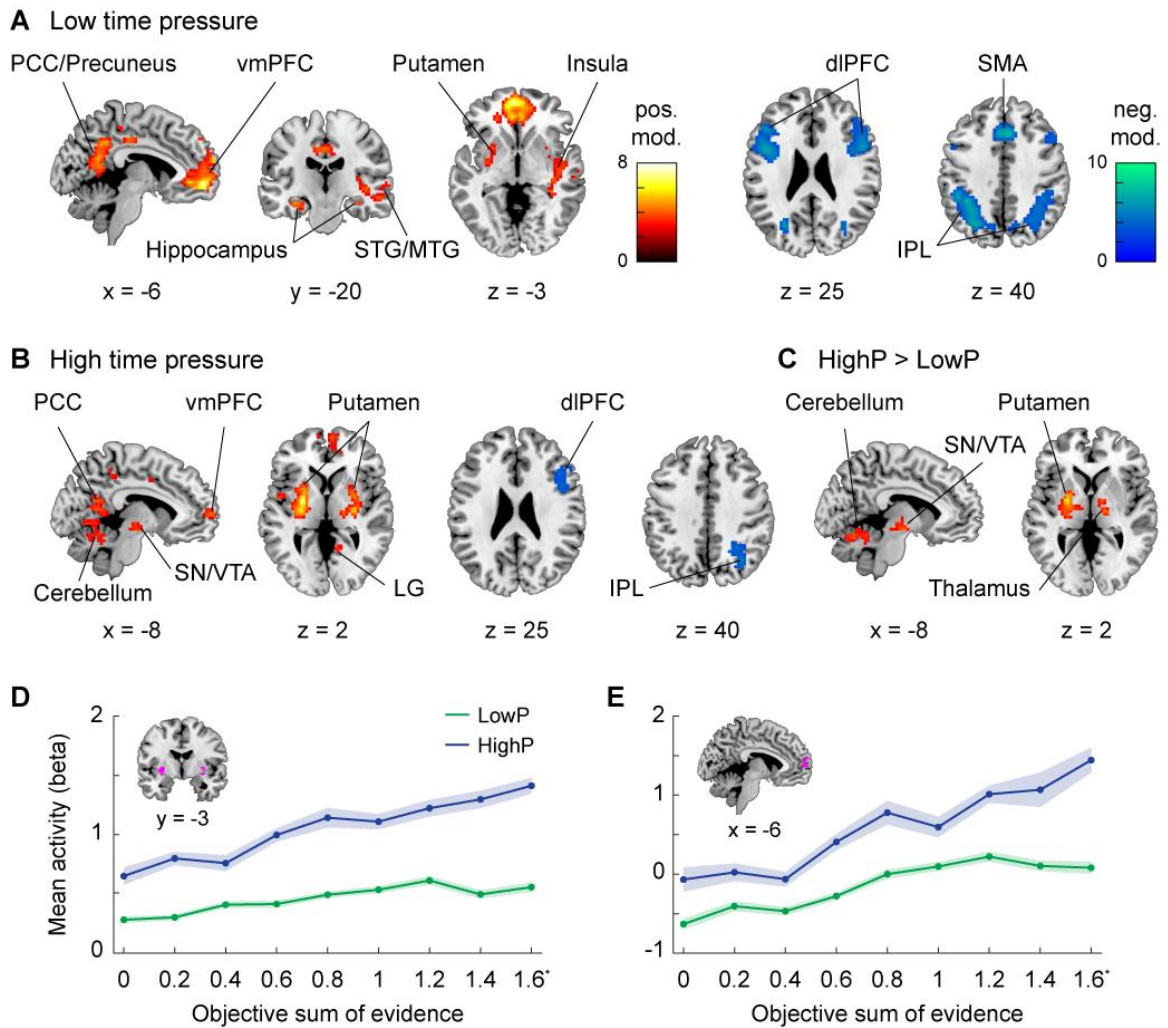


Figure 15: Modulation of Neural Activity by SoE. (A) Significant clusters under LowP. (B) Significant clusters under HighP. Positive modulation (red) indicates activity increasing with accumulation of evidence, whereas negative modulation (blue) represents activity increasing with decreasing SoE, and hence, greater uncertainty or decision difficulty. (C) Brain areas demonstrating significantly greater positive SoE modulation in the HighP compared to the LowP phase. STG = superior temporal gyrus; MTG = middle temporal gyrus; LG = lingual gyrus. All maps were whole-brain corrected to $p < 0.05$. (D) Mean activity as a function of objective SoE in the putamen. (E) Mean activity as a function of objective SoE in the vmPFC. Objective SoE of 1.6* includes trials with SoE levels ranging from 1.6 to 2. Shaded area indicates SEM.

Table 6: Activation Clusters for Positive Subjective SoE Modulation.

| Region | Peak MNI (<i>x, y, z</i>) | Peak <i>t</i> | Cluster size (voxels) |
|---|--------------------------------|------------------|--------------------------|
| Low time pressure | | | |
| L/R. Medial frontal gyrus/anterior cingulate | (-6, 56, -7) | 7.33 | 725 |
| L/R. Precuneus/posterior cingulate | (-9, -52, 38) | 5.86 | 308 |
| R. Superior temporal gyrus/insula/ hippocampus/putamen | (54, -28, 14) | 5.66 | 771 |
| L. Superior/middle frontal gyrus | (-12, 47, 41) | 5.35 | 155 |
| L. Superior/middle temporal gyrus/insula | (-60, -58, 20) | 5.25 | 388 |
| L. Parahippocampal gyrus/putamen/ hippocampus | (-33, -19, -16) | 4.60 | 107 |
| L. Cingulate gyrus | (-9, -25, 41) | 4.74 | 117 |
| R. Cuneus | (18, -85, 20) | 4.39 | 111 |
| High time pressure | | | |
| L. Putamen/insula/parahippocampal gyrus/ midbrain/cerebellum | (-27, -4, 2) | 6.87 | 1088 |
| R. Putamen/insula/parahippocampal gyrus | (30, -16, 5) | 5.51 | 389 |
| L. Posterior cingulate | (-12, -58, 14) | 5.04 | 113 |
| L/R. Postcentral/precentral gyrus/precuneus | (27, -52, 62) | 4.76 | 377 |
| R. Posterior cingulate/lingual gyrus/ cerebellum | (15, -55, 11) | 4.66 | 199 |
| L/R. Medial frontal gyrus | (-9, 65, 5) | 4.35 | 103 |
| L/R. Supplementary motor area | (-3, -13, 44) | 4.34 | 99 |
| R. Precentral/postcentral gyrus | (27, -25, 44) | 3.92 | 96 |
| High time pressure > Low time pressure | | | |
| L. Cerebellum/lingual gyrus | (-24, -61, -22) | 6.20 | 541 |
| L. Putamen/thalamus | (-24, -4, 2) | 5.76 | 241 |
| R. Thalamus/midbrain | (18, -16, -1) | 4.93 | 151 |

Table 7: Activation Clusters for Negative Subjective SoE Modulation.

| Region | Peak MNI (<i>x, y, z</i>) | Peak <i>t</i> | Cluster size (voxels) |
|----------------------------------|--------------------------------|------------------|--------------------------|
| Low time pressure | | | |
| L/R. Medial frontal gyrus | (-3, 11, 50) | 9.21 | 294 |
| L. Middle/inferior frontal gyrus | (-45, 8, 29) | 7.71 | 374 |
| L. Inferior parietal lobule | (-36, -43, 38) | 7.35 | 576 |
| R. Middle/inferior frontal gyrus | (51, 23, 32) | 5.88 | 387 |
| R. Inferior parietal lobule | (36, -58, 47) | 5.38 | 465 |
| High time pressure | | | |
| R. Inferior parietal lobule | (33, -64, 41) | 4.80 | 110 |
| R. Middle/inferior frontal gyrus | (48, 29, 20) | 4.59 | 99 |

5.2.2.3 Ruling out practice effects

We performed the time pressure manipulation in a block-wise, sequential manner to avoid spillover effect of severe time pressure influencing decision behavior on subsequent low pressure blocks. This experimental design, however, raises a concern that changes in behavior and cortical activity across experimental phases we report could be influenced by task practice. At the behavioral level, practice effects are defined as an increase in accuracy and decrease in response time (Ashby, Turner, & Horvitz, 2010; Kelly & Garavan, 2005; Schneider & Shiffrin, 1977). As reported in the previous section (see *Task performance*), both decision accuracy and RT remained constant within each the LowP and HighP phase during the scan ($ps > 0.05$). In addition, contrary to the assumption of improving performance with continued practice, we observed significant impairment in choice performance when moving from the low to high time pressure

condition. Hence, behaviorally, our within- and between-phase observations do not support a practice effect hypothesis.

At the neural level, practice is usually associated with an overall decrease in BOLD activation of the cognitive control and attentional network (i.e., prefrontal and parietal areas) as the task becomes automatized (see Kelly and Garavan, 2005 for a review). Some studies, additionally, have reported enhanced post-learning subcortical processing (Doyon et al., 2009; Lehericy et al., 2005; Van Turennout, Bielmowicz, & Martin, 2003), although the results are rather mixed (see Ashby et al., 2010 for a detailed discussion). To examine changes in cortical and subcortical activity over time, we divided our experiment into four different sets (120 trials/set, 2 sets/phase), each consisting of a complete set of unique stimulus combinations, and conducted an additional GLM analysis. If practice effects are present, then we would expect to observe gradual BOLD signal decrease in cortical regions, especially in the control network, and steady increase in subcortical regions. Hence, we focused our analyses on seven ROIs sensitive to the subjective SoE modulation: 1) The control network (dlPFC, SMA, and IPL) and the vmPFC defined from the LowP phase SoE modulation, and 2) the putamen, thalamus/midbrain, and cerebellum clusters identified from the HighP > LowP contrast (see Tables 6 and 7). Specifically, to investigate overall changes in trial-induced activation over time, we extracted and averaged the percent signal changes and the parameter estimates, β_{Trial} , in the seven ROIs.

Results from 2 (sets) \times 2 (LowP/HighP phases) ANOVAs based on the average percent signal changes showed no significant effect of phase nor set in all seven ROIs ($ps > 0.05$), suggesting that overall BOLD activation remained stable in these regions throughout the experiment (see Figure 16 and Table 8 in Appendix B). The mean β_{Trial} , however, demonstrated significant effects of phase for cognitive control regions (dlPFC: $F_{(1,24)} = 24.02, p < 0.001$; SMA: $F_{(1,24)} = 15.66, p = 0.001$; IPL: $F_{(1,24)} = 26.44, p < 0.001$) as well as subcortical regions (putamen: $F_{(1,24)} = 4.59, p = 0.04$; thalamus/midbrain: $F_{(1,24)} = 5.97, p = 0.02$) and the cerebellum ($F_{(1,24)} = 23.23, p < 0.001$), all of which exhibited an overall increase in trial-evoked response during the HighP phase. No effect of set nor phase \times set interaction was found ($ps > 0.05$). In other words, we did not observe the gradual changes in activity predicted by the practice effect hypothesis, but rather a stepwise increase when moving from the low-pressure to the high-pressure phase.

Finally, to account for within-phase strategy variability influencing our neural data, we repeated decision strategy model selection using four sets (2 sets/phase) (Figure 17A, Appendix B). We observed no changes in the overall winning model (i.e., model with the highest exceedance probability) within each phase. There was, however, increased tendency of switching to Model 11 in the second set of the LowP phase, possibly a sign of satisficing behavior induced by extended exposure to the task. In the first set of the HighP phase, participants adopted an even more constricted cue space,

with exceedance probabilities spreading out to Model 11 (c_1, c_2, c_3), Model 5 (c_1, c_2), and Model 1 (c_1), thereby reflecting substantial behavioral changes triggered by the onset of severe time pressure. These time pressure induced changes in choice strategy are of course in accordance with our “drop-the-worst” proposal.

To ensure that this observed strategy variability does not strongly influence our neural results, we further contrasted fMRI data of the first set of the LowP phase (highest reliability in using the optimal model) versus the first set of the HighP phase (highest probability of adopting sub-optimal satisficing strategy) (see Figure 17B-D and Table 9 in Appendix B). Consistent with our previous report, under low time pressure, the vmPFC and cognitive control network were significantly modulated by subjective SoE, whereas under high time pressure, increased modulation of the subcortical regions was observed. Hence, although some variability in cue usage exists within each phase and participants, the overall picture in terms of neuroimaging results does not differ much when individual sets are being considered as compared to collapsing across sets within each phase.

In sum, both behavioral and neural evidence speak quite strongly against the possibility that the neural correlates of the phase-dependent decision-making shifts we report here are simply reflecting practice effects. The significant between-phase differences, in the absence of gradual behavioral and BOLD signal changes, are more likely to reflect strategy changes triggered by the time pressure manipulation.

5.3 Discussion

To characterize neural mechanisms underlying satisficing decision-making induced by time pressure, we tested participants on a multi-cue probabilistic classification task under conditions of low (1500 ms) and high (500 ms) pressure during fMRI. Specifically, we focused on investigating post-learning decision performance, after participants had formed an understanding of the predefined cue structure through trial-and-error feedback learning. Using subjective cue weights and SoE inferred from a variational Bayesian approach, we demonstrate that, under low time pressure, participants accurately ranked the cues and integrated all available information to achieve near-optimal performance. Under this condition, distributed fMRI patterns in the dorsal striatum, dlPFC, and PCC could be used to make linear predictions about the subjective cue weights. Additionally, the total amount of evidence, quantified by subjective SoE, was encoded in the vmPFC, PCC, putamen, and hippocampus, whereas decision uncertainty was represented in the fronto-parietal cognitive control network. Under high time pressure, by contrast, participants adopted a drop-the-worst satisficing strategy by discounting the least important cue from their decision process, which was accompanied by an overall deterioration of performance and cue information encoding in the brain. Despite this degradation of information processing, we found significantly greater subjective SoE modulation under high time pressure in the putamen, dopaminergic midbrain, thalamus, and cerebellum. These findings suggest that time

pressure triggers a shift from near-optimal decision-making dependent on both subcortical and fronto-parietal regions to satisficing decision-making characterized by a greater involvement of the midbrain, striatum, thalamus, and cerebellum.

Learning to perform probabilistic classification is thought to mainly rely on an interplay between the striatum and the hippocampus (Knowlton et al., 1996; Packard & Knowlton, 2002; Squire & Zola, 1996). The striatum serves an essential role in acquiring probabilistic cue-outcome associations in the weather prediction task (Knowlton et al., 1996; Poldrack et al., 2001) as well as other choice tasks that involve learning from reinforcement (Daw & Doya, 2006; Delgado, Nystrom, Fissell, Noll, & Fiez, 2000; J P O'Doherty, Dayan, Friston, Critchley, & Dolan, 2003; Schultz, 1998; Schultz et al., 1997). Stress, induced by a socially evaluated cold pressor test, has been shown to trigger increased engagement of striatum-dependent implicit learning and impair the use of the hippocampus-dependent declarative system during probabilistic classification learning (Lars Schwabe & Wolf, 2012). This shift to striatum-mediated learning has also been found when participants performed the weather prediction task under cognitive load (Foerde et al., 2006). In studies of the speed-accuracy tradeoff, increased activation was found in the striatum during trials with speed emphasis compared to trials with accuracy emphasis (Forstmann et al., 2008; van Veen, Krug, & Carter, 2008), suggesting that enhanced striatal activity may be critical in reducing inhibitory control and facilitating speeded responses (Bogacz, Wagenmakers, et al., 2010). Here, we

significantly expand these findings by directly demonstrating that increased time pressure decreases the involvement of frontal and parietal regions and shifts probabilistic multi-cue decision-making to the striatum, even after the completion of initial cue learning.

In addition to the striatum, the midbrain (SN/VTA), and thalamus demonstrated significantly enhanced encoding of subjective SoE under high time pressure. Dopaminergic neurons in the SN/VTA are associated with expected rewards (D'Ardenne, McClure, Nystrom, Leigh, & Cohen, 2008; Schultz et al., 1997) and control of motor responses (e.g., Hikosaka, 1989; Chevalier and Deniau, 1990; Mink, 1996) in decision-making. The thalamus serves as a critical relay structure between cortical and subcortical regions, facilitating information integration among the SN/VTA and striatum as well as prefrontal areas such as the vmPFC and dlPFC (Haber and Knutson, 2009). The vmPFC has been widely implicated in subjective value encoding (Daw et al., 2006; Kable & Glimcher, 2007; D. J. Levy & Glimcher, 2012; P. Read Montague & Berns, 2002; Padoa-Schioppa & Assad, 2006; Plassmann, O'Doherty, & Rangel, 2007) and has been shown to integrate value information during multi-attribute decision-making (Basten, Biele, Heekeren, & Fiebach, 2010; Hare, Malmaud, & Rangel, 2011; Kahnt et al., 2011). Building on these findings, we here show that vmPFC encodes integrated cue values. Under high time pressure, this SoE encoding was accompanied by enhanced modulation of the midbrain and striatal-thalamic circuit, suggesting that these regions may

particularly facilitate information integration and evaluation when a speeded response is required.

We also found a significant modulatory effect of decision uncertainty on a fronto-parietal network that is consistently recruited during top-down attentional control processes that assist goal-directed behavior under cognitively demanding conditions (Dosenbach et al., 2006; Duncan & Owen, 2000; Niendam et al., 2012; Wager et al., 2004). In a multi-cue decision environment where only a single dimension is relevant, this network is involved in selecting the relevant feature, effectively reducing the dimensionality of the problem (Niv et al., 2015). In the current study, the dlPFC encoded individual cue weights as well as combined SoE, indicating that this region may be directly involved in the cue integration process. This is consistent with prior reports demonstrating engagement of the dlPFC in evidence accumulation during perceptual decision-making (Gold & Shadlen, 2007; Heekeren, Marrett, & Ungerleider, 2008), and encoding of ambiguity or difficulty in multi-attribute decision-making (Kahnt et al., 2011; Krebs, Boehler, Roberts, Song, & Woldorff, 2012). Our results further demonstrate involvement of the fronto-parietal control network in a *feature-invariant* evidence integration processes, possibly guiding attention to cues according to their weighted importance and tracking uncertainty in choices.

Interestingly, the cerebellum was also involved in both cue value encoding and subjective SoE modulation under high time pressure. The cerebellum is traditionally

regarded as contributing to movement planning and execution (Itō, 1984), and its role in cognitive functions is not well understood. Due to this structure's extensive connections to the cerebral cortex, including prefrontal areas (Habas et al., 2009), it has been suggested that the cerebellum may contribute to higher-level cognition, such as learning from feedback, which can support fast and adaptive control of motor behavior (Buckner, 2013; Ito, 2008; Strick, Dum, & Fiez, 2009). Hence, in the current task, the cerebellum may carry decision-relevant signals to facilitate quick motor responses. Although it is difficult to determine the precise role of the cerebellum in satisficing in the current study, the apparent involvement of this structure in fast-paced probabilistic inference represents an interesting starting point for future studies.

Although the use of heuristics is ubiquitous in everyday life, most neuroimaging studies on decision-making have focused on rather simple scenarios, where all decision-relevant information is available and participants are assumed to use a uniform, optimal strategy. Hence, neural mechanisms underlying heuristic processes that are used to simplify complex decision problems are not well understood (Volz and Gigerenzer, 2012). Prior studies that did examine neural correlates of heuristic decisions have largely focused on memory-guided heuristics (Khader et al., 2011, 2015; Rosburg, Mecklinger, & Frings, 2011; Volz et al., 2006, 2010). Given this focus on cached decision strategies and extensive training of explicit task rules, this prior work did not necessarily address how people arrive at decisions in more complex and natural decision environments. To

overcome this problem, we used a large set of non-deterministic cue combinations, encouraging participants to actively integrate available information on each trial. Our study, therefore, provides novel evidence concerning the neural substrates of satisficing decision-making in the context of active, speeded cue integration.

We show that participants discounted the least important cue in their decision processes under high time pressure, consistent with our previous findings (Oh et al., 2016). Participants performed at chance when the least important cue was the only differentiating cue between the two compound stimuli (Fig. 2B), indicating that the use of the drop-the-worst heuristic is a strategic choice, rather than due to having insufficient time to evaluate evidence. Where in the brain might cue information be “dropped” when decisions are made under high time pressure? It is possible that the deterioration in cue processing occurs at the sensory level (i.e., Ho et al., 2012) or alternatively, that it takes place at a later stage of the decision-making process. Due to an insufficient number of trials where only a single cue differentiated the competing stimuli, our current paradigm could not provide an unambiguous answer to this question. Despite this limitation, it is clear from the present data that the dopaminergic midbrain, striatum, and cerebellum encode subjective SoE preferentially under time pressure, and thus are key structures for facilitating satisficing decision-making.

In conclusion, the current study shows that near-optimal performance in multi-cue probabilistic classification under low time pressure is supported by widely

distributed regions including both subcortical (striatum and hippocampus) and cortical (frontal and parietal) areas. Under high time pressure, by contrast, participants adopted the drop-the-worst satisficing strategy, which is characterized by increased involvement of the dopaminergic midbrain, thalamus, striatum, and cerebellum in mediating fast and good-enough decision-making.

6. General Discussion

How do people make decisions in a complex and uncertain world? This question has been a core research topic in many fields of decision science, and has led to diverse theories of decision making, ranging from utility maximization to satisficing. Nonetheless, research in decision neuroscience has been primarily focused on investigating cognitive and neural mechanisms underlying rather simple, small-world decision processes, based on an assumption that all subjects carry out the task by using a uniform (and optimal) decision strategy. Although this approach has laid the foundations for characterizing core concepts of decision making and associated neural correlates (see Section 1.5), it falls short of explaining how complex decisions are made in a large-world setting, in which people exhibit bounded rationality.

Research on satisficing and heuristics, on the other hand, has focused on cognitive processes underlying decision making under uncertainty, situations in which optimization is not feasible. Unfortunately, these two fields – decision neuroscience and satisficing – have remained largely separate disciplines with very little overlap (e.g., Volz & Gigerenzer, 2014). The studies described here, therefore, can be considered as one of a few pioneering attempts to bring the two fields together to develop a coherent, large-world model of human decision-making behavior and its underlying neural mechanisms. Specifically, Chapter 2 presents a novel probabilistic classification task and its accompanying set of strategy analyses based on variational Bayesian inference. Then,

through Chapters 3, 4 and 5, the task and analysis techniques are put to the test in three behavioral studies investigating shifts in decision strategies under split-second time pressure (Chapters 3 and 5) and under heightened anxiety induced by a threat-of-shock manipulation (Chapter 4). Finally, Chapter 5 examines the neural substrates mediating changes in decision-making under time pressure. This concluding chapter will review and synthesize the current empirical findings, discuss possible theoretical and practical implications, and consider limitations of current approaches as well as directions for future research.

6.1 Adaptive satisficing decision making

The primary goal of this dissertation is to characterize human satisficing behavior triggered by various externally (e.g., time) or internally (e.g., anxiety) driven stressors. Through a series of behavioral studies, we found that under low pressure, participants performed near-optimally, appropriately weighing and integrating all available cues. Faced with high pressure, however, participants adopted heuristics by considering only a subset of available decision-relevant information. Importantly, this switch to satisficing decision making was highly adaptive in nature. That is, to cope with increasing pressure, people systematically disregarded cue(s) with the least information value(s) (i.e., Drop-the-Worst), thereby maintaining “good-enough” decision accuracy by focusing exclusively on the most predictive cues.

6.1.1 Near-optimal decision making under low pressure

Under low (or in the absence of) pressure, we repeatedly showed that participants engage in near-optimal decision making (Chapters 3, 4, and 5). That is, as suggested by the linear relationship between subjective and objective cue weights, participants were able to learn the relative importance of cues through trial-and-error, and utilize all decision-relevant information in guiding their choices. Interestingly, this maximization strategy seems to serve as a default choice behavior during learning as well as low pressure conditions, especially when the task involved compound stimuli (cf., Section 3.3). This general tendency to employ compensatory strategies under low pressure is consistent with previous analyses from heuristics studies (e.g., Bröder, 2000; Glöckner & Betsch, 2008; Newell & Shanks, 2003; Pachur & Olsson, 2012) as well as probabilistic classification learning experiments (e.g., Gluck et al., 2002; Lagnado et al., 2006; Meeter et al., 2006). Hence, the present empirical findings seem to suggest that, under low pressure, people default to maximization strategies, weighting and combining all available evidence to achieve high decision accuracy. This conclusion, however, should be interpreted with caution as such decision behavior is highly dependent on the task structure.

The multi-cue probabilistic classification task used in all three empirical studies specifically models a compensatory decision environment, in which the highest-value cue (c_1) can be out-weighed by some combinations of less valuable cues (see Table 2).

Specifically, in this environment, the optimal cue integration strategy yields the best decision accuracy (see Figure 2). Previous studies have shown that if a decision task could be solved via multiple different strategies, each leading to varying levels of accuracy, people tend to employ the most accurate strategy (Juslin et al., 2003; Rieskamp & Otto, 2006). In addition, when all cues are presented on the screen, and therefore, costs for acquiring information are rather low, people are more likely to adopt a maximization strategy (Bröder, 2000; Rieskamp & Otto, 2006), especially when the number of cues are relatively small (i.e., four cues in our task; Pachur & Olsson, 2012). Furthermore, a linear cue structure with an evenly spaced weight distribution (cf., Mata, Schooler, & Rieskamp, 2007) as well as a lack of information redundancy (i.e., correlation between cues; cf., Dieckmann & Rieskamp, 2007) could also shape decision strategies toward maximization. Near-optimal decision making under low pressure, therefore, may have largely been fostered by the compensatory structure of our task environment. Nonetheless, this observation is in line with adaptive decision theories (e.g., Gigerenzer & Goldstein, 1996; Payne et al., 1988; Shah & Oppenheimer, 2008), which predict that people select a strategy that performs the best in a given environment (i.e., if a task environment is noncompensatory, people tend to conform to a matching noncompensatory strategy; see Rieskamp & Otto, 2006).

In Chapter 5, we further extend these results by documenting neural correlates of near-optimal choice behavior under low time pressure. That is, in addition to the

striatum and hippocampus – key subcortical regions mediating probabilistic classification learning (Knowlton et al., 1996; Poldrack et al., 2001) – decision making under low pressure involved the vmPFC and the fronto-parietal attentional control network, all of which linearly tracked trial-by-trial variations in the amount of evidence and decision uncertainty. The involvement of these cortical regions suggests that near-optimal decision making, at least in part, is guided by top-down executive processes to build a better internal task representation, by combining subjective evidence to predict an outcome (Daw et al., 2006; D. J. Levy & Glimcher, 2012; John P. O’Doherty et al., 2001) and selectively updating cue knowledge through feedback (Akaishi et al., 2016; Niv et al., 2015) (for additional discussion, see Sections 5.3 and 6.4).

In sum, given a compensatory task environment, we report that people readily adopt the near-optimal decision strategy under low pressure, which is mediated by widely distributed cortical and subcortical regions in the brain. That is, our neuroimaging results suggest that decision-relevant information is encoded in both the stimulus-response habit learning system (the striatum) as well as the central executive-dependent goal-directed learning system (the prefrontal and parietal cortex) (Daw, Gershman, Seymour, Dayan, & Dolan, 2011; O’Doherty, Lee, & McNamee, 2015), indicating a possible interplay between multiple learning systems to optimally utilize available cue information to guide decision making.

6.1.2 Satisficing under externally and internally induced pressures

Under increasing pressure, we found that participants adopt the Drop-the-Worst heuristic by ignoring the least important pieces of information in their decision processes. Interestingly, both externally (i.e., time) and internally (i.e., anxiety) induced pressures led to the same satisficing behavior, suggesting that this shift to the Drop-the-Worst heuristic is not a condition-specific behavior but potentially a universal satisficing strategy people employ to cope with various stressors. Through Chapters 3 and 4, we showed that this strategy shift under severe pressure is highly adaptive in nature. First, as discussed previously, the Drop-the-Worst heuristic achieves good-enough performance by considering only a subset of cues with the highest validity (i.e., compare accuracies for Models 1, 5, 11, and *Opt* in Figure 2). That is, although sub-optimal, the Drop-the-Worst produces the maximal decision accuracy within a given set size limit (i.e., number of cues considered). Second, the application of the Drop-the-Worst heuristic is highly flexible. In Chapter 3, we found that with increasing time pressure, participants shift their strategies from Model 11 to Model 5, and then to Model 1, adaptively decreasing the number of cues considered. This in turn, speeds up the decision process to enable a choice output within the externally imposed time limit, while minimizing loss in accuracy. Lastly, the adoption of the Drop-the-Worst heuristic is dependent on the amount of pressure. Extending from the previous observations, we found that the size of the set of cues being considered decreases in accordance with

increasing time pressure, indicating that the intensity of heuristics is scaled by the amount of pressure. In Chapter 4, we additionally showed that the mere presence of pressure does not necessarily imply that participants would satisfice. That is, we observed shifts in decision strategy only in those subjects who displayed an objective physiological response (i.e., increase in skin conductance level) to the threat-of-shock manipulation.

6.1.3 Limitations of the current approach

In all three empirical studies, we examined shifts in decision strategies based on incremental changes to stressors without any counterbalancing across phases, and therefore, one could argue that the observed switch to noncompensatory heuristics is induced by the effect of learning and experience. Previous studies have reported that, in contrast to novices who tend to adopt a compensatory strategy, experts tend to rely more heavily on one-reason (or lexicographic) decision making in solving a highly trained decision problem (Garcia-Retamero & Dhimi, 2009; Pachur & Marinello, 2013). Although this is certainly a possibility, especially since we provided feedback throughout all phases of the experiment, our behavioral and neural evidence suggests otherwise. Specifically, in Chapter 4, participants who exhibited low to moderate anxiety response continued to employ the optimal cue integration strategy throughout the task. In addition, these groups of participants were more likely to exhibit enhanced concordance between their implicit strategy and explicit knowledge over the course of

the experiment. Taken together, these behavioral results indicate that the continued learning effect may facilitate formation of better explicit task representation, but does not necessarily encourage satisficing decision making in our task paradigm. These observations are consistent with findings from the weather prediction task, which indicate increased employment of a maximization strategy over the course of learning (Gluck et al., 2002; Lagnado et al., 2006; Meeter et al., 2006). This experience-based improvement is also associated with enhanced explicit task knowledge and self-insight to participants' own judgment processes (Lagnado et al., 2006).

In addition to the behavioral data, our fMRI results in Chapter 5 showed that there is no neural evidence supporting the task practice effects (see Section 5.2.2.3). In fact, satisficing decision making under severe time pressure was primarily characterized by a significant enhancement of the subcortical dopaminergic areas in tracking trial-by-trial variations in decision evidence, which suggests a reversal to the habit learning system, often emphasized in acquiring probabilistic cue-outcome associations (Knowlton et al., 1996; Poldrack et al., 2001). It is, however, difficult to dissociate whether the observed cortical-to-subcortical shift is an inherent mechanism supporting satisficing in general, or is driven by the specific nature of the task. For instance, neuroimaging data on memory-guided heuristics offer rather mixed results, emphasizing the role of the medial PFC (recognition heuristic; Volz et al., 2006) or the dlPFC (Take-the-Best heuristic; Khader et al., 2011, 2015) in mediating the use of

heuristics. The probabilistic task structure (see Shohamy, Myers, Kalanithi, & Gluck, 2008) as well as strict time pressure enforcing speeded response (see Bogacz, Wagenmakers, Forstmann, & Nieuwenhuis, 2010) in the current task, therefore, may have preferentially increased reliance on the subcortical dopaminergic system. Thus, questions such as how satisficing is achieved through task-specific neural processes, and whether there exists a task-independent region triggering satisficing decision making, remain open to future investigation.

6.2 Theoretical and methodological implications

Prior studies on bounded rationality have primarily focused on characterizing processes underlying memory- and cue-based heuristics, as reviewed in Section 1.3. Specifically, choices made based on multiple pieces of information have been extensively investigated either through exemplar-based mechanisms (memory-based; e.g., Juslin et al., 2003; Lamberts, 2000; Nosofsky & Palmeri, 1997) or through lexicographic rules such as the Take-the-Best heuristic (cue-based; e.g., Bergert & Nosofsky, 2007; Gigerenzer & Goldstein, 1996; Rieskamp & Otto, 2006). Many of these models, however, fall short of capturing uncertainty in decision making, which is one of the core challenges people face in real-life choices (see Platt & Huettel, 2008). That is, manipulations such as presenting explicit cue validities (e.g., Bröder, 2000; Payne et al., 1988; Rieskamp & Hoffrage, 2008), providing deterministic feedback (e.g., Bergert & Nosofsky, 2007; Juslin et al., 2003; Lamberts, 1995), employing a serial search paradigm (e.g., Glöckner &

Betsch, 2008; Pachur et al., 2014; Payne et al., 1996), and/or conducting an extensive task training (e.g., Khader et al., 2011; Pachur & Olsson, 2012) play a big role in reducing uncertainty in cue information, thereby facilitating (or even eliminating) the arduous process of acquiring the statistical structure of the decision environment through trial-and-error. These experimental paradigms, thus, do not necessarily mirror real-life scenarios in which decisions are made based on incomplete information or approximate knowledge of the information structure.

The multi-cue probabilistic classification task combined with variational Bayesian inference techniques, therefore, provides a novel methodological framework to examine changes in decision strategies in a controlled, but naturalistic environment. That is, the task allows for modelling each individual's internal representation of the task structure developed through statistical learning, and quantifying usage of cue information under various experimental conditions. Furthermore, the Bayesian inference process does not require "probe trials", where certain choices are mapped on to a specific strategy (e.g., Bergert & Nosofsky, 2007; Lee & Cummins, 2004), or a separate strategy tracking system such as the MouseLab program (Payne et al., 1988). Thus, instead of relying on an artificial decision environment, the task could be flexibly modified to utilize diverse types of stimuli and trials. Therefore, the methods presented in Chapter 2 can be applied to investigate various decision problems under uncertainty, potentially involving parameters such as large number of cues, high information

redundancy, or high dispersion of cue weights. Hence, integrating the current task with well-established decision problems may greatly help advance our understanding of satisficing decision making in the real world.

The most important theoretical implication of present results is that people do not heavily rely on one-reason heuristics but rather opt to consider as many cues as possible within a given resource limit. Thus, when all decision-relevant information can be examined simultaneously, it is unlikely that people engage in a serial search heuristic, which examines cues in the order of their validities and makes a choice based on a single clever cue (e.g., Gigerenzer & Goldstein, 1996). On the contrary, our findings suggest that participants are more likely to set appropriate decision thresholds according to the task condition, and integrate cues included within that satisfactory cutoff levels to produce choices. This Drop-the-Worst heuristic model is consistent with Simon (1955)'s initial formulation of satisficing, and makes similar choice predictions as evidence accumulation models such as the drift diffusion model (e.g., Ratcliff, 1978). Our findings, therefore, significantly advance previous models of heuristics and enhance our understanding of decision making in real life, all of which in turn could assist both healthy and clinical populations in making better choices.

6.3 Practical implications

Initiated by the heuristics and biases program (e.g., Tversky & Kahneman, 1974), understanding human decision making behavior in large worlds, especially

characterizing situations in which rationality fails, has given rise to behavioral research that investigates techniques to indirectly shape people's choice behavior. Programs such as nudging (Thaler & Sunstein, 2009) or boosting (Grune-Yanoff & Hertwig, 2016) have generated various approaches to intervene in decision making processes of laypeople as well as professionals to improve (or simply change) choice outcomes, by manipulating the decision environment and/or enriching individual's knowledge and decision-making skills. Theories generated by these programs have been applied to designing diverse policies, many of which successfully induced desired behavioral outcomes (e.g., Alemanno, 2012; Guthrie, Mancino, & Lin, 2015; Hanks, Just, Smith, & Wansink, 2012), although long-term positive effects of such policies in improving overall decision making abilities are still debatable (Grune-Yanoff & Hertwig, 2016). Hence, understanding specific cognitive processes underlying suboptimal decisions may help develop approaches to target and enhance choice behavior that could generate lasting, positive changes.

Another important issue in decision neuroscience, especially in clinical research, is preventing repeated failures in decision making, which could result in dire consequences. A clearer understanding of cognitive and neural mechanisms mediating complex real-world decision making, therefore, is of utmost importance for advancing neuropsychiatric research and achieving progress in the treatment of clinical symptoms related to such deficits. For instance, anxiety disorders can be comorbid with other

problems such as pathological gambling, alcohol and substance abuse, and nicotine dependence (Petry, Stinson, & Grant, 2005). Although the exact cause of these comorbidities is elusive, deficits in decision-making in these patients could partly be explained by impaired feedback-based learning (Petzold et al., 2010) and attentional control (Bar-Haim et al., 2007), which may in turn lead to inappropriate or maladaptive application of choice strategies. Indeed, if a common behavioral or neural trend of these deficits could be identified, it might be possible to develop diagnostic and training tools to assist patients in making better decisions based on simple heuristics.

Lastly, models of human satisficing behavior have the potential of greatly advancing the control strategies for engineered autonomous systems. The vast majority of current control systems are designed based on maximization strategies, such as Bayesian inference, optimal control, and optimal detection theory, all of which require perfect knowledge of the decision environment (e.g., Ferrari & Cai, 2009). Thus, to achieve the desired level of performance, autonomous agents are often programmed to perform exhaustive search and extensive computations. Unfortunately, these optimal strategies are susceptible to failures when some of the information is incomplete or highly uncertain, combined with limited time or computational resources (e.g., Czerlinski et al., 1999; Martignon et al., 2008; Wübben & Wangenheim, 2008). Satisficing strategies (e.g., Simon & Kadane, 1975), therefore, can serve as an ideal alternative, which will complement the limitations of optimal algorithms, especially under

challenging environmental pressures. Hence, theoretical advancement of human satisficing behavior would enable the development of precise mathematical models of satisficing, which could be implemented in future autonomous agents.

6.4 Directions for future research

A key concept that has been repeatedly emphasized throughout this dissertation is an adaptive exploitation of the environmental structure, which enables successful satisficing decision making. The present findings, however, do not fully examine this concept due to the use of fixed, independent cue weights. In the real world, there are strong statistical regularities in the environment as well as high autocorrelation amongst information, which tend to boost performance of simple heuristics (Hutchinson & Gigerenzer, 2005). In addition, in some circumstances, acquiring new information can be costly both time- and effort-wise, which also fosters good-enough decision making. Although several studies have examined choice behavior under such scenarios (e.g., Dieckmann & Rieskamp, 2007; Rieskamp & Hoffrage, 2008), the majority of them have only tested the applicability of a handful of predefined strategy models (i.e., maximization versus Take-the-Best). To fully understand satisficing behavior, it is necessary to examine information usage through a systematic manipulation of the decision environment, based on an extensive set of possible decision models.

Expanding the decision model space to incorporate memory-based heuristics as well as mixture models (e.g., memory + partial cue integration) can also bring new

insights to understanding satisficing. These models will be especially useful in characterizing decision processes for solving familiar, well-practiced problems. Additionally, improving analytical methods to track dynamic changes in strategies as a function of time can provide crucial evidence for answering the question of how people learn to use heuristics. That is, in Chapter 5, we observed changes in strategies within the experimental phase through set-by-set analysis (Appendix B, Figure 17), which provided us with hints that participants adopt more stringent and exploratory heuristics before finding the best available strategy. Hence, modelling choice behavior on a continuous timescale could potentially produce rich evidence describing how heuristics evolve through learning and experience.

Finally, continuing investigations uncovering neural correlates of satisficing will help advance our understanding of human decision behavior in general, through detection of commonalities and differences across decision tasks, and across healthy and clinical populations. Specifically, questions such as how satisficing is triggered and achieved (e.g., is there a brain region or system that initiates or controls satisficing?), and how multiple cue information is processed in the brain (e.g., where does the cue information get dropped?; how are cues processed – serial vs. parallel?) remain largely unanswered. In Chapter 5, we provided some initial neural evidence based on fMRI data, which points to a critical role of subcortical dopaminergic regions in mediating satisficing. Further studies using various techniques to measure and/or manipulate brain

activity will shine light on more fine-grained neural mechanisms of satisficing decision making.

6.5 Conclusion

In this dissertation, I propose and develop a novel approach to examine satisficing decision making, focusing on shifts in strategies triggered by various external and internal pressures. Although satisficing decision making underlies most of the choices we make in real-life, research in cognitive neuroscience has rarely focused on the topic and therefore, there remain a great number of open questions in understanding how satisficing is learned, triggered, and achieved. Here, through a series of empirical studies, I report that people satisfice via adaptively employing the Drop-the-Worst heuristic, provided with a compensatory environment constructed using multiple probabilistic cues. Extending the task to incorporate additional aspects of large-world decision scenarios, combined with innovative analysis techniques and neuroimaging methods will allow for further exploration and conceptualization of cognitive and neuro-computational mechanisms of satisficing decision making.

Appendix A

Post-experimental survey

1. Please rank the cues from the most informative (1) to the least (4).

Use C = color, S = shape, T = contour, O = line orientation.

For example, if you thought color cues (blue/red) were the most reliable indicator for winning a given trial and/or if you've made decisions mostly based on color cues, color would be the most informative cue.

(1) _____ (2) _____ (3) _____ (4) _____

Questions 2 and 3 used in experiments included in Chapters 3:

2. In the slow phase, how many cues, on average, did you take into account before making decisions?

- 1 cue 2 cues 3 cues all cues

3. In the fast phase, how many cues, on average, did you take into account before making decisions?

- 1 cue 2 cues 3 cues all cues

Questions 2 and 3 used in the experiment included in Chapter 4:

2. In the first (neutral) phase, how many cues, on average, did you take into account before making decisions?


- 1 cue 2 cues 3 cues all cues

3. In the second (shock) phase, how many cues, on average, did you take into account before making decisions?

- 1 cue 2 cues 3 cues all cues

4. Please provide your best guess to answer the questions below.

i) Color



- Which stimulus has a higher probability of winning?
 left right
- Based on your choice above, how likely do you think you will win this trial?
 50% 60% 70% 80% 90% 100%

ii) Shape



- Which stimulus has a higher probability of winning?
 left right
- Based on your choice above, how likely do you think you will win this trial?
 50% 60% 70% 80% 90% 100%

iii) Contour



- Which stimulus has a higher probability of winning?
 left right
- Based on your choice above, how likely do you think you will win this trial?
 50% 60% 70% 80% 90% 100%

iv) Orientation



- Which stimulus has a higher probability of winning?
 left right
- Based on your choice above, how likely do you think you will win this trial?
 50% 60% 70% 80% 90% 100%

Appendix B

Supplementary information for Chapter 5

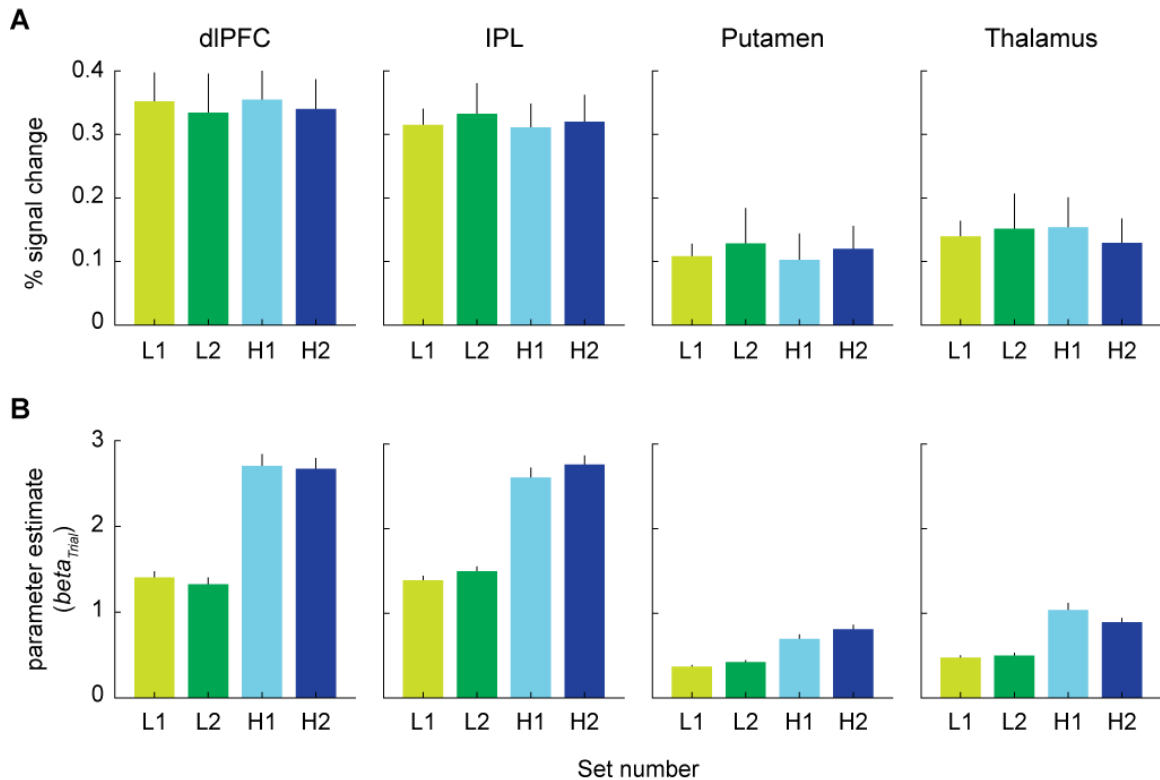


Figure 16: Changes in Trial-Induced BOLD Activity over Time in the dIPFC, IPL, Putamen, and Thalamus/Midbrain ROIs. (A) Average percent signal changes. (B) Average parameter estimates, β_{Trial} . A full statistical report is included in Table 8. Numbers on x-axis corresponds to the set number within the low (L) and high (H) time pressure phases. Error bars indicate SEM.

Table 8: Results from 2 (LowP/HighP Phase) x 2 (First/Second Set) Repeated-Measures ANOVAs in Selected Regions of Interest.

| Region | Phase | Set | Phase × Set |
|---|---------|------|-------------|
| Average percent signal change | | | |
| L/R. Medial frontal gyrus (vmPFC) | 1.44 | 0.15 | 0.01 |
| L/R. Medial frontal gyrus (SMA) | 0.05 | 1.47 | 0.39 |
| L/R. Middle frontal gyrus (dlPFC) | 0.01 | 0.27 | 0.004 |
| L/R. Inferior parietal lobule (IPL) | 0.12 | 0.23 | 0.02 |
| L. Putamen | 0.07 | 0.39 | 0.002 |
| R. Thalamus/midbrain | 0.01 | 0.03 | 0.51 |
| L. Cerebellum | 0.11 | 0.40 | 0.16 |
| Average parameter estimate, β_{Trial} | | | |
| L/R. Medial frontal gyrus (vmPFC) | 3.36 | 0.22 | 0.06 |
| L/R. Medial frontal gyrus (SMA) | 15.66* | 0.96 | 0.002 |
| L/R. Middle frontal gyrus (dlPFC) | 24.02** | 0.07 | 0.02 |
| L/R. Inferior parietal lobule (IPL) | 26.44** | 0.45 | 0.01 |
| L. Putamen | 4.59* | 0.42 | 0.05 |
| R. Thalamus/midbrain | 5.97* | 0.12 | 0.46 |
| L. Cerebellum | 23.23** | 0.65 | 0.48 |

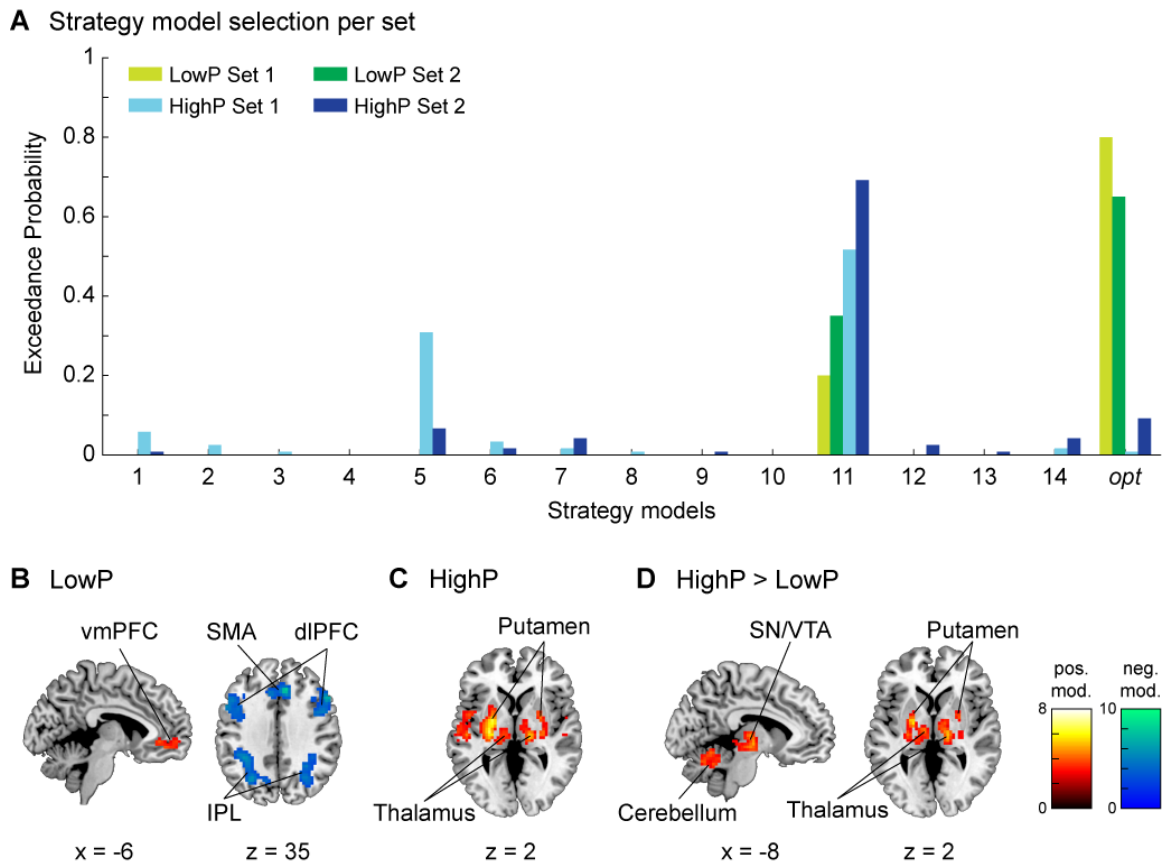


Figure 17: Strategy Model Selection and fMRI Data Analyses per Set. (A) Exceedance probabilities of each strategy model per set. **(B)** Brain regions significantly modulated by subjective SoE under the low time pressure (LowP). **(C)** Brain regions significantly modulated by subjective SoE under the high time pressure (HighP). **(D)** Areas demonstrating significantly greater positive SoE modulation in the HighP compared to the LowP phase. All neuroimaging analyses were performed using only the first set of each phase. All maps were whole-brain corrected to $p < 0.05$.

Table 9: Activation Clusters for Subjective SoE Modulation Using the First Set of Each Phase.

| Region | Peak MNI (<i>x, y, z</i>) | Peak <i>t</i> | Cluster size (voxels) |
|---|--------------------------------|------------------|--------------------------|
| Low time pressure, positive modulation | | | |
| L/R. Medial frontal gyrus/anterior cingulate | (0, 56, -1) | 4.34 | 156 |
| Low time pressure, negative modulation | | | |
| L/R. Medial frontal gyrus | (9, 26, 38) | 8.47 | 380 |
| L. Middle/inferior frontal gyrus | (-45, 29, 23) | 7.78 | 478 |
| R. Middle/inferior frontal gyrus | (51, 11, 26) | 7.13 | 619 |
| L. Inferior parietal lobule | (-45, -40, 41) | 6.55 | 474 |
| R. Inferior parietal lobule | (27, -61, 38) | 5.25 | 367 |
| L. Fusiform gyrus | (-45, -52, -13) | 4.96 | 90 |
| High time pressure, positive modulation | | | |
| L. Putamen/thalamus/parahippocampal gyrus/insula/amygdala/midbrain | (-30, -7, -4) | 7.44 | 1111 |
| R. Putamen/insula/parahippocampal gyrus/amygdala | (24, 11, -10) | 5.79 | 325 |
| R. Cerebellum | (18, -52, -25) | 5.46 | 129 |
| L. Cerebellum/lingual gyrus | (-18, -61, -28) | 5.41 | 336 |
| L/R. Cingulate gyrus | (0, -13, 44) | 4.91 | 109 |
| R. Postcentral gyrus | (42, -19, 53) | 4.77 | 105 |
| R. Superior temporal gyrus | (51, -19, 8) | 4.41 | 131 |
| High time pressure, negative modulation | | | |
| L. Precuneus/inferior parietal lobule | (-30, -64, 41) | 5.33 | 137 |
| High time pressure > Low time pressure, positive modulation | | | |
| L/R. Putamen/thalamus/parahippocampal gyrus/midbrain | (-21, -7, 2) | 5.70 | 738 |
| L/R. Cerebellum | (18, -52, -25) | 5.36 | 425 |
| R. Precentral/postcentral gyrus | (33, -19, 53) | 4.91 | 137 |

References

- Akaishi, R., Kolling, N., Brown, J. W., & Rushworth, M. (2016). Neural Mechanisms of Credit Assignment in a Multitue Environment. *Journal of Neuroscience*, 36(4), 1096–1112. <http://doi.org/10.1523/JNEUROSCI.3159-15.2016>
- Alemanno, A. (2012). Nudging Smokers - The Behavioural Turn of Tobacco Risk Regulation. *European Journal of Risk Regulation*, 3(1), 32–42.
- Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école américaine [Rational behavior under risk: Criticism of the postulates and axioms of the American school]. *Econometrica*, 21(4), 503–546.
- Ashby, F. G., Turner, B. O., & Horvitz, J. C. (2010). Cortical and basal ganglia contributions to habit learning and automaticity. *Trends in Cognitive Sciences*, 14(5), 208–215. <http://doi.org/10.1016/j.tics.2010.02.001>
- Balci, F., Simen, P., Niyogi, R., Saxe, A., Hughes, J. A., Holmes, P., & Cohen, J. D. (2011). Acquisition of decision making criteria: reward rate ultimately beats accuracy. *Attention, Perception, & Psychophysics*, 73(2), 640–657. <http://doi.org/10.3758/s13414-010-0049-7>
- Ballew, C. C., & Todorov, A. (2007). Predicting political elections from rapid and unreflective face judgments. *Proceedings of the National Academy of Sciences of the United States of America*, 104(46), 17948–53. <http://doi.org/10.1073/pnas.0705435104>
- Bar-Haim, Y., Lamy, D., Pergamin, L., Bakermans-Kranenburg, M. J., & van IJzendoorn, M. H. (2007). Threat-related attentional bias in anxious and nonanxious individuals: A meta-analytic study. *Psychological Bulletin*, 133(1), 1–24. <http://doi.org/10.1037/0033-2909.133.1.1>
- Barlow, H. B. (1961). Possible Principles Underlying the Transformations of Sensory Messages. In W. A. Rosenblith (Ed.), *Sensory Communication* (pp. 216–234). Cambridge, MA: MIT Press. <http://doi.org/10.7551/mitpress/9780262518420.003.0013>
- Bartra, O., McGuire, J. T., & Kable, J. W. (2013). The valuation system: A coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. *NeuroImage*, 76, 412–427. <http://doi.org/10.1016/j.neuroimage.2013.02.063>
- Basten, U., Biele, G., Heekeren, H. R., & Fiebach, C. J. (2010). How the brain integrates

- costs and benefits during decision making. *Proceedings of the National Academy of Sciences of the United States of America*, 107(50), 21767–21772.
<http://doi.org/10.1073/pnas.0908104107>
- Bergert, F. B., & Nosofsky, R. M. (2007). A response-time approach to comparing generalized rational and take-the-best models of decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33(1), 107–29.
<http://doi.org/10.1037/0278-7393.33.1.107>
- Bernoulli, D. (1954). Exposition of a new theory on the measurement of risk. *Econometrica*, 22(1), 23–36.
- Böckenholt, U., & Kroeger, K. (1993). *The effect of time pressure in multiattribute binary choice tasks.* (O. Svenson & A. J. Maule, Eds.) *Time pressure and stress in human judgment and decision making.* Boston, MA: Springer US. <http://doi.org/10.1007/978-1-4757-6846-6>
- Bogacz, R. (2014). Speed-Accuracy Trade-Off. In D. Jaeger & R. Jung (Eds.), *Encyclopedia of Computational Neuroscience* (pp. 1–4). New York, NY: Springer New York.
<http://doi.org/10.1007/978-1-4614-7320-6>
- Bogacz, R., Hu, P. T., Holmes, P. J., & Cohen, J. D. (2010). Do humans produce the speed-accuracy trade-off that maximizes reward rate? *Quarterly Journal of Experimental Psychology* (2006), 63(5), 863–891. <http://doi.org/10.1080/17470210903091643>
- Bogacz, R., Wagenmakers, E. J., Forstmann, B. U., & Nieuwenhuis, S. (2010). The neural basis of the speed-accuracy tradeoff. *Trends in Neurosciences*, 33(1), 10–16.
<http://doi.org/10.1016/j.tins.2009.09.002>
- 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), 1332–1346. <http://doi.org/10.1037//0278-7393.26.5.1332>
- Bröder, A. (2003). Decision making with the “adaptive toolbox”: Influence of environmental structure, intelligence, and working memory load. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29(4), 611–625.
<http://doi.org/10.1037/0278-7393.29.4.611>
- Bromberg-Martin, E. S., Matsumoto, M., & Hikosaka, O. (2010). Dopamine in Motivational Control: Rewarding, Aversive, and Alerting. *Neuron*, 68(5), 815–834.
<http://doi.org/10.1016/j.neuron.2010.11.022>

- Buckner, R. L. (2013). The cerebellum and cognitive function: 25 years of insight from anatomy and neuroimaging. *Neuron*, 80(3), 807–815.
<http://doi.org/10.1016/j.neuron.2013.10.044>
- Cain, M. S., Dunsmoor, J. E., LaBar, K. S., & Mitroff, S. R. (2011). Anticipatory anxiety hinders detection of a second target in dual-target search. *Psychological Science : A Journal of the American Psychological Society / APS*, 22, 866–871.
<http://doi.org/10.1177/0956797611412393>
- Chevalier, G., & Deniau, J. M. (1990). Disinhibition as a basic process in the expression of striatal functions. *Trends in Neurosciences*, 13(7), 277–280.
- Clark, L., Li, R., Wright, C. M., Rome, F., Fairchild, G., Dunn, B. D., & Aitken, M. R. F. (2012). Risk-avoidant decision making increased by threat of electric shock. *Psychophysiology*, 49(10), 1436–1443. <http://doi.org/10.1111/j.1469-8986.2012.01454.x>
- Clithero, J. a, & Rangel, A. (2013). Informatic parcellation of the network involved in the computation of subjective value. *Social Cognitive and Affective Neuroscience*, 9(9), 1289. <http://doi.org/10.1093/scan/nst106>
- Cohen, J. (1992). A Power Primer. *Quantitative Methods in Psychology*, 112(1), 155–159.
<http://doi.org/10.1037/0033-2909.112.1.155>
- Cornwell, B. R., Baas, J. M. P., Johnson, L., Holroyd, T., Carver, F. W., Lissek, S., & Grillon, C. (2007). Neural responses to auditory stimulus deviance under threat of electric shock revealed by spatially-filtered magnetoencephalography. *NeuroImage*, 37(1), 282–289. <http://doi.org/10.1016/j.neuroimage.2007.04.055>
- Curtis, C. E., & D’Esposito, M. (2003). Persistent activity in the prefrontal cortex during working memory. *Trends in Cognitive Sciences*, 7(9), 415–423.
[http://doi.org/10.1016/S1364-6613\(03\)00197-9](http://doi.org/10.1016/S1364-6613(03)00197-9)
- Czerlinski, J., Gigerenzer, G., & Goldstein, D. G. (1999). How good are simple heuristics? In G. Gigerenzer & P. M. Todd (Eds.), *Simple heuristics that make us smart* (pp. 97–118). New York, NY: Oxford University Press.
- D’Ardenne, K., McClure, S. M., Nystrom, Leigh, E., & Cohen, J. D. (2008). BOLD responses reflecting dopaminergic signals in the human ventral tegmental area. *Science*, 319, 1264–1267. <http://doi.org/10.1126/science.1229223>
- Daw, N. D., & Doya, K. (2006). The computational neurobiology of learning and reward. *Current Opinion in Neurobiology*, 16(2), 199–204.

<http://doi.org/10.1016/j.conb.2006.03.006>

Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., & Dolan, R. J. (2011). Model-based influences on humans' choices and striatal prediction errors. *Neuron*, 69(6), 1204–15. <http://doi.org/10.1016/j.neuron.2011.02.027>

Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, 441(June), 876–879. <http://doi.org/10.1038/nature04766>

Delgado, M. R., Nystrom, L. E., Fissell, C., Noll, D. C., & Fiez, J. A. (2000). Tracking the hemodynamic responses to reward and punishment in the striatum. *Journal of Neurophysiology*, 84(6), 3072–3077. [http://doi.org/10.1016/0166-2236\(90\)90107-1](http://doi.org/10.1016/0166-2236(90)90107-1)

Dias-Ferreira, E., Sousa, J. C., Melo, I., Morgado, P., Mesquita, A. R., Cerqueira, J. J., ... Sousa, N. (2009). Chronic stress causes frontostriatal reorganization and affects decision-making. *Science*, 325, 621–625.

Dieckmann, A., & Rieskamp, J. (2007). The influence of information redundancy on probabilistic inferences. *Memory & Cognition*, 35(7), 1801–1813.

Dosenbach, N. U. F., Visscher, K. M., Palmer, E. D., Miezin, F. M., Wenger, K. K., Kang, H. C., ... Petersen, S. E. (2006). A Core System for the Implementation of Task Sets. *Neuron*, 50(5), 799–812. <http://doi.org/10.1016/j.neuron.2006.04.031>

Doyon, J., Bellec, P., Amsel, R., Penhune, V., Monchi, O., Carrier, J., ... Benali, H. (2009). Contributions of the basal ganglia and functionally related brain structures to motor learning. *Behavioural Brain Research*, 199(1), 61–75. <http://doi.org/10.1016/j.bbr.2008.11.012>

Drugowitsch, J. (2013). Variational Bayesian inference for linear and logistic regression. *arXiv:1310.5438*, 1(3), 1–13.

Drugowitsch, J., DeAngelis, G., Klier, E., Angelaki, D., & Pouget, A. (2014). Optimal multisensory decision-making in a reaction-time task. *eLife*, 3, 1–19. <http://doi.org/10.7554/eLife.03005>

Duncan, J., & Owen, A. M. (2000). Common regions of the human frontal lobe recruited by diverse cognitive demands. *Trends in Neurosciences*, 23(10), 475–483. [http://doi.org/10.1016/S0166-2236\(00\)01633-7](http://doi.org/10.1016/S0166-2236(00)01633-7)

Dunsmoor, J. E., Mitroff, S. R., & LaBar, K. S. (2009). Generalization of conditioned fear

- along a dimension of increasing fear intensity. *Learning & Memory*, 16(7), 460–9. <http://doi.org/10.1101/lm.1431609>
- Easterbrook, J. A. (1959). The effect of emotion on cue utilization and the organization of behavior. *Psychological Review*, 66(3), 183–201. <http://doi.org/10.1037/h0047707>
- Edwards, W. (1954). The theory of decision making. *Psychological Bulletin*, 51(4), 380–417.
- Eysenck, M. W., Derakshan, N., Santos, R., & Calvo, M. G. (2007). Anxiety and cognitive performance: Attentional control theory. *Emotion*, 7(2), 336–353. <http://doi.org/10.1037/1528-3542.7.2.336>
- Ferrari, S., & Cai, C. (2009). Information-driven search strategies in the board game of CLUE. *IEEE Transactions on Systems, Man, and Cybernetics. Part B, Cybernetics : A Publication of the IEEE Systems, Man, and Cybernetics Society*, 39(3), 607–25. <http://doi.org/10.1109/TSMCB.2008.2007629>
- Fiorillo, C. D., Newsome, W. T., & Schultz, W. (2008). The temporal precision of reward prediction in dopamine neurons. *Nature Neuroscience*, 11(8), 966–973. <http://doi.org/10.1038/nn.2159>
- Fiorillo, C. D., Tobler, P. N., & Schultz, W. (2003). Discrete Coding of Reward Dopamine Neurons. *Science*, 299, 1898–1902.
- Fischer, J., & Whitney, D. (2014). Serial dependence in visual perception. *Nature Neuroscience*, 17(5), 738–743. <http://doi.org/10.1038/nn.3689>
- Fishburn, P. C. (1974). Lexicographic Orders, Utilities and Decision Rules: A Survey. *Management Science*, 20(11), 1442–1471. <http://doi.org/10.1287/mnsc.20.11.1442>
- Foerde, K., Knowlton, B. J., & Poldrack, R. A. (2006). Modulation of competing memory systems by distraction. *Proceedings of the National Academy of Sciences of the United States of America*, 103(31), 11778–83. <http://doi.org/10.1073/pnas.0602659103>
- Forstmann, B. U., Dutilh, G., Brown, S., Neumann, J., von Cramon, D. Y., Ridderinkhof, K. R., & Wagenmakers, E.-J. (2008). Striatum and pre-SMA facilitate decision-making under time pressure. *Proceedings of the National Academy of Sciences of the United States of America*, 105(45), 17538–42. <http://doi.org/10.1073/pnas.0805903105>
- Friedman, M. (1953). *Essays in positive economics. The Philosophy of Economics: An Anthology*. University of Chicago Press.

- Friston, K., Holmes, A., Worsley, K., Polin, J., Frith, C., & Frackowiak, R. (1994). Statistical parametric maps in functional imaging: a general linear approach. *Human Brain Mapping* 2, 2, 189–210.
- Friston, K., Litvak, V., Oswal, A., Razi, A., Stephan, K., van Wijk, B., ... Zeidman, P. (2015). Bayesian model reduction and empirical Bayes for group (DCM) studies. *NeuroImage*. <http://doi.org/10.1016/j.neuroimage.2015.11.015>
- Garcia-Retamero, R., & Dhami, M. K. (2009). Take-the-best in expert-novice decision strategies for residential burglary. *Psychonomic Bulletin & Review*, 16(1), 163–169. <http://doi.org/10.3758/PBR.16.1.163>
- Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology*, 62, 451–82. <http://doi.org/10.1146/annurev-psych-120709-145346>
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: models of bounded rationality. *Psychological Review*, 103(4), 650–69.
- Gigerenzer, G., & Todd, P. M. (1999). Fast and frugal heuristics: the adaptive toolbox. In *Simple Heuristics That Make Us Smart* (pp. 3–34). Oxford: Oxford University Press. <http://doi.org/10.1177/1354067X0171006>
- Glöckner, A., & Betsch, T. (2008). Multiple-reason decision making based on automatic processing. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, 34(5), 1055–1075. <http://doi.org/10.1037/0278-7393.34.5.1055>
- Gluck, M., & Bower, G. (1988). From conditioning to category learning: an adaptive network model. *Journal of Experimental Psychology. General*, 117(3), 227–47.
- Gluck, M., Shohamy, D., & Myers, C. (2002). How do people solve the “weather prediction” task?: individual variability in strategies for probabilistic category learning. *Learning & Memory*, 9, 408–418. <http://doi.org/10.1101/lm.45202>
- Gluth, S., Rieskamp, J., & Büchel, C. (2014). Neural evidence for adaptive strategy selection in value-based decision-making. *Cerebral Cortex*, 24, 2009–2021. <http://doi.org/10.1093/cercor/bht049>
- Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, 30, 535–574. <http://doi.org/10.1146/annurev.neuro.29.051605.113038>
- Goldstein, D. G., & Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristic. *Psychological Review*, 109(1), 75–90.

<http://doi.org/10.1037//0033-295X.109.1.75>

- Gould, I. C., Nobre, A. C., Wyart, V., & Rushworth, M. F. S. (2012). Effects of decision variables and intraparietal stimulation on sensorimotor oscillatory activity in the human brain. *The Journal of Neuroscience*, *32*(40), 13805–18. <http://doi.org/10.1523/JNEUROSCI.2200-12.2012>
- Green, S. R., Kragel, P. A., Fecteau, M. E., & LaBar, K. S. (2014). Development and validation of an unsupervised scoring system (Autonomate) for skin conductance response analysis. *International Journal of Psychophysiology*, *91*(3), 186–193. <http://doi.org/10.1016/j.ijpsycho.2013.10.015>
- Grillon, C., & Ameli, R. (1998). Effects of threat of shock, shock electrode placement and darkness on startle. *International Journal of Psychophysiology*, *28*(3), 223–231. [http://doi.org/10.1016/S0167-8760\(97\)00072-X](http://doi.org/10.1016/S0167-8760(97)00072-X)
- Grillon, C., Morgan, C. A., Davis, M., & Southwick, S. M. (1998). Effect of darkness on acoustic startle in Vietnam veterans with PTSD. *American Journal of Psychiatry*, *155*(6), 812–817. <http://doi.org/10.1176/ajp.155.6.812>
- Grueschow, M., Polania, R., Hare, T. A., & Ruff, C. C. (2015). Automatic versus Choice-Dependent Value Representations in the Human Brain. *Neuron*, *85*(4), 874–885. <http://doi.org/10.1016/j.neuron.2014.12.054>
- Grune-Yanoff, T., & Hertwig, R. (2016). Nudge Versus Boost: How Coherent are Policy and Theory? *Minds and Machines*, *26*, 149–183. <http://doi.org/10.1007/s11023-015-9367-9>
- Grupe, D. W., & Nitschke, J. B. (2013). Uncertainty and anticipation in anxiety: an integrated neurobiological and psychological perspective. *Nature Reviews*, *14*(July), 488–501. <http://doi.org/10.1038/nrn3524>
- Guthrie, J., Mancino, L., & Lin, C.-T. J. (2015). Nudging Consumers toward Better Food Choices: Policy Approaches to Changing Food Consumption Behaviors. *Psychology & Marketing*, *32*(5), 501–511. <http://doi.org/10.1002/mar>
- Habas, C., Kamdar, N., Nguyen, D., Prater, K., Beckmann, C. F., Menon, V., & Greicius, M. D. (2009). Distinct cerebellar contributions to intrinsic connectivity networks. *Journal of Neuroscience*, *29*(26), 8586–8594. <http://doi.org/10.1523/JNEUROSCI.1868-09.2009>
- Haber, S. N., & Knutson, B. (2009). The Reward Circuit: Linking Primate Anatomy and

- Human Imaging. *Neuropsychopharmacology*, 35(10), 4–26.
<http://doi.org/10.1038/npp.2009.129>
- Hanks, A. S., Just, D. R., Smith, L. E., & Wansink, B. (2012). Healthy convenience: Nudging students toward healthier choices in the lunchroom. *Journal of Public Health (United Kingdom)*, 34(3), 370–376. <http://doi.org/10.1093/pubmed/fds003>
- Hare, T. A., Malmaud, J., & Rangel, A. (2011). Focusing attention on the health aspects of foods changes value signals in vmPFC and improves dietary choice. *Journal of Neuroscience*, 31(30), 11077–87. <http://doi.org/10.1523/JNEUROSCI.6383-10.2011>
- Haynes, J. D., Sakai, K., Rees, G., Gilbert, S., Frith, C., & Passingham, R. E. (2007). Reading Hidden Intentions in the Human Brain. *Current Biology*, 17(4), 323–328. <http://doi.org/10.1016/j.cub.2006.11.072>
- Heekeren, H. R., Marrett, S., & Ungerleider, L. G. (2008). The neural systems that mediate human perceptual decision making. *Nature Reviews Neuroscience*, 9, 467–479. <http://doi.org/10.1038/nrn2374>
- Hikosaka, O. (1989). Role of Basal Ganglia in Initiation of Voluntary Movements. In *Dynamic interactions in neural networks: Models and data* (pp. 153–167). Springer New York.
- Ho, T., Brown, S., van Maanen, L., Forstmann, B. U., Wagenmakers, E.-J., & Serences, J. T. (2012). The optimality of sensory processing during the speed-accuracy tradeoff. *The Journal of Neuroscience*, 32(23), 7992–8003. <http://doi.org/10.1523/jneurosci.0340-12.2012>
- Hoyer, W. D., & Brown, S. P. (1990). Effects of Brand Awareness on Choice for a Common, Repeat-Purchase Product. *Journal of Consumer Research*, 17(2), 141. <http://doi.org/10.1086/208544>
- Hu, K., Bauer, A., Padmala, S., & Pessoa, L. (2012). Threat of bodily harm has opposing effects on cognition. *Emotion*, 12(1), 28–32. <http://doi.org/10.1037/a0024345>
- Hutchinson, J. M. C., & Gigerenzer, G. (2005). Simple heuristics and rules of thumb: where psychologists and behavioural biologists might meet. *Behavioural Processes*, 69(2), 97–124. <http://doi.org/10.1016/j.beproc.2005.02.019>
- Ito, M. (2008). Control of mental activities by internal models in the cerebellum. *Nature Reviews Neuroscience*, 9(Box 1), 304–313. <http://doi.org/10.1038/nrn2332>

- Itō, M. (1984). *The cerebellum and neural control*. Raven Pr.
- Jacoby, L. L., & Brooks, L. R. (1984). Nonanalytic Cognition: Memory, Perception, and Concept Learning. *Psychology of Learning and Motivation*, 18(C), 1–47.
[http://doi.org/10.1016/S0079-7421\(08\)60358-8](http://doi.org/10.1016/S0079-7421(08)60358-8)
- Johnson, E., & Payne, J. (1986). The Decision to Commit a Crime: An Information-Processing Analysis. In D. B. Cornish & R. V. Clarke (Eds.), *The Reasoning Criminal: Rational Choice Perspectives on Offending* (pp. 170–185). New York: Springer.
- Juslin, P., Olsson, H., & Olsson, A.-C. (2003). Exemplar effects in categorization and multiple-cue judgment. *Journal of Experimental Psychology. General*, 132(1), 133–156.
<http://doi.org/10.1037/0096-3445.132.1.133>
- Juslin, P., & Persson, M. (2002). PROBABILITIES from EXemplars (PROBEX): A “lazy” algorithm for probabilistic inference from generic knowledge. *Cognitive Science*, 26(5), 563–607. [http://doi.org/10.1016/S0364-0213\(02\)00083-6](http://doi.org/10.1016/S0364-0213(02)00083-6)
- Kable, J. W., & Glimcher, P. W. (2007). The neural correlates of subjective value during intertemporal choice. *Nature Neuroscience*, 10(12), 1625–1633.
<http://doi.org/10.1038/nn2007>
- Kable, J. W., & Glimcher, P. W. (2009). The neurobiology of decision: consensus and controversy. *Neuron*, 63(6), 733–45. <http://doi.org/10.1016/j.neuron.2009.09.003>
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Economic Society*, 49(2), 263–292.
- Kahnt, T., Heinzle, J., Park, S. Q., & Haynes, J.-D. (2011). Decoding different roles for vmPFC and dlPFC in multi-attribute decision making. *NeuroImage*, 56(2), 709–715.
<http://doi.org/10.1016/j.neuroimage.2010.05.058>
- Kass, R., & Raftery, A. (1995). Bayes factors. *Journal of the American Statistical Association*, 90(430), 773–795.
- Keinan, G. (1987). Decision making under stress: Scanning of alternatives under controllable and uncontrollable threats. *Journal of Personality and Social Psychology*, 52(3), 639–644. <http://doi.org/10.1037/0022-3514.52.3.639>
- Kelley, H., & Friedman, D. (2002). Learning To Forecast Price. *Economic Inquiry*, 40(4), 556–573. <http://doi.org/10.1093/ei/40.4.556>

- Kelly, A. M. C., & Garavan, H. (2005). Human functional neuroimaging of brain changes associated with practice. *Cerebral Cortex*, *15*(8), 1089–1102.
<http://doi.org/10.1093/cercor/bhi005>
- Khader, P. H., Pachur, T., Meier, S., Bien, S., Jost, K., & Rösler, F. (2011). Memory-based decision-making with heuristics: evidence for a controlled activation of memory representations. *Journal of Cognitive Neuroscience*, *23*(11), 3540–54.
http://doi.org/10.1162/jocn_a_00059
- Khader, P. H., Pachur, T., Weber, L. A., & Jost, K. (2015). Neural Signatures of Controlled and Automatic Retrieval Processes in Memory-based Decision-making. *Journal of Cognitive Neuroscience*, *28*(1), 69–83. <http://doi.org/10.1162/jocn>
- Kira, S., Yang, T., & Shadlen, M. N. (2015). A Neural Implementation of Wald's Sequential Probability Ratio Test. *Neuron*, *85*, 861–873.
<http://doi.org/10.1016/j.neuron.2015.01.007>
- Kizilbash, A. H., Vanderploeg, R. D., & Curtiss, G. (2002). The effects of depression and anxiety on memory performance. *Archives of Clinical Neuropsychology : The Official Journal of the National Academy of Neuropsychologists*, *17*(1), 57–67.
[http://doi.org/10.1016/S0887-6177\(00\)00101-3](http://doi.org/10.1016/S0887-6177(00)00101-3)
- Knight, F. H. (1921). *Risk, uncertainty and profit*. New York: Sentry Press.
- Knowlton, B., Mangels, J., & Squire, L. (1996). A neostriatal habit learning system in humans. *Science*, *273*(5280), 1399–402.
- Knowlton, B., Squire, L., & Gluck, M. (1994). Probabilistic classification learning in amnesia. *Learning & Memory*, *1*, 106–120. <http://doi.org/10.1101/lm.1.2.106>
- Knutson, B., Adams, C. M., Fong, G. W., & Hommer, D. (2001). Anticipation of increasing monetary reward selectively recruits nucleus accumbens. *The Journal of Neuroscience*, *21*, RC159. <http://doi.org/20015472> [pii]
- Krebs, R. M., Boehler, C. N., Roberts, K. C., Song, A. W., & Woldorff, M. G. (2012). The involvement of the dopaminergic midbrain and cortico-striatal-thalamic circuits in the integration of reward prospect and attentional task demands. *Cerebral Cortex*, *22*(3), 607–15. <http://doi.org/10.1093/cercor/bhr134>
- Kriegeskorte, N., Goebel, R., & Bandettini, P. (2006). Information-based functional brain mapping. *Proceedings of the National Academy of Sciences of the United States of America*, *103*, 3863–3868. <http://doi.org/10.1073/pnas.0600244103>

- Kruschke, J. K. (1992). ALCOVE: an exemplar-based connectionist model of category learning. *Psychological Review*, 99(1), 22–44. <http://doi.org/10.1037/0033-295X.99.1.22>
- Lagnado, D., Newell, B., Kahan, S., & Shanks, D. (2006). Insight and strategy in multiple-cue learning. *Journal of Experimental Psychology: General*, 135(2), 162–183. <http://doi.org/10.1037/0096-3445.135.2.162>
- Lamberts, K. (1995). Categorization under time pressure. *Journal of Experimental Psychology: General*, 124(2), 161–180. <http://doi.org/10.1037/0096-3445.124.2.161>
- Lamberts, K. (2000). Information-accumulation theory of speeded categorization. *Psychological Review*, 107(2), 227–260. <http://doi.org/10.1037//0033-295X.107.2.227>
- Lee, M. D., & Cummins, T. D. R. (2004). Evidence accumulation in decision making: unifying the “take the best” and the “rational” models. *Psychonomic Bulletin & Review*, 11(2), 343–352. <http://doi.org/10.3758/BF03196581>
- Lehéricy, S., Benali, H., Van de Moortele, P.-F., Péligrini-Issac, M., Waechter, T., Ugurbil, K., & Doyon, J. (2005). Distinct basal ganglia territories are engaged in early and advanced motor sequence learning. *Proceedings of the National Academy of Sciences of the United States of America*, 102(35), 12566–12571. <http://doi.org/10.1073/pnas.0502762102>
- Levy, D. J., & Glimcher, P. W. (2012). The root of all value: a neural common currency for choice. *Current Opinion in Neurobiology*, 22(6), 1027–1038. <http://doi.org/10.1016/j.conb.2012.06.001>
- Levy, I., Snell, J., Nelson, A. J., Rustichini, A., & Glimcher, P. W. (2010). Neural representation of subjective value under risk and ambiguity. *Journal of Neurophysiology*, 103, 1036–1047. <http://doi.org/10.1152/jn.00853.2009>
- Lykken, D. T., & Venables, P. H. (1971). Direct measurement of skin conductance: A proposal for standardization. *Psychophysiology*, 8(5), 656–672. <http://doi.org/10.1111/j.1469-8986.1971.tb00501.x>
- Ma, W. J., & Jazayeri, M. (2014). Neural Coding of Uncertainty and Probability. *Annual Review of Neuroscience*, 37, 205–220. <http://doi.org/10.1146/annurev-neuro-071013-014017>
- Macdonald, E. K., & Sharp, B. M. (2000). Brand Awareness Effects on Consumer Decision Making for a Common, Repeat Purchase Product: *Journal of Business Research*, 48(98), 5–15. [http://doi.org/10.1016/S0148-2963\(98\)00070-8](http://doi.org/10.1016/S0148-2963(98)00070-8)

- Maldjian, J. A., Laurienti, P. J., & Burdette, J. H. (2004). Precentral gyrus discrepancy in electronic versions of the Talairach atlas. *NeuroImage*, *21*(1), 450–455. <http://doi.org/10.1016/j.neuroimage.2003.09.032>
- Maldjian, J. A., Laurienti, P. J., Kraft, R. A., & Burdette, J. H. (2003). An automated method for neuroanatomic and cytoarchitectonic atlas-based interrogation of fMRI data sets. *NeuroImage*, *19*(3), 1233–1239. [http://doi.org/10.1016/S1053-8119\(03\)00169-1](http://doi.org/10.1016/S1053-8119(03)00169-1)
- Marewski, J. N., Gaissmaier, W., Schooler, L. J., Goldstein, D. G., & Gigerenzer, G. (2010). From recognition to decisions: extending and testing recognition-based models for multialternative inference. *Psychonomic Bulletin & Review*, *17*(3), 287–309. <http://doi.org/10.3758/PBR.17.3.287>
- Martignon, L., Katsikopoulos, K. V., & Woike, J. K. (2008). Categorization with limited resources: A family of simple heuristics. *Journal of Mathematical Psychology*, *52*(6), 352–361. <http://doi.org/10.1016/j.jmp.2008.04.003>
- Martin, M., Williams, R. M., & Clark, D. M. (1991). Does anxiety lead to selective processing of threat-related information? *Behaviour Research and Therapy*, *29*(2), 147–160. [http://doi.org/10.1016/0005-7967\(91\)90043-3](http://doi.org/10.1016/0005-7967(91)90043-3)
- Mata, R., Schooler, L. J., & Rieskamp, J. (2007). The Aging Decision Maker: Cognitive Aging and the Adaptive Selection of Decision Strategies. *Psychology and Aging*, *22*(4), 796–810. <http://doi.org/10.1093/acprof:oso/9780199744282.003.0022>
- Maule, A. J. (1994). A componential investigation of the relation between structural modelling and cognitive accounts of human judgement. *Acta Psychologica*, *87*(2–3), 199–216. [http://doi.org/10.1016/0001-6918\(94\)90051-5](http://doi.org/10.1016/0001-6918(94)90051-5)
- McClure, S. M., Berns, G. S., & Montague, P. R. (2003). Temporal prediction errors in a passive learning task activate human striatum. *Neuron*, *38*, 339–346.
- Meeter, M., Myers, C., Shohamy, D., Hopkins, R., & Gluck, M. (2006). Strategies in probabilistic categorization: results from a new way of analyzing performance. *Learning & Memory*, *13*, 230–239. <http://doi.org/10.1101/lm.43006>
- Meeter, M., Radics, G., Myers, C., Gluck, M., & Hopkins, R. (2008). Probabilistic categorization: How do normal participants and amnesic patients do it? *Neuroscience & Biobehavioral Reviews*, *32*(2), 237–248. <http://doi.org/10.1016/j.neubiorev.2007.11.001>

- Michael, E., de Gardelle, V., Nevado-Holgado, A., & Summerfield, C. (2015). Unreliable Evidence: 2 Sources of Uncertainty During Perceptual Choice. *Cerebral Cortex*, 25(4), 937–947. <http://doi.org/10.1093/cercor/bht287>
- Mink, J. W. (1996). The basal ganglia: Focused selection and inhibition of competing motor programs. *Progress in Neurobiology*, 50(4), 381–425. [http://doi.org/10.1016/S0301-0082\(96\)00042-1](http://doi.org/10.1016/S0301-0082(96)00042-1)
- Montague, P. R., & Berns, G. S. (2002). Neural economics and the biological substrates of valuation. *Neuron*, 36, 265–284. [http://doi.org/10.1016/S0896-6273\(02\)00974-1](http://doi.org/10.1016/S0896-6273(02)00974-1)
- Montague, P. R., Dayan, P., & Sejnowski, T. J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *The Journal of Neuroscience*, 16(5), 1936–1947. <http://doi.org/10.11156.635>
- Morris, G., Arkadir, D., Nevet, A., Vaadia, E., & Bergman, H. (2004). Coincident but distinct messages of midbrain dopamine and striatal tonically active neurons. *Neuron*, 43, 133–143. <http://doi.org/10.1016/j.neuron.2004.06.012>
- Murty, V. P., Shermohammed, M., Smith, D. V, Carter, R. M., Huettel, S. A., & Adcock, R. A. (2014). Resting state networks distinguish human ventral tegmental area from substantia nigra. *NeuroImage*, 100, 580–589. <http://doi.org/10.1016/j.neuroimage.2014.06.047>
- Newell, B. R., Lagnado, D. A., & Shanks, D. R. (2007). Challenging the role of implicit processes in probabilistic category learning. *Psychonomic Bulletin & Review*, 14(3), 505–511. <http://doi.org/10.3758/BF03194098>
- Newell, B. R., & Shanks, D. R. (2003). Take the best or look at the rest? Factors influencing “one-reason” decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29, 53–65. <http://doi.org/10.1037/0278-7393.29.1.53>
- Newell, B. R., & Shanks, D. R. (2004). On the role of recognition in decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30(4), 923–35. <http://doi.org/10.1037/0278-7393.30.4.923>
- Newell, B. R., Weston, N. J., & Shanks, D. R. (2003). Empirical tests of a fast-and-frugal heuristic: Not everyone “takes-the-best.” *Organizational Behavior and Human Decision Processes*, 91(1), 82–96. [http://doi.org/10.1016/S0749-5978\(02\)00525-3](http://doi.org/10.1016/S0749-5978(02)00525-3)
- Nichols, T., Brett, M., Andersson, J., Wager, T., & Poline, J. B. (2005). Valid conjunction inference with the minimum statistic. *NeuroImage*, 25(3), 653–660.

<http://doi.org/10.1016/j.neuroimage.2004.12.005>

- Niendam, T. A., Laird, A. R., Ray, K. L., Dean, Y. M., Glahn, D. C., & Carter, C. S. (2012). Meta-analytic evidence for a superordinate cognitive control network subserving diverse executive functions. *Cogn Affect Behav Neurosci*, *12*(2), 241–268. <http://doi.org/10.3758/s13415-011-0083-5>
- Niv, Y., Daniel, R., Geana, a., Gershman, S. J., Leong, Y. C., Radulescu, a., & Wilson, R. C. (2015). Reinforcement Learning in Multidimensional Environments Relies on Attention Mechanisms. *Journal of Neuroscience*, *35*(21), 8145–8157. <http://doi.org/10.1523/JNEUROSCI.2978-14.2015>
- Nosofsky, R. M., & Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review*, *104*(2), 266–300. <http://doi.org/10.1037/0033-295X.104.2.266>
- O’Doherty, J. P., Dayan, P., Friston, K., Critchley, H., & Dolan, R. J. (2003). Temporal difference models and reward-related learning in the human brain. *Neuron*, *28*, 329–337. [http://doi.org/10.1016/S0896-6273\(03\)00169-7](http://doi.org/10.1016/S0896-6273(03)00169-7)
- O’Doherty, J. P., Kringelbach, M. L., Rolls, E. T., Hornak, J., & Andrews, C. (2001). Abstract reward and punishment representations in the human orbitofrontal cortex. *Nature Neuroscience*, *4*(1), 95–102. <http://doi.org/10.1038/82959>
- O’Doherty, J. P., Lee, S. W., & Mcnamee, D. (2015). The structure of reinforcement-learning mechanisms in the human brain. *Current Opinion in Behavioral Sciences*, *1*, 94–100. <http://doi.org/10.1016/j.cobeha.2014.10.004>
- Oh, H., Beck, J. M., Zhu, P., Sommer, M. A., Ferrari, S., & Egner, T. (2016). Satisficing in split-second decision making is characterized by strategic cue discounting. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *42*(12), 1937–1956. <http://doi.org/10.1017/CBO9781107415324.004>
- Otto, A. R., Raio, C. M., Chiang, A., Phelps, E. A., & Daw, N. D. (2013). Working-memory capacity protects model-based learning from stress. *Proceedings of the National Academy of Sciences of the United States of America*, *110*(52), 20941–20946. <http://doi.org/10.1073/pnas.1312011110>
- Pachur, T., & Hertwig, R. (2006). On the psychology of the recognition heuristic: retrieval primacy as a key determinant of its use. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*(5), 983–1002. <http://doi.org/10.1037/0278-7393.32.5.983>

- Pachur, T., Hertwig, R., & Wolkewitz, R. (2014). The affect gap in risky choice: Affect-rich outcomes attenuate attention to probability information. *Decision, 1*(1), 64–78. <http://doi.org/10.1037/dec0000006>
- Pachur, T., & Marinello, G. (2013). Expert intuitions: How to model the decision strategies of airport customs officers? *Acta Psychologica, 144*(1), 97–103. <http://doi.org/10.1016/j.actpsy.2013.05.003>
- Pachur, T., & Olsson, H. (2012). Type of learning task impacts performance and strategy selection in decision making. *Cognitive Psychology, 65*(2), 207–240. <http://doi.org/10.1016/j.cogpsych.2012.03.003>
- Pachur, T., Todd, P. M., Gigerenzer, G., Schooler, L. J., & Goldstein, D. G. (2011). The recognition heuristic: a review of theory and tests. *Frontiers in Psychology, 2*(July), 147. <http://doi.org/10.3389/fpsyg.2011.00147>
- Packard, M. G., & Knowlton, B. J. (2002). Learning and Memory Functions of the Basal Ganglia. *Annual Review of Neuroscience, 25*, 563–93. <http://doi.org/10.1146/annurev.immunol.22.012703.104702>
- Padoa-Schioppa, C., & Assad, J. A. (2006). Neurons in the orbitofrontal cortex encode economic value. *Nature, 441*, 223–226. <http://doi.org/10.1038/nature04676>
- Payne, J., Bettman, J., & Johnson, E. (1988). Adaptive strategy selection in decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14*(3), 534–552. <http://doi.org/10.1037//0278-7393.14.3.534>
- Payne, J., Bettman, J., & Luce, M. (1996). When time is money: Decision behavior under opportunity-cost time pressure. *Organizational Behavior and Human Decision Processes, 66*(2), 131–152.
- Persky, J. (1995). Retrospectives : The Ethology of Homo Economicus. *The Journal of Economic Perspectives, 9*(2), 221–231.
- Petry, N. M., Stinson, F. S., & Grant, B. F. (2005). Comorbidity of DSM-IV Pathological Gambling and Other Psychiatric Disorders: Results From the National Epidemiologic Survey on Alcohol and Related Conditions. *The Journal of Clinical Psychiatry, 66*(5), 564–574. <http://doi.org/http://dx.doi.org/10.4088/JCP.v66n0504>
- Petzold, A., Plessow, F., Goschke, T., & Kirschbaum, C. (2010). Stress reduces use of negative feedback in a feedback-based learning task. *Behavioral Neuroscience, 124*(2), 248–255. <http://doi.org/10.1037/a0018930>

- Phelps, E. a, Lempert, K. M., & Sokol-Hessner, P. (2014). Emotion and Decision Making: Multiple Modulatory Neural Circuits. *Annual Review of Neuroscience*, 263–290. <http://doi.org/10.1146/annurev-neuro-071013-014119>
- Plassmann, H., O'Doherty, J., & Rangel, A. (2007). Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, 27(37), 9984–9988. <http://doi.org/10.1523/JNEUROSCI.2131-07.2007>
- Platt, M. L., & Huettel, S. A. (2008). Risky business: the neuroeconomics of decision making under uncertainty. *Nature Neuroscience*, 11(4), 398–403. <http://doi.org/10.1038/nn2062>
- Pohl, R. (2006). Empirical tests of the recognition heuristic. *Journal of Behavioral Decision Making*, 271(June), 251–271.
- Poldrack, R., Clark, J., Paré-Blagoev, E., Shohamy, D., Creso Moyano, J., Myers, C., & Gluck, M. (2001). Interactive memory systems in the human brain. *Nature*, 414, 546–550. <http://doi.org/10.1038/35107080>
- Poldrack, R., & Rodriguez, P. (2004). How do memory systems interact? Evidence from human classification learning. *Neurobiology of Learning and Memory*, 82(3), 324–32. <http://doi.org/10.1016/j.nlm.2004.05.003>
- Raguraman, a. P., & Padoa-Schioppa, C. (2014). Integration of Multiple Determinants in the Neuronal Computation of Economic Values. *Journal of Neuroscience*, 34(35), 11583–11603. <http://doi.org/10.1523/JNEUROSCI.1235-14.2014>
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2), 59–108. <http://doi.org/10.1037/0033-295X.85.2.59>
- Rhudy, J. L., & Meagher, M. W. (2000). Fear and anxiety: divergent effects on human pain thresholds. *Pain*, 84(1), 65–75. [http://doi.org/10.1016/S0304-3959\(99\)00183-9](http://doi.org/10.1016/S0304-3959(99)00183-9)
- Rieskamp, J. (2006). Perspectives of probabilistic inferences: Reinforcement learning and an adaptive network compared. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32(6), 1355–1370. <http://doi.org/10.1037/0278-7393.32.6.1355>
- Rieskamp, J., & Hoffrage, U. (2008). Inferences under time pressure: how opportunity costs affect strategy selection. *Acta Psychologica*, 127(2), 258–76. <http://doi.org/10.1016/j.actpsy.2007.05.004>

- Rieskamp, J., & Otto, P. E. (2006). SSL: a theory of how people learn to select strategies. *Journal of Experimental Psychology. General*, 135(2), 207–36.
<http://doi.org/10.1037/0096-3445.135.2.207>
- Rigoux, L., Stephan, K. E., Friston, K. J., & Daunizeau, J. (2014). Bayesian model selection for group studies - revisited. *NeuroImage*, 84, 971–85.
<http://doi.org/10.1016/j.neuroimage.2013.08.065>
- Robinson, O. J., Overstreet, C., Charney, D. R., Vytal, K., & Grillon, C. (2013). Stress increases aversive prediction error signal in the ventral striatum. *Proceedings of the National Academy of Sciences of the United States of America*, 110(10), 4129–33.
<http://doi.org/10.1073/pnas.1213923110>
- Robinson, O. J., Vytal, K., Cornwell, B. R., & Grillon, C. (2013). The impact of anxiety upon cognition: perspectives from human threat of shock studies. *Frontiers in Human Neuroscience*, 7(203), 1–21. <http://doi.org/10.3389/fnhum.2013.00203>
- Rosburg, T., Mecklinger, A., & Frings, C. (2011). When the brain decides: a familiarity-based approach to the recognition heuristic as evidenced by event-related brain potentials. *Psychological Science*, 22(12), 1527–34.
<http://doi.org/10.1177/0956797611417454>
- Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, 16(2), 225–237. <http://doi.org/10.3758/PBR.16.2.225>
- Sailer, U., Robinson, S., Fischmeister, F. P. S., König, D., Oppenauer, C., Lueger-Schuster, B., ... Bauer, H. (2008). Altered reward processing in the nucleus accumbens and mesial prefrontal cortex of patients with posttraumatic stress disorder. *Neuropsychologia*, 46(11), 2836–2844.
<http://doi.org/10.1016/j.neuropsychologia.2008.05.022>
- Savage, L. J. (1954). *The foundations of statistics*. New York, NY: Wiley.
- Schmitz, A., & Grillon, C. (2012). Assessing fear and anxiety in humans using the threat of predictable and unpredictable aversive events (the NPU-threat test). *Nature Protocols*, 7(3), 527–32. <http://doi.org/10.1038/nprot.2012.001>
- Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review*, 84(1), 1–66.
<http://doi.org/10.1037/0033-295X.84.1.1>

- Schooler, L. J., & Hertwig, R. (2005). How forgetting aids heuristic inference. *Psychological Review*, *112*(3), 610–28. <http://doi.org/10.1037/0033-295X.112.3.610>
- Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, *80*, 1–27. <http://doi.org/10.1007/s00429-010-0262-0>
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*(5306), 1593–1599. <http://doi.org/10.1126/science.275.5306.1593>
- Schwabe, L., & Wolf, O. T. (2009). Stress Prompts Habit Behavior in Humans. *Journal of Neuroscience*, *29*(22), 7191–7198. <http://doi.org/10.1523/JNEUROSCI.0979-09.2009>
- Schwabe, L., & Wolf, O. T. (2012). Stress modulates the engagement of multiple memory systems in classification learning. *The Journal of Neuroscience*, *32*(32), 11042–11049. <http://doi.org/10.1523/JNEUROSCI.1484-12.2012>
- Schwabe, L., & Wolf, O. T. (2013). Stress and multiple memory systems: From “thinking” to “doing.” *Trends in Cognitive Sciences*, *17*(2), 68. <http://doi.org/10.1016/j.tics.2012.12.001>
- Shackman, A. J., Maxwell, J. S., McMenemy, B. W., Greischar, L. L., & Davidson, R. J. (2011). Stress Potentiates Early and Attenuates Late Stages of Visual Processing. *Journal of Neuroscience*, *31*(3), 1156–1161. <http://doi.org/10.1523/JNEUROSCI.3384-10.2011>
- Shackman, A. J., Sarinopoulos, I., Maxwell, J. S., Pizzagalli, D. A., Lavric, A., & Davidson, R. J. (2006). Anxiety selectively disrupts visuospatial working memory. *Emotion*, *6*(1), 40–61. <http://doi.org/10.1037/1528-3542.6.1.40>
- Shah, A. K., & Oppenheimer, D. M. (2008). Heuristics made easy: an effort-reduction framework. *Psychological Bulletin*, *134*(2), 207–22. <http://doi.org/10.1037/0033-2909.134.2.207>
- Shohamy, D., Myers, C. E., Kalanithi, J., & Gluck, M. A. (2008). Basal ganglia and dopamine contributions to probabilistic category learning. *Neuroscience and Biobehavioral Reviews*, *32*(2), 219–36. <http://doi.org/10.1016/j.neubiorev.2007.07.008>
- Shohamy, D., Myers, C. E., Onlaor, S., & Gluck, M. A. (2004). Role of the basal ganglia in category learning: how do patients with Parkinson’s disease learn? *Behavioral Neuroscience*, *118*(4), 676–686. <http://doi.org/10.1037/0735-7044.118.4.676>
- Shohamy, D., Myers, C., Grossman, S., Sage, J., Gluck, M., & Poldrack, R. (2004). Cortico-

- striatal contributions to feedback-based learning: converging data from neuroimaging and neuropsychology. *Brain*, 127, 851–859.
<http://doi.org/10.1093/brain/awh100>
- Simen, P., Contreras, D., Buck, C., Hu, P., Holmes, P., & Cohen, J. D. (2009). Reward rate optimization in two-alternative decision making: empirical tests of theoretical predictions. *Journal of Experimental Psychology. Human Perception and Performance*, 35(6), 1865–1897. <http://doi.org/10.1037/a0016926>
- Simon, H. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1), 99–118.
- Simon, H. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63(2), 129–138.
- Simon, H. (1990). Invariants of human behavior. *Annual Review of Psychology*, 41, 1–19.
- Simon, H., & Kadane, J. (1975). Optimal Problem-Solving. Search: All-or-None Solutions. *Artificial Intelligence*, 6(1975), 235–247.
- Smith, P. L., & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends in Neurosciences*, 27(3), 161–168. <http://doi.org/10.1016/j.tins.2004.01.006>
- Speekenbrink, M., Lagnado, D. A., Wilkinson, L., Jahanshahi, M., & Shanks, D. R. (2010). Models of probabilistic category learning in Parkinson's disease: Strategy use and the effects of L-dopa. *Journal of Mathematical Psychology*, 54(1), 123–136.
<http://doi.org/10.1016/j.jmp.2009.07.004>
- Spielberger, C. D. (1983). *Manual for the State-Trait Anxiety Inventory*. Palo Alto, CA: Consulting Psychologists Press.
- Squire, L. R., & Zola, S. M. (1996). Structure and function of declarative and nondeclarative memory systems. *Proceedings of the National Academy of Sciences of the United States of America*, 93(24), 13515–13522. <http://doi.org/10.1073/pnas.93.24.13515>
- Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J., & Friston, K. J. (2009). Bayesian model selection for group studies. *NeuroImage*, 46(4), 1004–17.
<http://doi.org/10.1016/j.neuroimage.2009.03.025>
- Strick, P. L., Dum, R. P., & Fiez, J. A. (2009). Cerebellum and nonmotor function. *Annu Rev Neurosci*, 32, 413–434. <http://doi.org/10.1146/annurev.neuro.31.060407.125606>

- Summerfield, C., & Tsetsos, K. (2015). Do humans make good decisions? *Trends in Cognitive Sciences*, 19(1), 27–34. <http://doi.org/10.1016/j.tics.2014.11.005>
- Svenson, O., & Maule, A. J. (Eds.). (1993). *Time Pressure and Stress in Human Judgment and Decision Making*. Boston, MA: Springer US. <http://doi.org/10.1007/978-1-4757-6846-6>
- Thaler, R. H., & Sunstein, C. R. (2009). *Nudge: Improving Decisions About Health, Wealth, and Happiness*. New York: Penguin Books.
- Thomas, L. a, & LaBar, K. S. (2008). Fear relevancy, strategy use, and probabilistic learning of cue-outcome associations. *Learning & Memory (Cold Spring Harbor, N.Y.)*, 15(10), 777–784. <http://doi.org/10.1101/lm.1048808>
- Tobler, P. N., Fiorillo, C. D., & Schultz, W. (2005). Adaptive coding of reward value by dopamine neurons. *Science*, 307(March), 1642–1645. <http://doi.org/10.1126/science.1105370>
- Tobler, P. N., O'Doherty, J. P., Dolan, R. J., & Schultz, W. (2007). Reward value coding distinct from risk attitude-related uncertainty coding in human reward systems. *Journal of Neurophysiology*, 97, 1621–1632. <http://doi.org/10.1152/jn.00745.2006>
- Tobler, P. N., & Weber, E. U. (2013). *Valuation for Risky and Uncertain Choices. Neuroeconomics: Decision Making and the Brain: Second Edition*. Elsevier Inc. <http://doi.org/10.1016/B978-0-12-416008-8.00009-7>
- Tversky, A., & Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science*, 185(4157), 1124–1131. <http://doi.org/10.1126/science.185.4157.1124>
- Van Turennout, M., Bielowicz, L., & Martin, A. (2003). Modulation of neural activity during object naming: Effects of time and practice. *Cerebral Cortex*, 13(4), 381–391. <http://doi.org/10.1093/cercor/13.4.381>
- van Veen, V., Krug, M. K., & Carter, C. S. (2008). The neural and computational basis of controlled speed-accuracy tradeoff during task performance. *Journal of Cognitive Neuroscience*, 20(11), 1952–1965. <http://doi.org/10.1162/jocn.2008.20146>
- Venkatraman, V., Payne, J. W., Bettman, J. R., Luce, M. F., & Huettel, S. a. (2009). Separate Neural Mechanisms Underlie Choices and Strategic Preferences in Risky Decision Making. *Neuron*, 62(4), 593–602. <http://doi.org/10.1016/j.neuron.2009.04.007>
- Volz, K. G., & Gigerenzer, G. (2012). Cognitive Processes in Decisions Under Risk are not the Same as in Decisions Under Uncertainty. *Frontiers in Neuroscience*, 6, 105.

<http://doi.org/10.3389/fnins.2012.00105>

- Volz, K. G., & Gigerenzer, G. (2014). The brain is not “as-if” -- Taking stock of the neuroscientific approach on decision making. In T. D. Papageorgiou, G. I. Christopoulos, & S. M. Smirnakis (Eds.), *Advanced Brain Neuroimaging Topics in Health and Disease - Methods and Applications* (pp. 573–603). InTech.
- Volz, K. G., Schooler, L. J., Schubotz, R. I., Raab, M., Gigerenzer, G., & von Cramon, D. Y. (2006). Why you think milan is larger than modena: neural correlates of the recognition heuristic. *Journal of Cognitive Neuroscience*, *18*(11), 1924–36. <http://doi.org/10.1162/jocn.2006.18.11.1924>
- Volz, K. G., Schooler, L. J., & von Cramon, D. Y. (2010). It just felt right: the neural correlates of the fluency heuristic. *Consciousness and Cognition*, *19*(3), 829–37. <http://doi.org/10.1016/j.concog.2010.05.014>
- Von Neumann, J., & Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton University Press. Princeton.
- Vulkan, N. (2000). An economist’s perspective on probability matching. *Journal of Economic Surveys*, *14*(1), 101–118. <http://doi.org/10.1111/1467-6419.00106>
- Vytal, K. E., Cornwell, B. R., Letkiewicz, A. M., Arkin, N. E., & Grillon, C. (2013). The complex interaction between anxiety and cognition: insight from spatial and verbal working memory. *Frontiers in Human Neuroscience*, *7*(March), 1–11. <http://doi.org/10.3389/fnhum.2013.00093>
- Wager, T. D., Jonides, J., & Reading, S. (2004). Neuroimaging studies of shifting attention: A meta-analysis. *NeuroImage*, *22*(4), 1679–1693. <http://doi.org/10.1016/j.neuroimage.2004.03.052>
- Wang, H., & Song, M. (2011). Ckmeans.1d.dp: Optimal k-means Clustering in One Dimension by Dynamic Programming. *The R Journal*, *3*, 29–33.
- Wheeler, M. E., Woo, S. G., Ansel, T., Tremel, J. J., Collier, A. I., Velanova, K., ... Yang, T. (2015). The strength of gradually accruing probabilistic evidence modulates brain activity during a categorical decision. *Journal of Cognitive Neuroscience*, *27*(4), 705–719. <http://doi.org/10.1162/jocn>
- Whittlesea, B. W. (1993). Illusions of familiarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*(6), 1235–1253. <http://doi.org/10.1037//0278-7393.19.6.1235>

- Wilson, T. D., & Schooler, J. W. (1991). Thinking too much: introspection can reduce the quality of preferences and decisions. *Journal of Personality and Social Psychology*, 60(2), 181–192. <http://doi.org/10.1037/0022-3514.60.2.181>
- Wübben, M., & Wangenheim, F. V. (2008). Instant customer base analysis: Managerial heuristics often “get it right.” *Journal of Marketing*, 72(3), 82–93. <http://doi.org/10.1509/jmkg.72.3.82>
- Yang, T., & Shadlen, M. N. (2007). Probabilistic reasoning by neurons. *Nature*, 447(7148), 1075–80. <http://doi.org/10.1038/nature05852>
- Young, C. B., Raz, G., Everaerd, D., Beckmann, C. F., Tendolkar, I., Hendler, T., ... Hermans, E. J. (2017). Dynamic Shifts in Large-Scale Brain Network Balance As a Function of Arousal. *Journal of Neuroscience*, 37(2), 281–290. <http://doi.org/10.1523/JNEUROSCI.1759-16.2017>

Biography

Hanna Oh (also known as Hanna Oh Descher) was born in Seoul, South Korea in March 1986. She received a bachelor's degree in interdisciplinary engineering with bioengineering concentration from The Cooper Union for the Advancement of Science and Art in 2009. She then pursued graduate studies in South Korea, and received a master's degree in brain and cognitive sciences from Seoul National University in 2012. She entered a doctoral program in cognitive neuroscience at Duke University in fall 2012, and received a master's degree in psychology and neuroscience from Duke University in 2016. Under the supervision of Dr. Tobias Egner, Hanna's doctoral projects focus on investigating the mechanisms underlying satisficing decision making.

PUBLICATIONS

- Oh-Descher, H.**, Tanaka, H., LaBar, K. S., Ferrari, S., Sommer, M. A., & Egner, T. (submitted). Anticipatory anxiety promotes satisficing during multi-cue probabilistic decision making.
- Oh-Descher, H.**, Beck, J. M., Ferrari, S., Sommer, M. A., & Egner, T. (submitted). Probabilistic inference under time pressure leads to a cortical-to-subcortical shift in decision evidence integration.
- Oh, H.**, Beck, J. M., Zhu, P., Sommer, M. A., Ferrari, S., & Egner, T. (2016). Satisficing in split-second decision making is characterized by strategic cue discounting. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 42(12), 1937-1956.
- Keenan, J. P., **Oh, H.**, & Amati, F. (2011). An overview of self-awareness and the brain. In J. Decety, & J. T. Cacioppo (Eds.), *The Oxford handbook of social neuroscience* (pp. 314-324). New York, NY: Oxford University Press.

Amati, F., **Oh, H.**, Kwan, V. S., Jordan, K., & Keenan, J. P. (2010). Overclaiming and the medial prefrontal cortex: A transcranial magnetic stimulation study. *Cognitive Neuroscience*, 1(4), 268-276.