Title
Perspectives on Cognitive Modeling of Adaptive Behavior

Author
Mistry, Percy K

Publication Date
2018

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA,
IRVINE

Perspectives on cognitive modeling of adaptive behavior

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Psychology

by

Percy K Mistry

Dissertation Committee:
Professor Michael D. Lee, Chair
Assistant Professor Jennifer S. Trueblood
Professor Mark Steyvers

2018

DEDICATION

To my lovely wife, Dilber, and my awesome kids, Tanaisha and Diana!!

TABLE OF CONTENTS

LIST OF FIGURES

Page

LIST OF TABLES

ACKNOWLEDGMENTS

The research in this document was made possible by all of the guidance and support I received from my advisor, committee members, friends, and colleagues at the University of California, Irvine. I would especially like to thank my advisors, Jennifer Trueblood and Michael Lee, for mentoring and supporting me through the graduate program.

CURRICULUM VITAE

Percy K Mistry

EDUCATION

Doctor of Philosophy in Psychology — **June 2018**
University of California, Irvine — *Irvine, CA*

Masters of Arts in Psychology — **December 2015**
University of California, Irvine — *Irvine, CA*

PGDCM (MBA) in Systems, Finance — **March 2003**
Indian Institute of Management Calcutta — *Kolkata, India*

Bachelor of Engineering in Electronics Engineering — **June 2001**
University of Mumbai — *Mumbai, India*

ABSTRACT OF THE DISSERTATION

Perspectives on cognitive modeling of adaptive behavior

By

Percy K Mistry

Doctor of Philosophy in Psychology

University of California, Irvine, 2018

Professor Michael D. Lee, Chair

Many experimental and statistical paradigms collect and analyze behavioral data under steady-state assumptions. Such paradigms may have low external validity, since real-world decisions are often made in situations when people are still learning and adapting, or their beliefs and preferences are in a state of flux, or where the decision environment is constantly changing. I focus on experimental and real-world paradigms that represent some form of adaptive behavior. Under such situations, factoring in structural adaptivity into cognitive modeling frameworks can improve their descriptive and predictive performance, and allow us to make superior inferences about the underlying latent cognitive processes. Towards this endeavor, first, a novel probabilistic framework is proposed for representation of heuristic strategies of information search and multi-attribute choice within a model of learning. Second, a novel adaptive reference point mechanism is introduced, and I show how this can be used in a variety of different tasks and applications, and be incorporated into cognitive frameworks to structurally capture the adaptive process. This mechanism provides superior predictive performance along with psychologically meaningful parameter inference in tasks ranging from signal detection, bandit problems, judgments about estimating true values, consumption behavior, probability tracking, and expectation formation. Third, I show how cognitive modeling can be applied to adaptive population behavior in the real world. Real world data is used to make inferences about the latent processes involved in aspects such as the changing nature of cycles of violence, and the impact of tax policies on changing consumption patterns. Incorporating cognitive

modeling frameworks and psychological insights to constrain econometric models yields possibly simpler models, but with directly interpretable parameter inferences. Fourth, cognitive modeling approaches are used to design choice architecture frames to create behavioral nudges within a risk allocation paradigm, where people change their behavior based on representational, rather than meaningful changes in their environment.

Chapter 1

Introduction

1.1 PART I: Adaptive strategy switching and learning

In the first part, I motivate a discussion on finding novel probabilistic methods to represent decision heuristics and strategies in multi-attribute choice decisions. In chapter 2, I provide a theoretical framework for developing such probabilistic representations of decision strategies. In chapter 3, I report experimental results from 2 multi-attribute decision tasks. In chapter 4, I implement a set of cognitive models based on traditional and probabilistic strategies, including incorporating a cognitive model of latent strategy selection and switching. I show that these novel probabilistic models along with latent learning provide better descriptive and predictive capabilities, and can be used to make improved inferences about the underlying cognitive processes involved in information search and aggregation. In chapter 5, I propose a measurement framework to evaluate the potential effectiveness of a consideration set of strategies a priori, based on the experimental design, but before looking at the data.

1.2 PART II: Adaptive reference point formation

In the second part, in chapter 6, I propose a novel adaptive reference point mechanism, and propose that this mechanism can be ubiquitously applied to many different types of tasks, and incorporated within many existing but varied cognitive modeling approaches. In chapter 7, I show the modeling results and inferences made by incorporating such mechanisms. I show how incorporating such a mechanism improves descriptive and predictive capabilities in a variety of experimental tasks including signal detection, price judgment, consumption behavior, probability tracking, and expectation formation tasks. I show how robust inferences about individual differences can be made using this as a focal point in various cognitive frameworks.

1.3 PART III: Adaptive behavior at a population level

In the third part, I focus on real world data that reflects decision making of populations, and show that incorporating cognitive modeling frameworks and psychological insights to constrain econometric models yields possibly simpler models, but with directly interpretable parameter inferences. This is especially true when the underlying phenomenon involves some form of adaptive behavior. I apply this approach in chapter 8, to modeling violent behavior during the Second Intifada, and show that this approach can make reasonably strong predictions about the near future based on the past history of violence. The model can characterize the latent build up of violence, in terms of retaliatory versus repetitive behavior, and in terms of long term internalization versus short term dynamics. It provides a separate characterization for each political period during the Second Intifada. In chapter 9, I propose a detailed theoretical framework to adapt sufficient statistics based approaches used in econometrics to model consumption behavior. I show how such analytically obtained sufficient statistics are only valid under certain cognitive assumptions, and propose that additional cognitive variables are required to reach the correct conclusions about how people re-

spond to aspects such as changes in the tax rate and government policies. I provide a very basic application of this modeling framework to alcohol consumption in the US as a proof of concept.

1.4 PART IV: Inducing adaptive behavior by manipulation of choice architecture

All of the above address adaptive behavior and cognitive modeling of such behavior under different situations. The focus is on highlighting that accounting for adaptive behavior within cognitive modeling frameworks makes for a better quality of inferences, and more powerful models in terms of their descriptive and predictive capabilities. In the last and fourth part, I highlight experimental work that attempts to induce a change in risky behavior of participants, by changing the choice representation, that is a form of a behavioral *nudge*. In chapter 10, I show that people's adaptation is not homogeneous, and that perceptions of their own risk attitudes influence the direction in which people might be nudged. I show that behavior under such risky choices is best explained by cognitive models that account for choice allocation in a segregated rather than aggregate manner.

Part I

Adaptive Strategy Switching

Chapter 2

Theoretical Considerations

2.1 Introduction to multi-attribute decision tasks

Real world decision making tasks span a lot of different domains, including individual (e.g. what financial investments should be bought or sold, and when), professional (e.g. what lab tests should be run for a medical diagnosis), and institutional (e.g. what tax policy changes should be made in a certain economic environment), amongst others. A large subset of such decisions tasks are characterized by the fact that they are often repeated over time, although the frequency of repetitions can vary significantly. The objective is often to select an option that is expected to optimize a particular criterion, or set of criteria (e.g. maximize returns from investments). Uncertainty about how different options are related to the criteria make the decisions non-trivial. There are however usually a large set of known information cues related to the options. These cues have causal or correlational relationships with the criteria, and hence represent varying degrees of predictive or diagnostic information. Acquiring this information may be associated with varying degrees of cost (time, effort, or money). Since the number of potential cues is usually quite large, most decisions rely on selective information processing. Whilst the value of cues is known, the nature of relation-

ship (strength of causal or correlational influence) between each cue and the criteria usually carries a high degree of uncertainty. Effective decision making often involves reducing this uncertainty by learning through experience and feedback. Finally, this cue-outcome relationship often changes over time. Such change may be gradual or abrupt, and may or may not preserve the order of strength of influence that various cues have, in terms of being able to predict the desired outcome.

An example is selecting between investing in two companies in order to maximize the criteria of returns on investment. There is always uncertainty over which investments are likely to perform well over a particular time frame. However there are multiple possible information cues, such as the previous sales growth, the outstanding debt, capital reinvestment, change in management, and so on, which can help people select between options. These cues are causally or otherwise related to the expected returns on investment, some more strongly than others. Over time, an individual may learn, for example, that previous sales growth does not really influence future returns, and learn to ignore this information. Aspects such as a change in management may require privately rather than publicly available knowledge, and would represent high-cost information cues. So the individual begins to rely on the remaining two cues to predict which investments to make. It is possible however, that the relative importance of these cues changes over time. For example, outstanding debt may have a much larger influence on outcomes in a year when the interest rates have increased significantly. To study how people select and rely on different sources of information in order to make such decisions, many of these decision making task properties are captured in multi-attribute decision making (MADM) tasks.

An MADM task involves selecting one or more information cues (the attributes), and assimilating information from such cues in order to make a decision. The cue attributes typically incorporate some information about the outcome of each option. The statistical relationship between cue attributes and choice outcomes may be probabilistic or deterministic, may assume different functional forms, and in an experimental task, may either be known to participants or need to be learned through experience. Performance in such a task is measured by the accuracy of the decision, but

also by how well the relationships between cues and outcome are learned, whether the used cue search and order patterns are effective, and how efficiently multiple sources of information are combined. Uncertainty is captured in different ways in such tasks. In probabilistic tasks, the relationship between cues and criterion is defined probabilistically, as a *cue validity*. As an example, a cue validity of 0.8 implies that the cue correctly predicts the outcome criterion about 80% of the time. In such tasks (Rieskamp and Otto (2006); Lee et al. (2014)), the validities are typically, but not always, provided to the participants and the uncertainty lies in the probabilistic nature of the relationships. There is a different class of tasks (Bröder and Schiffer (2006a); Bröder et al. (2013)), where the criterion value is determined as some weighted average of the different cues. The *cue weights* are analogous to the cue validities, but are deterministic, with some random noise added to the overall outcome. Apart from the noise, the uncertainty in these tasks arises from the fact that the cue weights are typically not communicated to participants, but need to be learned over time.

2.1.1 Basic notation

For a particular multi-attribute decision task, we define n_A as the number of attributes available, and n_O as the number of choice options. Thus, there are $n_A n_O$ unique cues in the choice task. The key observable measures for the t^{th} decision include which cues are searched (s_t), and which choice option is selected (y_t). Note that this is not a comprehensive list of observable measures, but the most typical. Observed behavior x_t is defined as $x_t = s_t \cap y_t$. Under a paradigm where each individual cue can either be searched or not, there are $N_s = 2^{n_A n_O}$ unique cue search patterns possible. Thus, a particular search pattern s_j on a particular trial t is identified as $s_{j,t} \in \{s_j : j \in [1 : N_s]\}$. Similarly, we have $y_{k,t} \in \{y_k : k \in [1 : n_O]\}$.

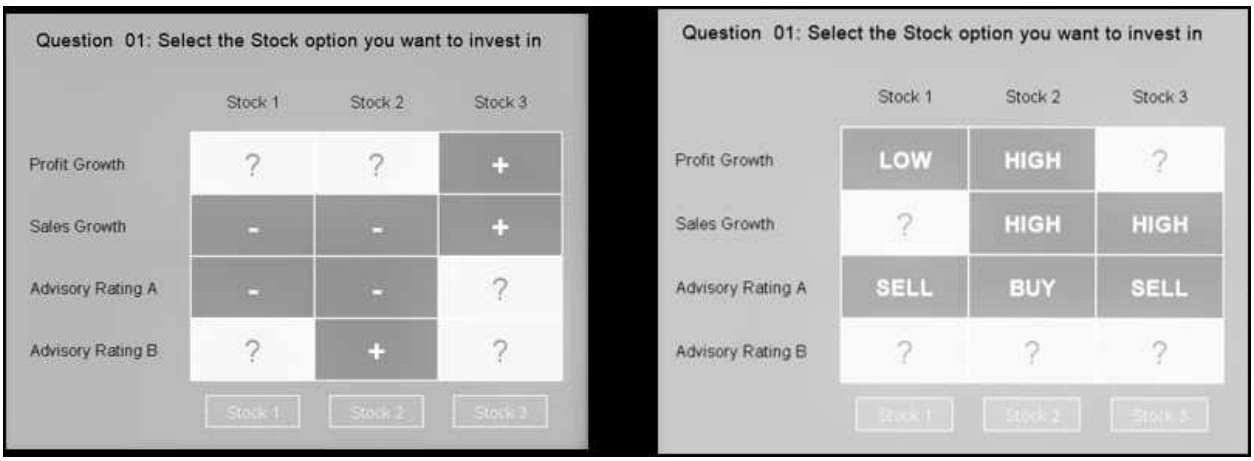

Figure 2.1: Example of 2 similar multi-attribute decision making tasks. The four rows correspond to information attributes and the three columns are the choice options. Each attribute by option cue can be individually acquired for a cost. At the end of the trial, feedback is provided on the rewards or penalties associated with the selected and foregone choices.

2.1.2 An example

Figure 2.1 provides an example of a multi-attribute decision task. Here, on each trial, there are 3 possible choice options (Stocks 1,2, or 3), so $n_O = 3$. There are 4 possible information cues (profit growth, sales growth, advisory ratings A and B), so $n_A = 4$. In this example, the information cues can take binary values (+/-, or High/Low or Buy/Sell). The statistical relationship between the different cues and the reward for each option is unknown to the participant, and need to be learned through experience. On any trial, there are $N_s = 2^{n_A n_O} = 4096$ possible search patterns, depending on which combination of the 12 cues has been selected.

2.1.3 Changing environmental contingencies

Some studies have manipulated the underlying cue-outcome relationship across trials or blocks (Bröder and Schiffer (2006a); Lee et al. (2014)), representing the change that may occur in naturalistic descision environments. Specifically, one way of incorporating this change is to vary the environments between compensatory and non-compensatory structures. These represent the nature

of statistical relationships between cues and the outcome criterion. A perfectly compensatory environment has all cue weights (or validities) equal. This of course, is too stringent and unrealistic. More broadly defined, in compensatory environments, most of the cue weights or validities are in a similar range, so that different cues can offset, or compensate for, each other. Thus, optimal strategies in such an environment should take into account as many cues as possible and aggregate them in an appropriate manner. A non-compensatory environment is one where cues are sequentially dominant. Here, the cue weight (or validity) of the most dominant cue is so high, that all the remaining cues put together cannot effectively offset the information from the dominant cue. In case the most dominant cue does not discriminate between options, the next best cue is similarly dominant compared to all remaining cues, and so on. Optimal strategies in such environments would involve identifying the most dominant discriminating cues and relying solely on those, in order to save any information acquisition costs associated with collecting more information.

2.2 Heuristic strategies

A large body of research focuses on both descriptive and prescriptive accounts of behavior in multi-attribute decision tasks (and their real world counterparts) using fast and frugal heuristics (Gigerenzer and Todd (1999)). Gigerenzer and Gaissmaier (2011) defined heuristics as *"a strategy that ignores part of the information, with the goal of making decisions more quickly, frugally, and/or accurately than more complex methods."*. They highlighted three key building blocks towards a theoretical framework for how heuristics are constructed:

1. Search rules: How do people explore the search space?
2. Stopping rules: When do people stop searching?
3. Decision rules: How do people make a final decision?

The key aspects of this framework are that specific heuristics or strategies are particularly efficient and accurate under certain environmental conditions, and that people are usually quite good at identifying the most appropriate heuristics for a given environment. The most commonly analyzed heuristics involve compensatory heuristics such as weighted average (WA) and tallying, or non-compensatory ones such as take-the-best (TTB). In tallying, all informations cues in the consideration set are equally weighted, and the option that has the highest number of favorable cues is selected. In WA, the cues are weighted based on their relative importance, and the weighted average of all cues is used to select a particular option. In TTB, the most important cue is selected from the consideration set, and the option which is most favored by this single cue is selected. If the most important cue cannot distinguish between the options, the next most important cue is selected, and so on. Each of these strategies can be differentially beneficial under different environmental conditions, but importantly, also dictates different patterns of information search and assimilation. Under conditions when the environmental structure changes over time (non-stationary), this posits that people will eventually detect changes and adopt a new, more appropriate strategy over time. The mechanisms for change detection, learning and adaptation of different strategies have been examined in tasks with non-stationary environments (e.g. Gluth et al. (2013); Rieskamp (2008); Lee et al. (2014)). These, and similar studies have provided some insights into the mechanics of adaptive information search and decision making behavior. There is evidence that people do learn environmental contingencies, can detect changes and adaptively switch strategies, but not necessarily in an ideal manner. It may be useful to think of strategy switching in operant conditioning terms, that is, learning new strategies in appropriate environments, extinction of learned strategies on encountering a change in the environment, and spontaneous recovery or renewal of previously abandoned strategies when recognizing previously encountered environmental conditions. However, attempts to develop a robust cognitive model of adaptive strategy switching and learning have met with limited success. In the next few sections, we propose that this is partly due to the structure of the inference problem and the way heuristics are traditionally defined.

2.3 Challenges in cognitive modeling of strategy learning

Existing methods to formalize adaptive learning and strategy selection propose reinforcement-learning of strategies (Rieskamp and Otto (2006); Rieskamp (2008); Erev and Barron (2005)) or rational metareasoning (Lieder and Griffiths (2015)). In such approaches, an important step is to infer the generating process that leads to sequential learning and selection of strategies. The learning model is implemented by simultaneously making inferences about two different latent processes: (a) which heuristics people have used on a particular problem or trial, and (b) how people update their belief about the effectiveness of this, or any counterfactual heuristic in their consideration set. This effectiveness is based on some measure of accuracy, cost, time, effort, constraint satisfaction, or a combination of such measures. When feedback is available, some measure of accuracy is an especially important factor. Since a cognitive model will sequentially infer how strategies are being used and learned, poor quality of inferences in step (a) will sequentially propagate, resulting in magnified errors in later trials. Thus, a poor quality of inferences about the latent heuristic process being used will significantly impact the quality of inferences made about the learning process. Note that this problem does not arise in typical implementations of cognitive learning in other domains, where the locus of learning is almost exclusively attributed to *observed* behaviors, rather than *latent* processes. We illustrate this difference below. First, we define the use of a latent heuristic strategy as h_m, where $h_m \in \{h_i : i \in [1 : N_h]\}$, where N_h is the number of distinct strategies in the consideration set.

A typical simple learning model where *observable actions* are reinforced can be specified as equation 2.1, where $q_{hm,t}$ is the value accorded to action h_m after trial t, r_t is the obtained reward at time t, η is the sensitivity to reward value r_t, or the learning rate, and $I_{h_{m,t}}$ is an indicator that records whether we observed a subject implementing action h_m on trial t. Determination of $I_{h_{m,t}}$ in this case is trivial.

$$q_{h_{m,t}} = I_{h_{m,t-1}}\left((1-\eta)\, q_{h_{m,t-1}} + \eta\, r_{t-1}\right) + (1 - I_{h_{m,t-1}})\, q_{h_{m,t-1}} \tag{2.1}$$

There are of course models of counterfactual learning, but the key issue is recognizing the difference between actions that were implemented and not implemented, on any particular trial t. In learning models of heuristic strategies however, h_m is not an observable action that is reinforced, but a latent cognitive strategy. Inference about which latent strategy was used on a particular trial, and thus determination of $I_{h_{m,t}}$, is non-trivial, since there is often a many-to-many mapping between latent strategies and observable behaviors, which is manifested in terms of the varying information search patterns and choice selection.

2.3.1 Determination of $I_{h_{m,t}}$ where h_m is latent

The inference problem here is to infer which latent strategy h_m resulted in the observable behavior x_t, for each individual decision, or on each trial of the experiment. A brief survey of methods used to infer this in existing models of learning and strategy selection reveal the following commonly used approaches:

1. *Choice:* By simply checking if the final observed choices are compatible with the heuristic h_m, but ignoring the information search patterns (e.g. Rieskamp (2008)). This method is often the most common approach used in literature, but essentially ignores the information search patterns, which are a key component of how the heuristics are defined.
2. *Minimum:* Based on whether the choices are compatible, and whether all the necessary minimum required cues proposed by the heuristic have been searched, but allowing for the search of any extra cues that are not required by a heuristic (e.g. Rieskamp and Otto (2006)).
3. *Exact:* Based on whether the choices are compatible, and whether the observed cue search pattern *exactly matches* the search pattern proposed by a heuristic.
4. *Process:* By analyzing aspects such as response time or process tracing (e.g. Bergert and Nosofsky (2007)). This however requires making additional assumptions about the implementation of the heuristic and the distributions of process variables under different heuristics.

The approach used may infer either that:

1. None of the strategies in the consideration set were applied on a particular trial: $I_{h_{m,t}} = 0 \; \forall \; h_m$, which we denote as *no match*.
2. Multiple strategies in the consideration set satisfy the inference criteria: $I_{h_{m,t}} = 1$ for more than one strategy h_m, which we denote as *multiple match*.
3. Exactly one of the strategies in the consideration set satisfies the inference criteria: $I_{h_{m,t}} = 1$ for exactly one strategy h_m, and is 0 for all others, which we denote as a *unique match*.

In table 2.1 we show a brief meta-analysis of 11 experiments (2 of which are new experiments presented in detail in this thesis, and 6 of which are explicitly modeled, using a new approach to cognitive modeling introduced in this chapter) involving multi-attribute decision making, in table 2.1. For each experiment, we take the most commonly applied consideration set in literature - a set of two heuristics, take-the-best (TTB) and weighted average (WA) - based on the true cue validities used in the experiments. We use the *choice*, *minimum*, and *exact* inference methods described above, and report the percentage of trials in each experiment, on which these inference methods result in a no match, multiple match, or unique match situation. The table also shows the average percentage of cues acquired across all participants within the experiment. Across these 10 experiments, *unique* matches were found on between 8% to 50% using the choice only method, between 10% to 70% using the minimum cues methods, and 5% to 74% using the exact match method. Generally, experiments with higher cue use (as we go down the list in the table) have a higher level of unique identifiability.

Now revisiting equation 2.1, the probability of using a particular strategy h_m on trial t is assumed to be given by a softmax function, as in equation 2.2, where there are N_h possible strategies in the consideration set, and θ is a consistency parameter.

$$p(h_{m,t}) = \frac{e^{\theta \; q_{h_{m,t}}}}{\Sigma_{j=1:N_h}\left(e^{\theta \; q_{h_{j,t}}}\right)} \tag{2.2}$$

Table 2.1: Meta-analysis showing the proportion of trials on which inference about TTB and WA yielded either no match, a match for both heuristics, and a unique match. The analysis is conducted using three of the popular inference methods - choice, minimum cues, and exact cues. For each study and each method, the table shows the percentage of trials on which neither strategy could be identified (no match), use of strategies was not discernible (multiple match), and where strategies could be uniquely identified. Studies are ranked in the order of increasing cue use (first data column).

		(1) Choice			(2) Minimum			(3) Exact		
Experiment	%cues	No match	Multiple	*Unique*	No match	Multiple	*Unique*	No match	Multiple	*Unique*
Mistry and Trueblood (2015) - 2	31%	15%	42%	*44%*	90%	0%	*10%*	95%	0%	*5%*
Newell and Shanks (2003)	32%	18%	74%	*8%*	71%	8%	*20%*	83%	1%	*16%*
Walsh and Gluck (2016)	42%	11%	73%	*17%*	79%	3%	*18%*	86%	1%	*13%*
Bröder and Schiffer (2006b) - 1	42%	24%	52%	*24%*	63%	2%	*35%*	81%	0%	*19%*
Bröder and Schiffer (2006b) - 2	42%	24%	52%	*24%*	64%	2%	*34%*	81%	0%	*19%*
Mistry and Trueblood (2015) - 1	48%	24%	55%	*21%*	77%	7%	*16%*	88%	0%	*11%*
Lee et al. (2014) - 3	57%	4%	56%	*41%*	30%	4%	*66%*	70%	0%	*30%*
Lee et al. (2014) - 4	57%	4%	56%	*41%*	35%	4%	*61%*	76%	0%	*24%*
Lee et al. (2014) - 2	65%	1%	73%	*26%*	19%	11%	*70%*	69%	0%	*31%*
Lee et al. (2014) - 1	84%	2%	73%	*26%*	9%	37%	*54%*	41%	1%	*58%*
Rieskamp and Otto (2006)	98%	2%	48%	*50%*	3%	46%	*51%*	22%	4%	*74%*

Note that in the case of *no match*, the probabilities for the use of any particular strategy will remain the same as on the previous trial. In the case of a consideration set of two strategies with multiple match, the $q-$values for both strategies will increase by the same amount, resulting in a minuscule change in the probabilities of using each strategy, unless the existing q values are very small, in which case the model will begin to predict equal use of either strategy. In either case, the learning inferred by the model is extremely minimal, and the inference process is inefficient. *Essentially, the model infers any kind of meaningful learning by individuals only when a unique match is observed.* As seen in table 2.1, using existing heuristics and inference methods, this happens in only a small percentage of the trials, especially in the complex experimental designs, where a greater degree of learning is actually plausible. Although this is clearly visible in the reinforcement learning case demonstrated above, this issue is ubiquitous regardless of the learning model employed, and will remain a challenge for other learning models including Bayesian and instance-based learning.

2.3.2 Probabilistic identification of the latent locus of learning h_m: The need to redefine strategies

Rather than use the choice, minimum, or exact matching methods for identifying the latent locus of learning h_m, we propose using Bayesian inference. Here, in order to model an individual's belief updating process for trial t, we first calculate the posterior probability $p(h_{m,t-1}|x_{jk,t-1})$ that a strategy h_m was used on trial $t-1$, after having observed the behavioral outcome $x_{jk,t-1}$. We then replace $I_{h_{m,t-1}}$ with $p(h_{m,t-1}|x_{jk,t-1})$ in equation 2.1. In the simplest case, x_{jkt} is the choice selected, denoted as y_{kt}. More realistically, it also includes the information cue search patterns, denoted as s_{jt}. Thus $x_{jkt} = y_{kt} \cap s_{jt}$.

For ease of exposition, we drop the subscript t for the t^{th} observation. The posterior probability can be written as:

$$p(h_m|x_{jk}) = \frac{p(x_{jk}|h_m)p(h_m)}{p(x_{jk})} = \frac{p(y_k|s_j h_m)p(s_j|h_m)p(h_m)}{\Sigma_{i=1:N_h}\Big(p(y_k|s_j h_i)p(s_j|h_i)p(h_i)\Big)} \tag{2.3}$$

Here, $p(h_m)$ defines the prior probability of a strategy h_m being used, and is obtained from the underlying cognitive model, as in equation 2.2. Note that while this is a straightforward application of Bayes rule, most approaches that specify learning models do not take this into account. Effectively, whilst the models calculate the probability of an individual using a particular heuristic trial-by-trial, they do not apply this to the inference process, rather the models behave as if the prior probability of inferring which heuristic will be used is uniformly distributed on each trial, and dependent only on the information contained in the current trial.

It is important to observe here that $p(y_k|s_j h_i)$ is just a probabilistic representation of the *decision rule* of the heuristic h_i, given a particular set s_j of information cues searched. Similarly, $p(s_j|h_i)$ is a probabilistic representation of the *search and stop rules* of the heuristic h_i. However heuristic

strategies such as TTB, tallying, and WA are traditionally defined in a rule-based manner. This means that the probabilities corresponding to the information search rules will be 0 for all but one search combination in the possible search space of magnitude $2^{n_A n_O}$, if we insist on an exact search. On the other hand, if we adopt a minimum cues acquired approach for a lexicographic strategy such as TTB, which requires only a small proportion of the cues to be searched, this probability $p(s_j|h_i)$ will be a small but equal value for a large part of the possible search space, and zero for the remaining space. This will result in the same inference problem faced in section 2.3.1, where we land up inferring either no matches, or multiple matches with equal probability, for most trials.

The reason for this is because the rule-based heuristics do not have a well-defined error gradient defined. That is, there is no existing approach that defines how proximal a particular information search pattern is from the rule-based definition of the heuristic. A search pattern is classified in a binary manner, as either compatible or incompatible with the heuristic. Any search and decision pattern that is incompatible with any strategy within the consideration set is classified either as an application error, or a guessing strategy, with some uniform or normal error distribution across all incompatible search and decision patterns. Given that human decision making can be noisy yet structurally defined by the underlying heuristic strategies, we propose defining a more appropriate graded model of noise for the heuristic strategies. This is done by redefining the building blocks of heuristics proposed by Gigerenzer and Todd (1999) probabilistically within a multidimensional psychological space, that allows us to infer the proximity of different search and decision patterns, and define a structurally superior error model of behavior in the next section. To the best of our knowledge, this is a completely novel approach to defining such heuristic strategies. Essentially, we redefine distribution of information search patterns conditional on the use of a particular strategy, $p(s_j|h_i)$.

2.4 A generic framework for redefining heuristic strategies

2.4.1 Defining strategies probabilistically

Step 1: Representation of the search space

We propose redefining the information search and stop block of heuristics in an n-dimensional psychological feature space. Each individual search pattern can be represented as a point in this n-dimensional space. It is possible for more than one unique search pattern to lie on the same point in this reduced dimensional feature space. The features are extracted from the search patterns, and can be thought of as representing the cognitive primitives based on which people direct their information search. These features could be statistical properties of the search patterns (e.g. proportion of cues searched, sensitivity to validity, variability in cues searched across attributes, search density within selected attributes, etc.), process measures (e.g. time spent, search orders, across-options versus across-attribute search transitions), or other psychological constructs that define search behavior (e.g. confidence, effort, or contextual features). In any possible problem, we may consider using a subset of these (or other) features. Table 2.2 provides a representative list of features grouped into statistical, psychological, and context-based, however this is not an exhaustive list.

Table 2.2: Examples of defining features for cue search patterns

Search pattern based features	Psychological features	Contextual features
Proportion of cues searched	Expected effort	e.g. Profit
Proportion of most valid cues	Computation time	(as a cue to stock performance)
Variability across alternatives	Resulting confidence	e.g. Price
Variability across attributes	Cost sensitivity	(as a cue to sales)
Cue density within selected attributes		
Sensitivity to true cue validity		
Type 1 vs Type 2 transitions		

We denote the vector of n *selected* features as f, where $f : \{F_1, ...F_n\}$. Figure 2.2 shows the

hypothetical example of 3 patterns and a few selected statistical properties of these patterns. As an example, we take a rather simple case with two selected features, F_1 is the proportion of cues searched (ranges from 0 to 1), and F_2 is the sensitivity to true validity (measured as a weighted validity of all selected cues, normalized to range from 0 to 1). Figure 2.3 shows the 3 patterns of figure 2.2 in this 2-dimensional representation.

There are many possible choices in terms of what features to define in terms of the psychological representation of search patterns. While different feature choices may make sense in different applications, we propose that ideally, the selected features are proposed as cognitive primitives of information search, and should hence be psychologically interpretable. Secondly, the features should be relatively independent, to capture maximum information. Although the level of independence can be checked using a principal component analysis, we do not recommend using PCA to extract independent features which are not psychological interpretable, since the interpretation of prototypical heuristics in this space may get lost.

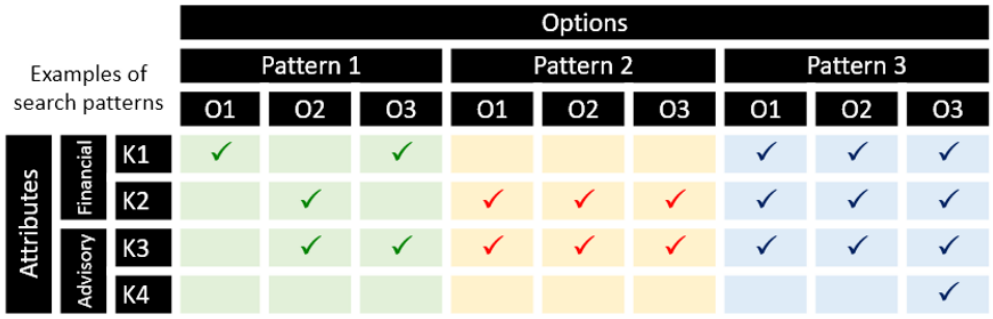

Examples of Features	Pattern 1	Pattern 2	Pattern 3
Proportion of cues searched	0.42	0.50	0.83
Sensitivity to Validity	0.67	0.50	0.60
Proportion of most valid cues searched	0.67	0.00	1.00
Variability across attributes	0.55	1.00	0.58
Proportion cues within searched attributes	0.56	1.00	0.83

Figure 2.2: Three hypothetical search patterns and some derived features for each pattern.

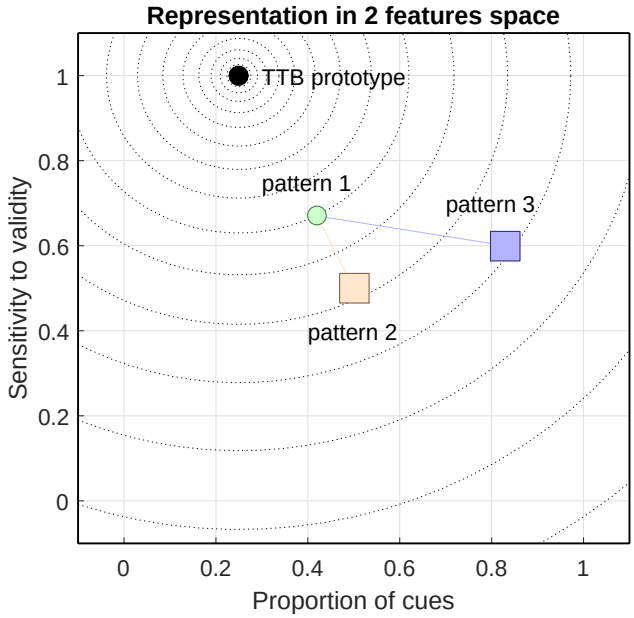

Figure 2.3: A hypothetical simple 2-dimensional representation of the 3 search patterns shown in figure 2.2, as well as a TTB prototype. The density of the radial lines around the TTB prototype represent the strength of this Gaussian TTB kernel at that point.

Step 2: Prototypical representation of strategies within this space

Behavior arising from the rule-based definition of the heuristic strategy to be considered, for instance, TTB, can be identified as a *prototypical point* in this n-dimensional space. This is straight forward for defining statistical properties, but may require some subjective estimates for aspects such as process measures or psychological constructs. For instance, the prototype for a particular strategy h_i is defined by $\bar{f}_i : \{F_1(s_i), ...F_n(s_i)\}$, where s_i is the search pattern obtained by strict application of the strategy rule. Here, $F_1(s_i)$ refers to the first feature based on the search pattern s_i, and so on. For a consideration set of multiple strategies, each strategy will have a different prototypical position in this feature space. Figure 2.3 shows the representation of the TTB prototype within the hypothetical example of the 2-dimensional representation, along with the 3 example patterns. This shows that *under this selected representation*, pattern 1 is closer to the TTB prototype, than the other two patterns. It is clear that the choice of features used to represent strategies via dimension reduction plays an important role in the classification of strategies.

Step 3: Kernel-based probabilistic classification of all possible search patterns

We formalize the distance measure as observed in the example in figure 2.3 between any prototypical heuristic and a search pattern, in any n-dimensional representation. Any behavioral search patterns (s_j) can be represented as an n-dimensional point, $f_j : \{F_1(s_j), ...F_n(s_j)\}$. We then define a Gaussian kernel K for each strategy h_i.

$$K(s_j, h_i) = \exp\left(\frac{-||f_j - \bar{f}_i||^2}{2\sigma_i^2}\right) \tag{2.4}$$

This kernel defines a similarity measure in the range $[0,1]$, between any search pattern s_j, and the prototypical search pattern s_i for strategy h_i. This measure reaches a with a maximum value of 1 when $f_j = \bar{f}_i$. The parameter σ_i defines how rapidly the similarity measure drops off as the search point moves away from the prototype for heuristic h_i in the feature space. Depending on the purpose, this parameter can be designed a priori or post hoc after looking at the data, to minimize the loss function. We then define a probability distribution for observing search patterns in this feature space, conditional on this prototypical strategy being used, as in equation 2.5. Note that this distribution depends on the parameter σ_i. Selecting an infinitesimally small value for σ_i essentially reduces this representation to the rule-based heuristic, since in this case, $p(s_j|h_i) = 1$ if s_j is the prototypical search pattern, and 0 otherwise. As we increase the value of σ_i, this defines an increasing radial error distribution in the feature space.

$$p(s_j|h_i) = \frac{K(s_j, h_i)}{\Sigma_l K(s_l, h_i)} \tag{2.5}$$

In figure 2.3, the density of the radial lines around the TTB prototype represent the strength of this Gaussian TTB kernel at that point.

Step 4: Inferring the latent locus of learning

We can plug in the posterior probability in equation 2.5 into equation 2.3. The posterior probabilities have a well-formed gradient, and can make far more discernible inferences about which heuristic strategy is being used on each trial.

$$p(h_m|x_{jk}) = \frac{p(y_k|s_j h_m)\dfrac{K(s_j,h_m)}{\Sigma_l K(s_l,h_m)}p(h_m)}{\Sigma_{i=1:N_h}\left(p(y_k|s_j h_i)\dfrac{K(s_j,h_i)}{\Sigma_l K(s_l,h_i)}p(h_i)\right)} \tag{2.6}$$

Here, all the terms are defined in terms of the probabilistic definition of strategies, except $p(h_i)$, which is the prior probability of a strategy being used on any particular trial. This effectively serves as a method for Bayesian regularization, and will especially influence the reward-learning mechanism when a particular single trial is ineffective in inferring the underlying latent strategy. This is calculated on a trial by trial basis, based on equation 2.2, and depends on the previous $q-$values of the strategies. These $q-$values are iteratively updated based on rewriting equation 2.1 as equation 2.7.

$$q_{h_{m,t}} = p(h_{m,t-1}|x_{jk,t-1})\big((1-\eta)\ q_{h_{m,t-1}} + \eta\ r_{t-1}\big) + (1 - p(h_{m,t-1}|x_{jk,t-1}))\ q_{h_{m,t-1}} \tag{2.7}$$

2.4.2 Specific examples of defining probabilistic strategies

Example 1: Single dimensional feature space

We start with the simplest example, using only one feature of cue search patterns. We base this on the experimental paradigm used in Lee et al. (2014). Here, participants on each trial have access to 9 different cue attributes for two choice options. The cue validities are known to participants and attributes can only be selected in order of their validity, and selecting an attribute reveals the

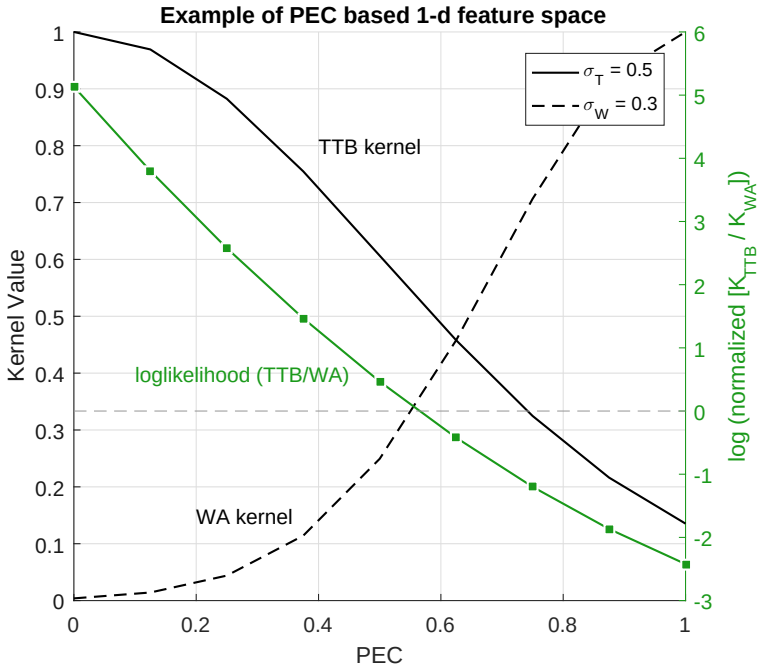

Figure 2.4: Kernels for asymmetrical σ values of $\sigma_{ttb} = 0.5$ and $\sigma_{wa} = 0.3$. The black lines show the kernel values as PEC increases from 0 to 1, and the green line shows how the log likelihood ratio of inferring TTB versus WA changes as PEC moves from 0 to 1.

value for both choice options. Thus, there are only 9 different possible cue search patterns, from selecting 1 attribute to 9 attributes. Traditional take-the-best (TTB) search patterns would imply selecting the minimum numbers of attributes required to discriminate between the two choices, and weighted average (WA) would imply selecting all the attributes. The original paper represents the proportion of extra cues (PEC) searched incremental to the first discriminating cue (FDC).

$$PEC = \frac{N_{cues} - FDC}{N_{total} - FDC} \tag{2.8}$$

In this case, the *PEC* provides a natural 1-dimensional psychological space, which varies discretely in the range [0,1]. Under the TTB prototype, $PEC_{ttb} = 0$, and under WA, $PEC_{wa} = 1$. Hence, for a particular search pattern s_j, we can write kernel densities for the TTB and WA prototypes as equations 2.9 and 2.10. Figure 2.4 visualizes such kernels for asymmetrical σ values of $\sigma_{ttb} = 0.5$ and $\sigma_{wa} = 0.3$. The black lines show the kernel values as PEC increases from 0 to 1, and the green line shows how the log likelihood ratio of inferring TTB versus WA changes as PEC moves from

0 to 1. The slope of this curve depends on the values of σ selected (or inferred from data).

$$K(s_j, h_{ttb}) = \exp\left(\frac{-||PEC_j||^2}{2\sigma^2_{ttb}}\right) \tag{2.9}$$

$$K(s_j, h_{wa}) = \exp\left(\frac{-||PEC_j - 1||^2}{2\sigma^2_{wa}}\right) \tag{2.10}$$

Example 2: Multidimensional feature space

Next, we consider the type of example demonstrated in section 2.1.2. We define a 3-dimensional psychological feature space, $f_j : [F_1(s_j), F_2(s_j), F_3(s_j)]$, so that any search pattern s_j can be expressed in terms of these features and represented as f_j.

F_1 : Proportion of cues searched

F_2 : Sensitivity to cue validity

F_3 : Variability in cues selected across attributes

We use the experimental structure of Bröder and Schiffer (2003), and take the cue search patterns suggested by TTB and WA, and calculate the feature vector corresponding to that search pattern, denoting these as $\bar{f}_{ttb}$ and $\bar{f}_{wa}$ respectively. For a cue space with $n_A = 4$ and $n_O = 3$, we get $\bar{f}_{ttb} = [0.25, 1.00, 0.50]$ and $\bar{f}_{wa} = [1.00, 0.25, 0.00]$. The squared Euclidean distance $||f_j - \bar{f}_{ttb}||^2$ is calculated as $\Sigma_{n=1:3}\left(F_n(s_j) - F_n(s_{ttb})\right)^2$, and similarly for WA. The kernel density (shown for TTB) is calculated as in equation 2.11.

$$K(s_j, h_{ttb}) = \exp\left(\frac{-\Sigma_{n=1:3}\left(F_n(s_j) - F_n(s_{ttb})\right)^2}{2\sigma^2_{ttb}}\right) \tag{2.11}$$

Example 3: Multidimensional feature space with skewed feature weights

Note that the kernel specification in example 2 essentially assumes equal weights on the three features. However, it is entirely feasible that individuals pay differential attention to these features. Theoretically, we can accommodate such individual differences by defining a mixture of kernels. Here the subscript p refers to the p^{th} individual, and w_{np} refers to the weight placed on the n^{th} feature by the p^{th} individual. These can be treated as free parameters during the inference process. This shows how this framework can be used to infer individual differences in attention to different aspects of the search space. Effectively, this skews the n-dimensional space by scaling each feature dimension by its corresponding weight, allowing for different scaling by individuals.

$$K(s_j, h_{pi}) = \Sigma_{n=1:3}\left[w_{np}\ \exp\left(\frac{-\left(F_n(s_j) - F_n(s_i)\right)^2}{2\sigma_i^2}\right)\right] \tag{2.12}$$

Example 4: Context specific feature space

The above examples all focused on prototypical representations of traditional heuristics such as TTB and WA, and generalizing the statistical properties of cue search patterns. However, since we generalize strategies in terms of a psychological feature space, we can define strategies in a context-specific manner. For instance in the example shown in section 2.1.2, two of the attributes were financial metrics, while the remaining two were advisory recommendations (Mistry and Trueblood (2015)). It is perfectly reasonable that people search the attribute space based on their own level of expertise (whether or not they are adept in evaluating financial metrics) and prior perceptions of trust in advisory recommendation. We can define features in terms of proportion of advisory cues searched compared to total cues searched. A *full-advisory* strategy will be represented by a feature value of 1, and a *full-financial* heuristic by a feature value of 0. The kernel and probability calculations can then proceed as in the previous examples.

2.5 Conclusions

I suggest that traditional rule-based heuristics need to be redefined in a probabilistic sense to provide a more meaningful analysis of the cognitive process involved in multi-attribute decision making, especially as it applies to learning and strategy switching. A prototypical rule-based heuristic such as take-the-best may be a reasonable approximation to an underlying cognitive process, however, most implementations assume that any deviation from the rule arises randomly, or based on a uniform distribution of error. Instead, we recast such rule-based heuristics probabilistically, so that the error gradient for deviation from the prototypical rule is well-defined. We thus propose a novel method for thinking about heuristic strategies in multi-attribute decisions.

The key contributions in this chapter include, identification of a significant challenge in the cognitive modeling of strategy learning and switching, the use of Bayesian regularization for identifying the latent locus of learning, and introduction of a novel probabilistic framework for defining heuristic strategies and classifying information search patterns. Chapter 3 reports some new experimental results within non-stationary multi-attribute decision tasks. Chapter 4 then uses the novel probabilistic approach introduced in this chapter for cognitive modeling of these tasks as well as other previous published studies.

Chapter 3

Experimental Results

3.1 Motivation

Learned strategies could be driven to extinction because they no longer provide the expected level of outcomes (or result in negative feedback), or because an alternate strategy is explored and shown to provide either superior outcomes or similar outcomes with lower costs. Routinization, a failure of extinction of strategies in the face of change, has been reported in various studies. It has been described as the phenomenon where people are good at learning initial environmental contingencies, but over time internalize these learned routines, thus responding slowly or ineffectively to subsequent changes (Bröder and Schiffer (2006a); Bröder et al. (2013)). Routinization has not been consistently observed, and seems to manifest in different ways. Some studies (Gluth et al. (2013)) have reported an initial lag in adaptation of strategies, some (Bröder and Schiffer (2006a)) found sustained lower performance in response to a change in the environment, and yet others (Lee et al. (2014); Racey et al. (2011)) found that people were in fact quite efficient in changing behavior to shifting environmental patterns. Routinization has been shown to be susceptible to moderation by the type of feedback provided (Bröder et al. (2013)), by the strength of the initial

routines and the surface properties of the task (Betsch et al. (2001)), and by time pressure (Betsch et al. (1999)). One obvious possible reason for routinization might be the failure to detect a change in the environmental structure or contingencies. There has however been evidence of routinization even when participants were provided hints regarding the timing of possible changes and additional incentives for correctly identifying optimal strategies (Bröder and Schiffer (2006a)).

3.1.1 Primitives of possible routinization effects

In our studies, we wanted to focus on what aspects of environmental structure and change affect routinization when individuals are made aware of the possible timing of changes in environmental structure. Most deterministic non-stationary MADM tasks (Bröder and Schiffer (2006a); Bröder et al. (2013)) involve a single change of environmental contingencies from compensatory to non-compensatory or vice versa. Most real world environments would involve multiple changes, including reverting back to previously experienced states from time to time. It is not clear in such situations, how the length of static periods and the frequency of change affect exploration and routinization. Further, renewal of cue search and combination strategies when previously encountered environments are extinguished and then presented again after intervening environmental changes has not been tested. We attempt to identify if such renewal takes place, and if so, what characteristics of environmental structure and change does it depend on? Specifically, is a failure to respond to change dependent on, (1) the degree of change in the environment (i.e. if the difference in performance of the incumbent and new strategy is not significant enough to change beliefs, (Betsch et al. (2001))), (2) the nature of the post-change environment (i.e. is it difficult identify optimal strategies in specific environments leading to a reversion to previous behavior), or (3) sub-optimal learning processes (i.e. selective belief updating (Wilson and Niv (2011)))?

In response to these questions, our first design manipulation is *routine*; consecutive routines are when blocks with similar environmental types are placed consecutively after each other, and alter-

nate routines are when compensatory and non-compensatory blocks are interleaved.

3.1.2 Asymmetries in strategy switching

Pearson et al. (2011) suggested that the incumbent strategy has an impact on the reward prediction error, and subsequently, on the change detection mechanism. This could lead to asymmetrical learning and adaptation, depending on what initial bias people have, and what environmental conditions they first encounter. Asymmetrical mechanisms are explored by Newell and Lee (2009); Lee et al. (2014). They found that people shift their strategies from a limited to an exhaustive search and vice versa, depending on the environmental conditions. However the process of switching between these information-search strategies is different. Feedback revealing a decline in accuracy, or lower confidence is proposed as the key reason for switching from limited to exhaustive search strategies. On the other hand, a switch back to limited search strategies is observed even when there is no adverse impact on accuracy or confidence. This is explained by means of an effort reduction mechanism. Bröder and Schiffer (2006a) suggested that people had an initial propensity to prefer compensatory strategies, and that while routinization effects existed for both compensatory to non-compensatory shifts and vice versa, the routinization effect was larger when people had to shift from compensatory to non-compensatory strategies. However, Rieskamp (2008) reanalyzed the same data using a reinforcement learning strategy selection model and proposed that participants in fact had a higher initial preference to use a non-compensatory (TTB) strategy in all the environments. Rather than demonstrating routinization of a compensatory strategy, the analysis suggests an initial preference for a TTB strategy which is gradually unlearned and shifts towards a compensatory strategy. The original analysis depended on performance based measures, comparing choice selection to optimal strategies. The reanalysis depended on process measures by factoring in the probabilities of using each type of strategy, learned on the basis of a reinforcement learning model.

Orthogonal to routine, we manipulate *starting condition*, where participants either encounter compensatory blocks first or non-compensatory blocks first. All participants however are presented with two compensatory and two non-compensatory blocks within-subject repeated measures.

3.1.3 Degree of change and positive versus negative cue contingencies

Most MADM tasks vary cue weights and validities across a range of positive values. There is evidence of confirmation bias, or positive test strategy, whereby people are prejudiced towards aspects that have previously produced positive results (Klayman and Ha (1987)). Le Pelley et al. (2011) reported results that cues with a high level of predictive power resulted in a higher attentional bias. Beesley et al. (2015) suggested that attention bias towards the exploitation of predictive cues was more robust than an attention bias towards exploratory behavior arising from increasing uncertainty about cues. Rolison et al. (2011) proposed that learning about cues that are negatively associated with outcomes is more difficult than learning about positively associated cues because such learning involves greater use of working memory capacity as well as deliberative attention and control processes. They proposed that positive and negative cue learning both involve explicit hypothesis testing, but only learning of negative cues seems to involve explicit application of beliefs to judgments. These assessments were made under stationary conditions. However, these have not been tested in non-stationary conditions involving a shift between cues that are only positively related to the outcome, versus cues which have a mixed positive and negative relationship. If the relationship were to hold, this would require that optimal strategy switching also involves a change in the deliberative control, attention and working memory related processes. Would renewal of previously learned strategies be faster when only positive cues are involved, with greater autonomous control, or when negative cues are involved, with greater working memory and deliberative control?

Hence, we vary the nature of the non-compensatory environments between experimental studies (subjects), to explore the impact of degree of change on routinization and learning mechanisms. In

study 1, all environmental conditions only involve positive cue-outcome relationships. In study 2, we test responses to shifts between positive-only versus positively and negatively associated cue structures. These shifts result in a larger degree of change in the environmental conditions than in study 1, and enables us to investigate how the difference in pre-change versus post-change signals affects learning, exploration and adaptive strategy switching.

3.2 Methods

3.2.1 Study 1

In this study, we varied the statistical contingencies linking information cues with outcome criteria between blocks within participants, to measure how well individuals could detect and respond to changes, and whether the nature of the change resulted in different levels of adaptivity, or lack thereof. 32 University of California Irvine undergraduate students participated in the experiment for course credit[1].

Stimulus

We used a multiple cue learning task, where participants had to make repeated forced choices between one of three options on the basis of a set of underlying cues. The cover story for the choice task was a hypothetical stock market game, in which participants had to choose between three financial stock options over 120 trials. Figure 3.1 shows how one trial might appear to a participant. In each trial, each of the options was associated with four binary cue attributes, namely past profit growth, sales growth, and recommendations from two independent advisors. The binary cue values were represented as a + or - for all the cues. On each trial, participants had

[1] 15 additional participants experienced a technical system failure during their session and were unable to complete the experiment

the option to individually acquire as many of the twelve possible cues (3 options X 4 cue attributes per option) in any order. Each choice option was associated with a reward, and acquiring cues cost participants 4% of the total gross positive rewards obtained for that trial, for each cue acquired. Once a participant selected a cue attribute, it remained visible throughout the trial. Whilst this did not allow us to calculate process tracing variables that depended on the time spent on each attribute, it also did not restrict participants to serial processing of cues, which has been proposed to put constraints on the choice of strategy (Glöckner and Betsch (2008)), and did not impose working memory constraints which have been known to affect cue interpretation (Rolison et al. (2011)).

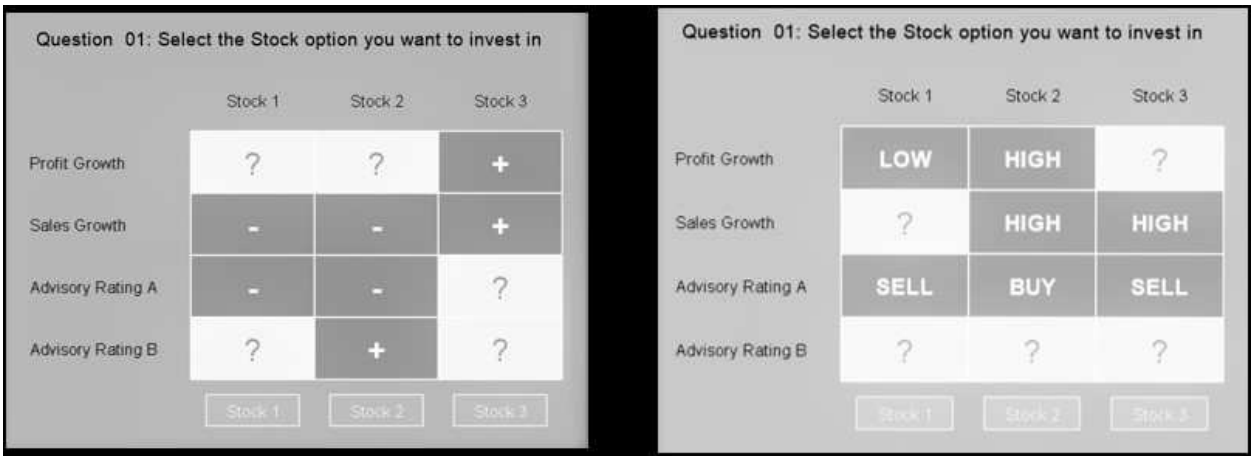

Figure 3.1: Stimulus

Procedure

The 120 trials were segregated into four blocks of 30 trials each, with each participant facing two compensatory (C-Block) and two non-compensatory (N-Block) blocks. There were no practice trials since we wanted to measure the learning of cue selection strategies, and wanted to make sure that participants were not biased by any information obtained from the practice trials. The block size was designed to be small (30 trials each) to manipulate the possible effects of routinization of decision strategies. The experiment had a factorial design with four different between-subject conditions based on the order in which the four blocks were presented (2 conditions depending

on whether the starting block had a compensatory (C) or non-compensatory (N) environment X 2 conditions where the four blocks alternated in environmental condition versus when the two same conditions were grouped together). In two conditions, both compensatory blocks and both non-compensatory blocks appear consecutively, whereas in the remaining two conditions, participants always encountered a different environmental condition in alternate blocks. Essentially, the four blocks in the four conditions can be represented as CCNN, CNCN, NCNC and NNCC. This allowed us to measure the interaction between routine length and starting conditions. We wanted to test for differences in adaptivity and performance between these conditions. A secondary objective of this task was also to measure how quickly preferences could change, and whether the relatively frequent environmental changes led to stronger or weaker learning. Compensatory and non-compensatory blocks differed in the relationship between the cue values (c_1 to c_4; encoded as either +1 or -1) and the resulting gross rewards (r) associated with these sets of cues.

$$r_C = 32c_1 + 26c_2 + 22c_3 + 20c_4 + U(-8,8) \tag{3.1}$$

$$r_N = 47c_1 + 25c_2 + 17c_3 + 10c_4 + U(-8,8) \tag{3.2}$$

The rewards for any choice option thus ranged from -108 to $+108$, depending on the cue configuration. A random component, drawn from a uniform (-8, 8) distribution was also incorporated into this relationship. It can be seen that r_N is somewhat, but not strictly, non-compensatory, since there are very few possible configurations where a combination of cues of lower validity can override the choices determined by a higher validity cue. On the other hand, r_C is somewhat, but not strictly compensatory. The cue weights are very comparable, but not all equal. This set of compensatory and non-compensatory validities was picked since they have been used in similar experiments in the past (Bröder and Schiffer (2006a)), and used to draw conclusions about the nature of learning and routinization. They involve strictly positive cue-outcome relationships (i.e. cue weights were all positive). After each trial, the gross rewards (r_{gross}) for all of the options were shown to the participant. They were also provided with the gross reward for the specific option they had chosen, as well as the costs incurred depending on the number of individual cues selected (n_{cues}), and the

net rewards(r_{net}). These costs were incurred only if the gross rewards were positive. Thus, the net rewards were calculated as in equation 3.3. Since a maximum of 12 unique cues could be selected, the maximum reduction from gross to net rewards was $0.52r_{gross}$. The objective of the task for participants was to maximize the net rewards remaining after any cue search related costs were deducted.

$$r_{net} = \big(1 - (0.04n_{cues})\big)r_{gross} \quad \text{if } r_{gross} > 0; \quad = r_{gross} \quad \text{otherwise} \tag{3.3}$$

For a single participant, the sequential ordering of cues from highest to lowest was maintained across blocks, and only the relative cue weights were changed (i.e. the cue with the highest weight in the compensatory blocks remained the cue with the highest validity in the non-compensatory blocks). The actual cue weights, the order of importance of the weights, or the nature of the environment (i.e. compensatory or non-compensatory) were not communicated to the participants. However, in all four conditions the participants were explicitly told that the underlying environment and relationships between cues and options would remain the same within a block, could change between blocks, and the start and end of each block were clearly demarcated.

In this task, a take-the-best (TTB) strategy always provided a higher net payoff (after factoring in information acquisition costs) in N-blocks (although a compensatory strategy such as WA provided the same gross payoff), and a weighted average (WA) strategy provided a higher payoff in C-blocks.

3.2.2 Study 2

34 University of California, Irvine undergraduates participated in the experiment for course credit.

Stimulus

The cover story and structure of the task, conditions and environmental blocks was very similar to study 1, with the following set of differences. Acquiring the cues cost participants 5% of the total gross positive rewards obtained for that trial, for each cue acquired. Instead of + and -, the cue values were represented as High or Low for the first two cue attributes (financial indicators) and as Buy or Sell for the other two (advisory recommendations). For calculating the rewards, High / Buy were encoded as +1 and Low / Sell encoded as -1. The gross rewards in C-blocks and N-blocks were calculated as:

$$r_C = 40\ c_1 + 37\ c_2 + 34\ c_3 + 3\ 1c_4 + noise(-8,8)$$

$$r_N = 78\ c_1 + 7\ c_2 - 21\ c_3 - 36\ c_4 + noise(-8,8)$$

These rewards for any choice option thus ranged from -150 to $+150$, depending on the cue configuration. A random component drawn from a U(-8,8) distribution was also incorporated into this relationship. It can be seen that r_N is strictly non-compensatory, since no combination of [c2, c3, c4] can override the reward determined by [c1], and similarly, no combination of [c2, c3] can override [c4]. On the other hand, r_C is; somewhat, but not strictly compensatory. The cue weights are comparable, but not all equal. The presence of negative cue weights was an important difference between the studies. Whilst the actual cue weights or order was not disclosed, participants were explicitly told that it was possible for cues to be negatively related to the options (this was justified within the paradigm, for example, one reason could be that a particular investment advisor was consistently incorrect). The final difference was that unlike in study 1, not only were the cue weights changed, but the sequential ordering of cues was also changed between C-blocks and N-blocks. This amplified the difference between C-blocks and N-blocks, to push participants towards a more deliberative cognitive effort. The rest of the procedure was the same as for study 1.

3.3 Behavioral Results

Table 3.1 shows the log Bayes factors (LBF) for a Bayesian repeated measures ANOVA (JASP-Team (2016)) that tests for the effect of block type (compensatory[C] versus non-compensatory[N]), starting blocks (C[includes CNCN/CCNN] vs N[includes NCNC/NNCC]), sequence (alternating[includes CNCN/NCNC] versus consecutive[includes CCNN/NNCC]), and interactions between these factors. The dependent behavioral variables include cues (proportion of total cues searched), attributes (proportion of unique attributes where at least one cue was searched), validity (average normalized sensitivity to validity), best (proportion of trials on which the best option was selected), worst (proportion of trials on which the worst option was selected), and reward (standardized reward scores based on net rewards after search costs). LBF > 1 indicates evidence in favor of including the factor and LBF < -1 in favor of the null (factor not significant). Larger values imply greater evidence of a difference. Table 3.2 shows the mean values for the dependent variables. The rows show the total values for the C and N blocks, but also the breakup between alternating (Alt) and consecutive (Cons) routines within each of these block types. The values in bold correspond to values that show a significant effect based on the Bayesian ANOVA analysis presented in table 3.1.

Table 3.1: Bayesian ANOVA for key behavioral results

	log(Bayes Factor)	Cues	Attributes	Validity	Best	Worst	Reward
Experiment 1	Block Type (C vs N)	-1.8	-2.0	**1.2**	**5.0**	**4.0**	**6.5**
	Starting (C vs N)	-0.9	-0.7	-1.3	-1.0	-1.0	-0.8
	Sequence (Alt vs Cons)	-0.9	-0.4	-0.5	-0.7	-0.5	-0.1
	Type X Start	-1.6	-2.1	-0.8	-0.9	-0.9	-0.5
	Type X Sequence	-1.7	-2.2	0.6	-0.7	-0.8	-0.3
	Start X Sequence	-1.0	-0.2	-0.8	-1.3	-0.7	-0.8
Experiment 2	Block Type (C vs N)	**2.8**	**4.4**	**8.3**	-1.0	-1.2	-0.9
	Starting (C vs N)	-1.0	-0.8	-0.9	-0.2	-0.5	-0.1
	Sequence (Alt vs Cons)	0.9	0.3	-1.3	-1.2	-1.2	-1.3
	Type X Start	-0.7	0.05	-0.9	0.3	-0.2	0.4
	Type X Sequence	**2.2**	**1.6**	-1.2	-1.4	-1.1	-1.7
	Start X Sequence	-1.1	-0.6	-0.7	-0.4	-1.1	-0.7

Table 3.2: Mean value of behavioral variables

Mean Values		Cues	Attributes	Validity	Best	Worst	Reward
Experiment 1	C - blocks	0.48	0.68	**0.44**	**0.68**	**0.11**	**0.70**
	C - Alt	0.48	0.66	0.43	0.69	0.09	0.72
	C - Cont	0.48	0.71	0.44	0.64	0.12	0.69
	N - blocks	0.49	0.68	**0.47**	**0.76**	**0.06**	**0.76**
	N - Alt	0.47	0.64	0.5	0.79	0.05	0.8
	N - Cont	0.51	0.72	0.45	0.72	0.08	0.74
Experiment 2	C - blocks	**0.33**	**0.51**	**0.47**	0.82	0.08	0.80
	C - Alt	**0.29**	**0.46**	0.47	0.80	0.11	0.79
	C - Cont	**0.37**	**0.56**	0.46	0.85	0.05	0.82
	N - blocks	**0.28**	**0.41**	**0.70**	0.83	0.09	0.78
	N - Alt	**0.29**	**0.43**	0.69	0.84	0.08	0.78
	N - Cont	**0.27**	**0.39**	0.71	0.83	0.09	0.78

There is evidence for a significant effect of the block type on performance metrics (best, worst, reward) in experiment 1, with better performance in N-blocks (best option 0.76, standardized reward 0.76) compared to C-blocks (best option 0.68, standardized reward 0.70). There is evidence for a significant effect of block type on search metrics (cues, attributes, sensitivity to validity) in experiment 2, with lower depth of search in C-blocks (proportion of attributes searched 0.51) versus N-blocks (proportion of attributes search 0.41), and significantly higher sensitivity to validity in N-blocks (0.70) versus C-blocks (0.47). Starting block type does not seem to have any significant influence in both experiments. The type of sequence has an interaction effect with the block type only in experiment 2, influencing the proportion of cues and attributes selected, with higher cue and attribute search in the continuous condition versus in the alternate condition for C-blocks, but lower cue and attribute search in the continuous condition versus in the alternate condition for N-blocks. Essentially, in the continuous condition, participants shift their information search patterns towards the more efficient mechanism (higher cue search in C-bocks and lower in N-blocks) more efficiently than they do in the alternating condition. Thus, higher frequency of change (alternating sequence) hampers effective adaptation of cue search patterns. In experiment 1, the shift in environment from C to N does without increased efficiency of search does not penalize the gross

rewards, only the information search costs. Thus, the magnitude of feedback provided in experiment 1 is not strong enough to warrant a change in the information search strategies in N-blocks.

3.3.1 Individual Differences

Figures 3.2 and 3.3 show the observed information search and performance metrics split by half-blocks, that is, in the first and second half of each of the 4 blocks, C-1 (first compensatory block), C-2 (second compensatory block), N-1 (first non-compensatory block), and N-2 (second non-compensatory block). The lines thus show increasing of decreasing trends within each block. The lines are color coded by condition, the thin dashed lines show individual performance, and the thick lines show the mean levels for each condition. It is clear that there are significant individual differences in terms of the search and performance metrics, and in terms of how people change their behavior within and across blocks.

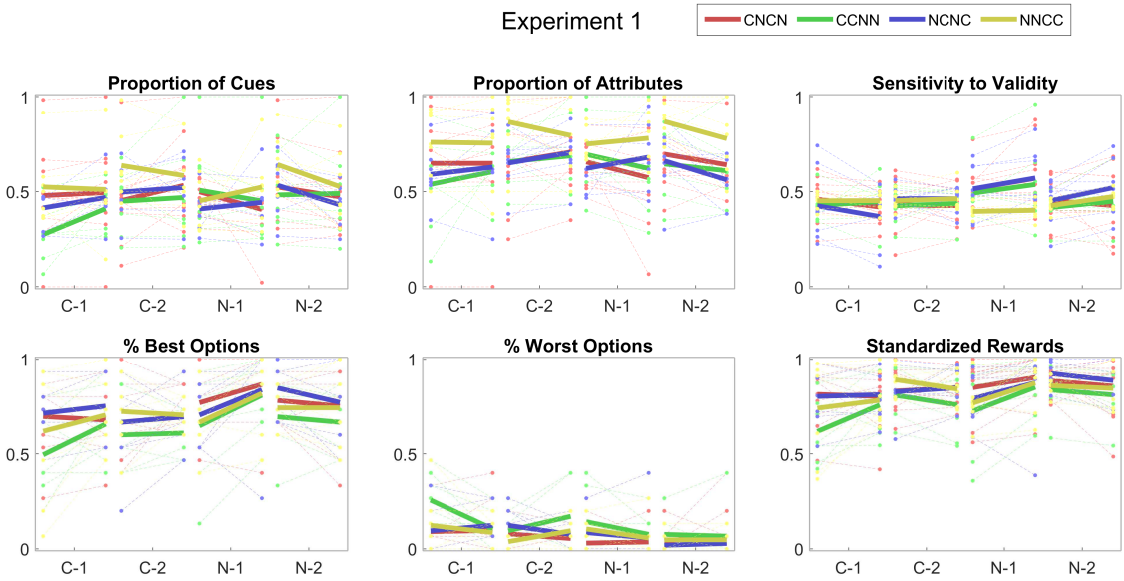

Figure 3.2: Experiment 1: Key behavioral results summarized in the first and second half of each block. Results are grouped by the type of block (C1-C2-N1-N2). Each dashed line is an individual participant and the thick lines are the mean over participants. The colors indicate the different between subject conditions.

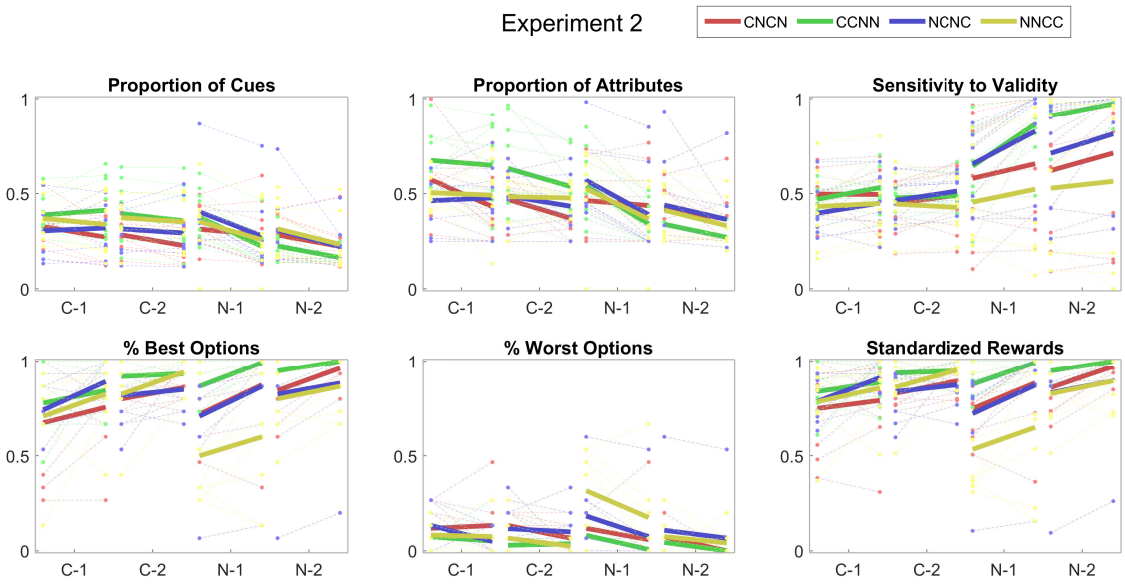

Figure 3.3: Experiment 2: Key behavioral results summarized in the first and second half of each block. Results are grouped by the type of block (C1-C2-N1-N2). Each dashed line is an individual participant and the thick lines are the mean over participants. The colors indicate the different between subject conditions.

3.3.2 Within-block learning

Tables 3.3 and 3.4 show the differences between 1st half and 2nd half of blocks, with significant differences highlighted in bold. In experiment 1, most differences are not significant. In experiment 2, there is very strong evidence for changes in both search and performance metrics, with stronger effects in N-blocks compared to C-blocks, indicating stronger learning effects than experiment 1 in the presence of greater reward dispersion, but also stronger learning effects within N-blocks compared to C-blocks.

3.3.3 Cue search patterns

Figures 3.4 to 3.7 illustrate the proportion of cues searched and sensitivity to true validity, two of the most important search features, over trials and split by condition. The thin lines show individual participants in the condition, and the bold lines show the mean. The blue background represents

Table 3.3: Bayesian ANOVA for key behavioral results - within block. This compares the performance in the first and second half of each of the blocks, grouped separately for C-blocks and N-blocks, to measure whether there is within-block learning, manifested in terms of significant within-block changes in search patterns or performance metrics. Values in bold indicate LBF > 1.

Half-Block	log(Bayes Factor)	Cues	Attributes	Validity	Best	Worst	Reward
Experiment 1	C-blocks	-0.09	-0.74	-0.96	-0.43	-1.10	-1.30
	N-blocks	0.74	**1.10**	**4.0**	0.83	-0.87	-0.89
Experiment 2	C-blocks	0.36	**2.60**	0.94	**4.93**	0.35	**3.96**
	N-blocks	**9.99**	**12.1**	**10.5**	**10.3**	**7.87**	**15.0**

Table 3.4: Mean values of key behavioral results grouped by values in the first half and second half of blocks, grouped separately for C-blocks and N-blocks. These values correspond to the Bayesian ANOVA shown in table 3.3. Values in bold correspond to LBF > 1.

Half-Block	Mean values	Cues	Attributes	Validity	Best	Worst	Reward
Experiment 1	C-block 1st Half	0.47	0.67	0.44	0.65	0.11	0.71
	C-block 2nd Half	0.50	0.69	0.43	0.69	0.10	0.70
	N-block 1st Half	0.51	**0.70**	**0.46**	0.74	0.06	0.76
	N-block 2nd Half	0.46	**0.65**	**0.49**	0.79	0.05	0.78
Experiment 2	C-block 1st Half	0.34	**0.54**	0.46	**0.78**	0.09	**0.77**
	C-block 2nd Half	0.32	**0.48**	0.48	**0.86**	0.07	**0.83**
	N-block 1st Half	**0.32**	**0.46**	**0.65**	**0.78**	**0.12**	**0.73**
	N-block 2nd Half	**0.23**	**0.36**	**0.76**	**0.89**	**0.05**	**0.84**

C-blocks and the yellow environment N-blocks. Experiment 2 shows better instances of learning to select fewer cues, with small exploration peaks towards the start of each block. This is not seen in experiment 1, where the reward schemes were not as diversified. Participants in experiment 2 also seem to learn the true cue validities better, showing higher sensitivity in N-blocks.

3.3.4 Identification of heuristic strategies

Figures 3.8 to 3.13 show what proportion of each trial can be identified uniquely as TTB or WA, is identified as a multiple match (Both), or a no match (neither), based on the 3 commonly used inference methods of matching choice only, choice matching plus minimum cues acquired, and choice match plus exact cues acquired. Using the first, choice only, most choices are compatible with both strategies, whereas using the latter information search information, most trials cannot

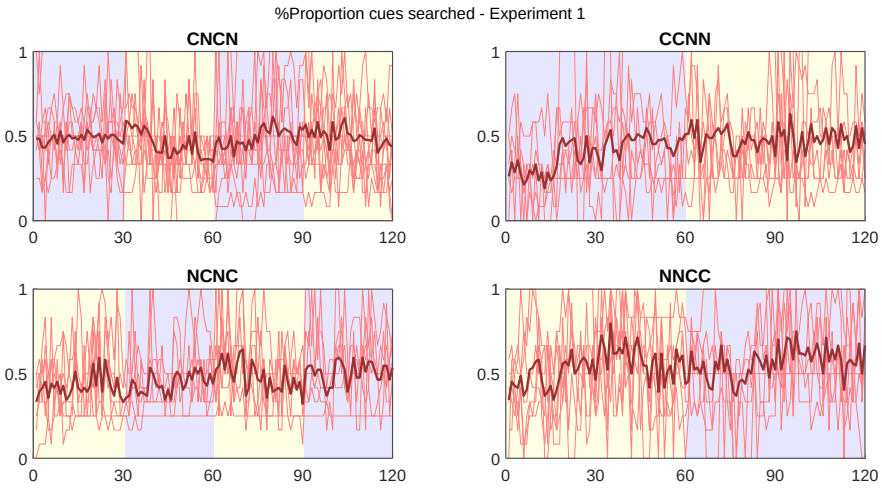

Figure 3.4: Experiment 1: Proportion of cues searched over trials, by condition. Thin lines are individuals, bold lines are the mean values. Blue background is C-blocks and yellow background is N-blocks.

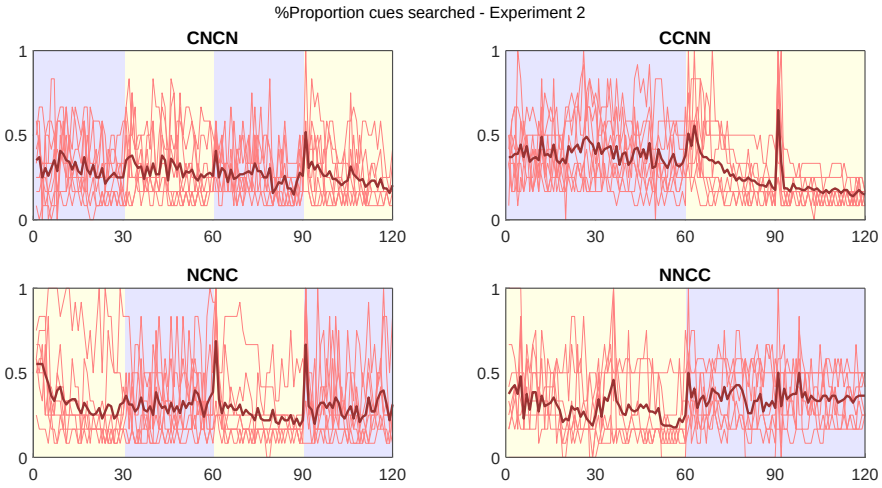

Figure 3.5: Experiment 2: Proportion of cues searched over trials, by condition. Thin lines are individuals, bold lines are the mean values. Blue background is C-blocks and yellow background is N-blocks.

be classified into any of the two strategies. This highlights the inference problem raised in the previous chapter.

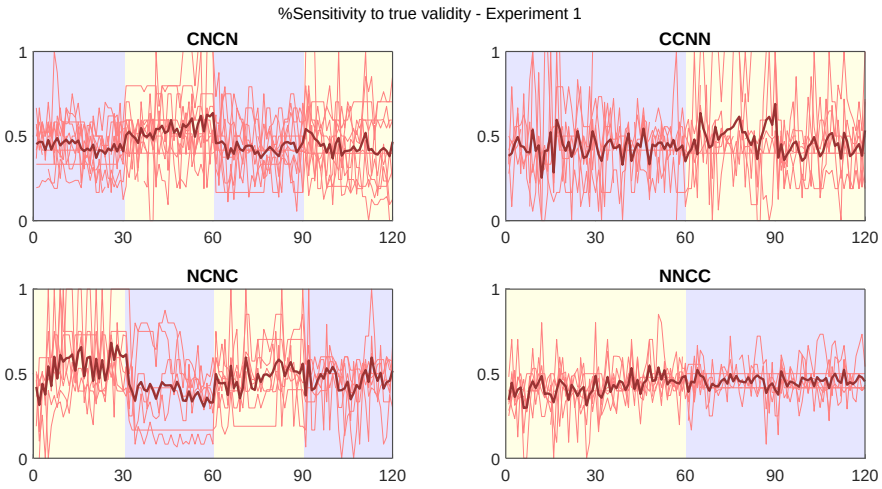

Figure 3.6: Experiment 1: Sensitivity to true validity, over trials, by condition. Thin lines are individuals, bold lines are the mean values. Blue background is C-blocks and yellow background is N-blocks.

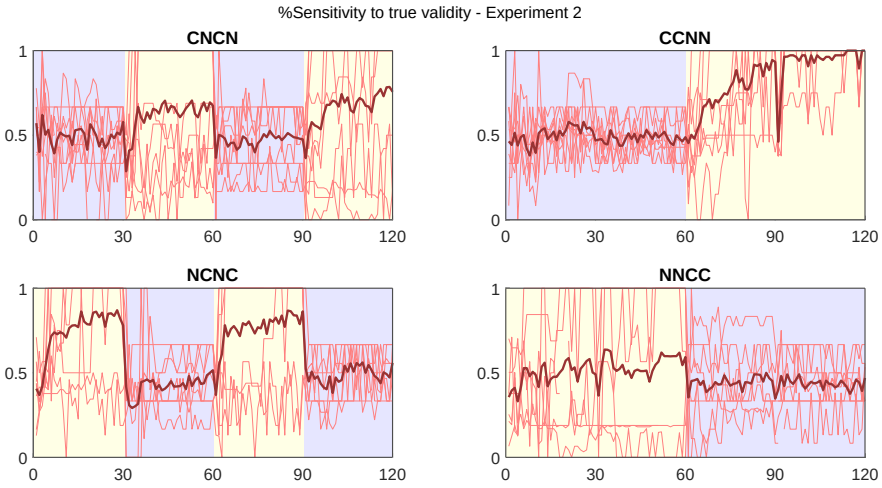

Figure 3.7: Experiment 2: Sensitivity to true validity, over trials, by condition. Thin lines are individuals, bold lines are the mean values. Blue background is C-blocks and yellow background is N-blocks.

3.4 Discussion

Both studies show an across-block learning effect, however this is much stronger in study 2. Since C-blocks in the two studies are quite similar, the contrast between N-blocks and C-blocks seems to be an important aspect of the task structure impacting the difference between the results observed in the two studies. This contrast between environmental conditions may be manifested in terms

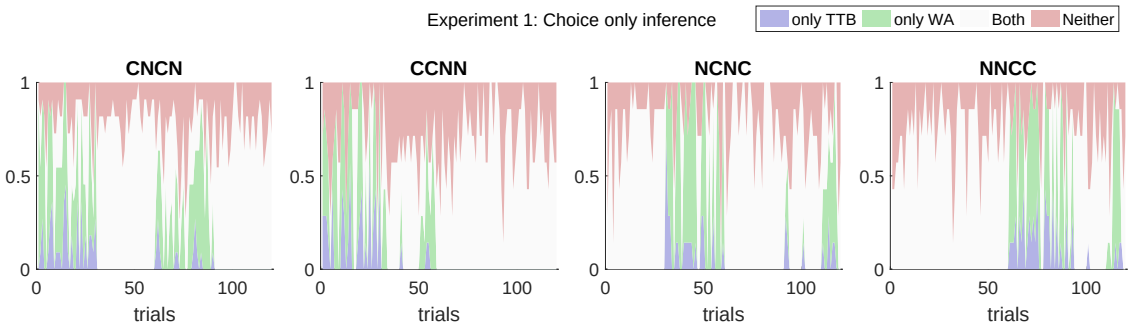

Figure 3.8: Choice matching

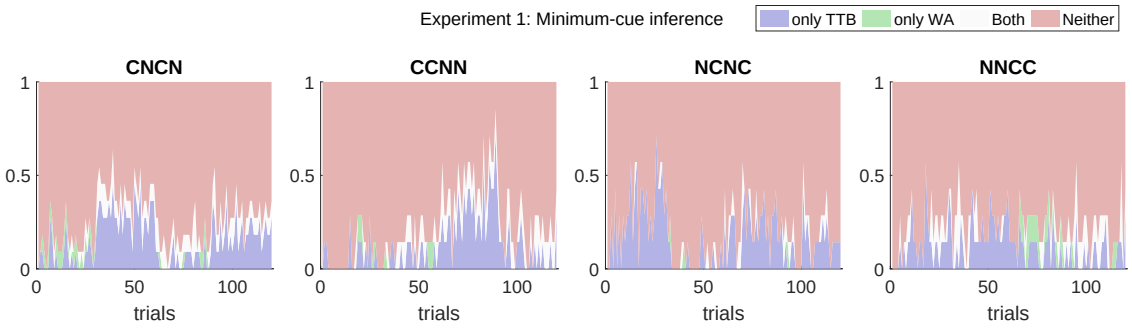

Figure 3.9: Minimum cue

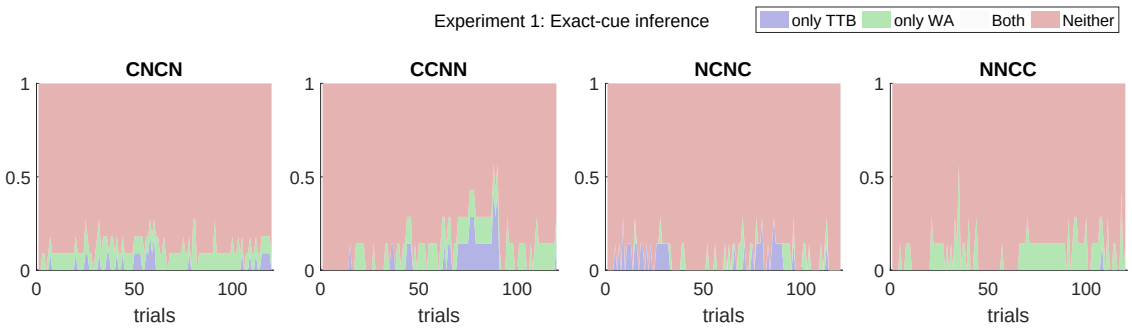

Figure 3.10: Exact cue

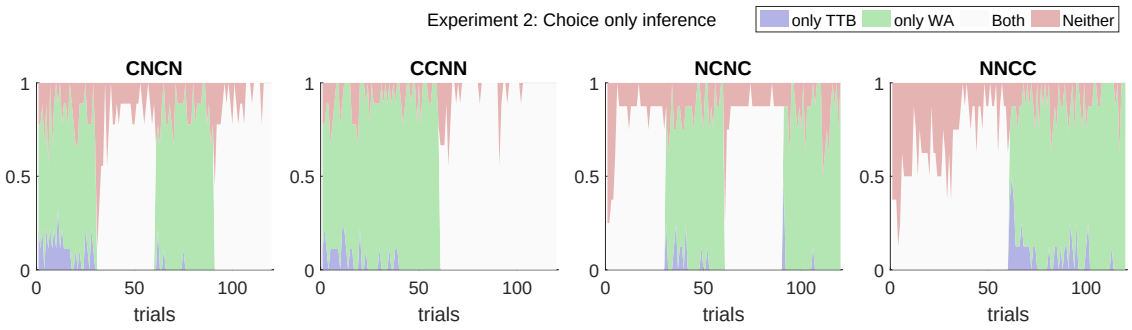

Figure 3.11: Choice matching

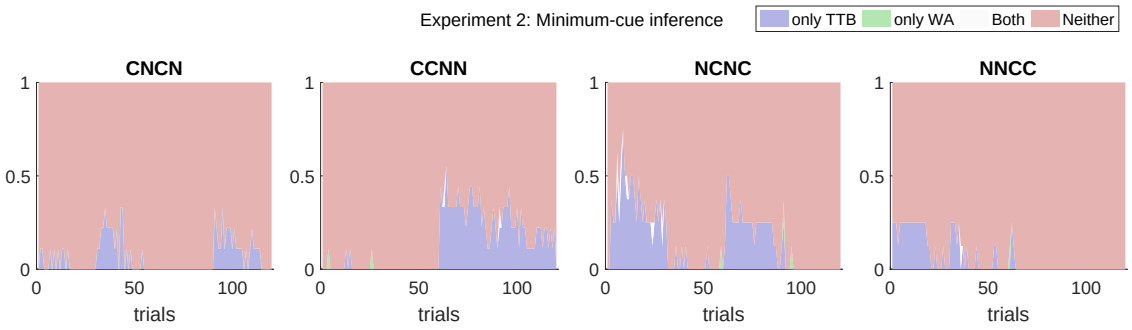

Figure 3.12: Minimum cue

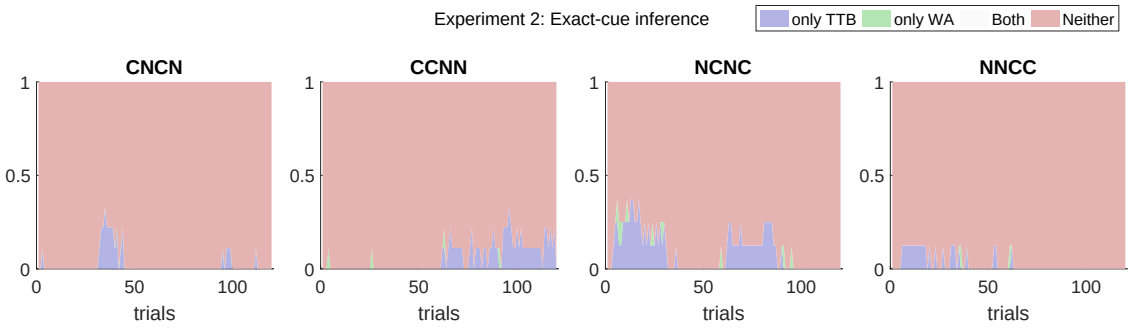

Figure 3.13: Exact cue

of the difference in level of rewards or feedback that complementary strategies receive in these environments. Task structure has a significant effect on the percentage of trials on which the best option is selected (0.72 for experiment 1 versus 0.82 for experiment 2) and on the standardized reward scores (0.73 for experiment 1 versus 0.79 for experiment 2), however this difference is only significant in C-blocks, with performance being significantly higher for the C-blocks in study 2 (best option selected 0.82 versus 0.68 for study 1). This is interesting considering that the C-blocks are more similar between the two studies compared to N-blocks. The higher C-block performance in study 2 is assumed to result from the greater contrast between C-blocks and N-blocks in this study.

There is another interesting aspect to the performance in N-blocks. In study 2, the first N-block performance is lower than that in study 1, but the second N-block performance is much higher. Study 2 shows an increase in the sensitivity to validity, reduction in the number of cues searched, and higher renewal of search patterns in the second N-block encountered. All of these point towards

strong learning and refinement of the strategies used in N-blocks over trials. We posit that this improved learning in N-blocks in study 2 is a result of the wider reward dispersion, a function of the higher dispersion in cue weights, including negative weights. However this same property also leads to a slightly lower N-block performance at the start of the first N-block in study 2 compared to study 1, since random exploration leads to lower performance under this configuration.

Task structure also has a significant effect on the percentage of cues selected, percentage of attributes selected, and sensitivity to validity. Importantly, study 1 participants showed an increase in the cue and attribute use across block sequences, whereas study 2 participants showed the opposite effect. The cue use was also similar between C-blocks and N-blocks in study 1, but was lower for N-blocks compared to C-blocks in study 2. Once again, the greater change in environmental contingencies impacted the cue search patterns. Importantly, the fact that N-block search entailed fewer cues implies that participants not only recognized that there was a change, but seemed to understand the non-compensatory nature of this changed environment.

The slight fall in performance in N-blocks in study 1 is accompanied by higher levels of cue use compared to N-blocks in study 2, but lower sensitivity to validity, compared to study 2. We thus reckon that the lower performance is not due to some form of routinization effect in study 1. Rather, participants seem to be aware of a change but find it difficult to recognize the correct dominant cues in N-blocks in study 1. They seem to sustain higher levels of exploration and consequently lower performance.

In terms of inferring which heuristics are being used, both minimum cue matching and exact cue matching methods of inference yield poor results, with a huge majority of trials in both experiments remaining unclassified as either TTB or WA.

Chapter 4

Cognitive modeling

4.1 Introduction

In this chapter, we fully specify the cognitive model of adaptive learning and strategy switching in multi-attribute choice, incorporating both traditionally defined and probabilistic strategies, as defined in chapter 2. The models are implemented as Bayesian graphical models. We show the relative descriptive and predictive performance of the models, the nature of inferences we can make, and identification of individual differences in adaptive decision making.

4.2 Data

4.2.1 Cognitive modeling of multi-dimensional feature space

For this, we use the two novel experiments described in chapter 3.

4.2.2 Cognitive modeling of single-dimensional feature space

For this, we use data from a secondary dataset, obtained from work published in Lee et al. (2014). In this study, participants answered 200 two-alternative forced choice multi-attribute decision problems, where they had to decide which of two objects was higher on a particular criterion. To make a decision participants could access information from between one to nine information cues on each trial. These cues provided binary information about each alternative. Cues could only be accessed in descending order of validity (known to participants) and selecting a cue provided information about both the alternatives. The experiment had a total of 200 trials which were subdivided into three blocks (unknown to the participant), with the statistical structure of the underlying environment changing between blocks so that some blocks required a more exhaustive search of cues (compensatory environment) whereas in some blocks a non-compensatory approach (making a

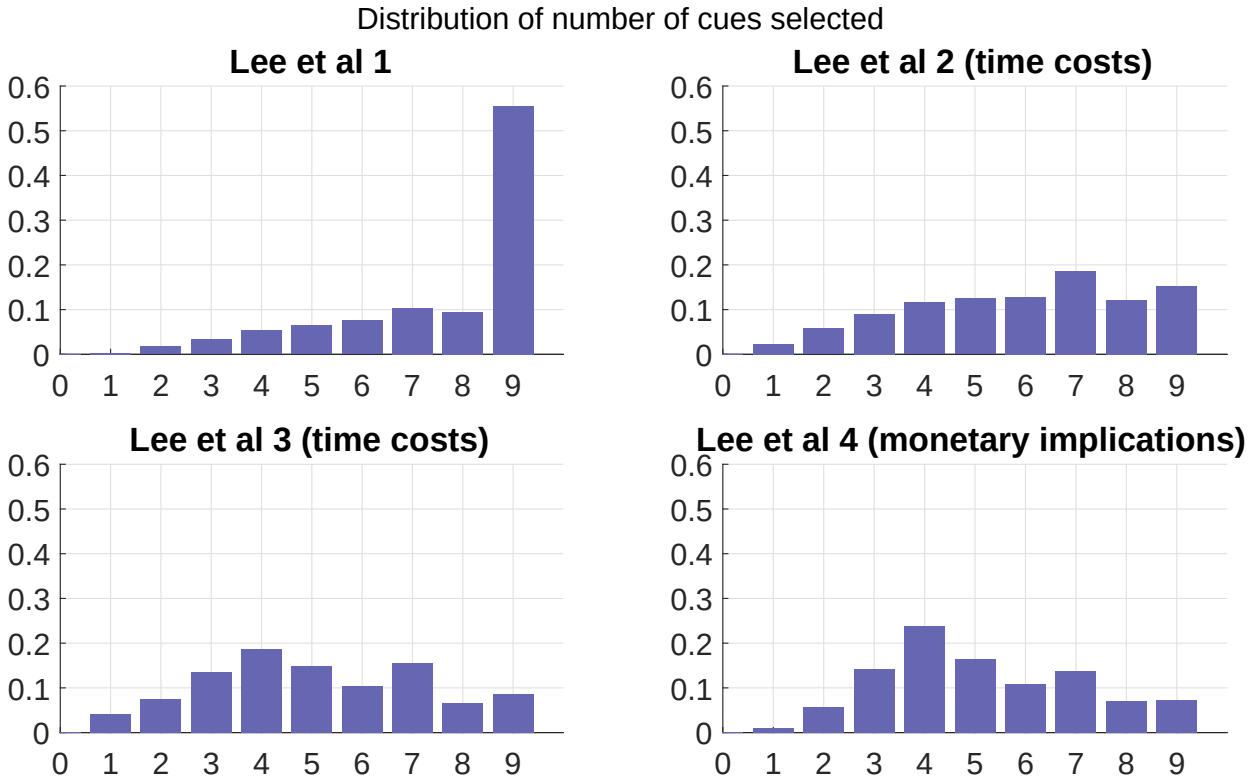

Figure 4.1: Distribution of number of cues selected (out of a maximum of 9, in each of the 4 experiments)

decision based on the first discriminating cue) was sufficient to achieve the same level of accuracy. The key measures of interest included the number of cues searched on each trial, how this changed along with the change in environment structure, and the resulting accuracy. We analyze 4 experiments from this study. The experiments are broadly similar, with the key differences being that experiment 4 has real monetary consequences for the participants, while experiments 2 and 3 have real time consequences (real delays in order to access more information), and experiment 1 has neither. Figure 4.1 summarizes the distribution of cue use in the 4 experiments. The data are skewed towards using almost all cues in the first experiment, but we observe that when real time and money come into play, cue use is more selective. There are significant individual differences, as well as changes from trial to trial, as can be observed from figure 4.2, which shows the cues used in each trial. The gray lines show individuals whereas the bold line shows the mean across all individuals. Apart from specifying the strategy based on PEC, the model specification is identical to that detailed in section 4.2. In this case, rather than set σ a priori, we infer it as a free parame-

Figure 4.2: Change in number of use on a trial by trial basis. The gray lines show individuals whereas the bold line shows the mean across all individuals.

ter for each individual, thus allowing for greater individual differences in the definition of what a strategy actually represents.

4.3 Model Specification

The choice made on trial t, denoted as y_t, is modeled as:

$$y_t \sim \text{Categorical}(p(y_{1,t}), ..., p(y_{n_O,t})) \tag{4.1}$$

$$p(y_{k,t}) = \Sigma_{i=1:N_h} \Sigma_{j=1:N_s} \left(p(y_{k,t}|s_{j,t}h_{i,t}) p(s_{j,t}|h_{i,t}) p(h_{i,t}) \right) \tag{4.2}$$

Probability of selecting a choice, conditional on a particular information search pattern and a par-

ticular strategy being used:

$$p(y_{k,t}|s_{j,t}h_{i,t}) \sim \text{ Decision rule for strategy } h_i \text{ (section 4.3.1 for details)} \tag{4.3}$$

Probability of selecting a particular information search pattern, conditional on a particular strategy being used:

$$p(s_{j,t}|h_{i,t}) \sim \text{ Information search and stop rule for strategy } h_i \text{ (section 4.3.1 for details)} \tag{4.4}$$

Probability of a particular strategy being used (generative component of the model) is modeled using a softmax action rule, governed by a free parameter θ at an individual level.

$$p(h_{m,t}) = \frac{e^{\theta\ q_{h_{m,t}}}}{\Sigma_{j=1:N_h}\left(e^{\theta\ q_{h_{j,t}}}\right)} \tag{4.5}$$

Calculation of q-values for each strategy, is as initially specified in equation 2.7.

$$q_{h_{m,t}} = p(h_{m,t-1}|x_{jk,t-1})\left((1-\eta)\ q_{h_{m,t-1}} + \eta\ r_{t-1}\right) + (1 - p(h_{m,t-1}|x_{jk,t-1}))\ q_{h_{m,t-1}} \tag{4.6}$$

Posterior probability of a particular strategy being used (discriminative component of the model, used only for generative modeling of the next trial):

$$p(h_{m,t-1}|x_{jk,t-1}) = \frac{p(y_{k,t-1}|s_{j,t-1}h_{m,t-1})p(s_{j,t-1}|h_{m,t-1})p(h_{m,t-1})}{\Sigma_{i=1:N_h}\left(p(y_{k,t-1}|s_{j,t-1}h_{i,t-1})p(s_{j,t-1}|h_{i,t-1})p(h_{i,t-1})\right)} \tag{4.7}$$

The parameter η reflects the learning rate, which is a free parameter for each individual.

$$\eta \sim \text{ Beta}(1,1) \tag{4.8}$$

The initial probability of selecting each strategy is modeled as a free parameter for each individual.

$$p_{h_{1:N_h,1}} \sim \text{ Dirichlet}\left(\left[\frac{1}{N_h}\right]_{1xN_h}\right) \tag{4.9}$$

$$q_{h_{m,1}} = p_{h_{m,1}} \tag{4.10}$$

4.3.1 Strategy specification

Traditional strategies - *Minimum*

To implement the model with traditional strategies, with inference based on whether or not the choice matches the choice predicted by the strategy, and whether the minimum cues required for implementing the strategy are acquired, we apply the below. The strategies are denoted *ttbm* and *wam* respectively. The search pattern corresponding to minimum cues required for implementing the two strategies are denoted as s_{ttb} and s_{wa} respectively. The number of unique search patterns that have at least the required information to implement these strategies are denoted as N_{ttbm} and N_{wam} respectively. To account for choice behavior that cannot be explained for any strategy, an application error rate ε is introduced as a free parameter at an individual level.

$$p(y_k|s_j,\, h_{ttbm}) \;= (1-\varepsilon) \text{ if } \mathrm{y_k} \text{ predicted by applying TTB; } \frac{\varepsilon}{\mathrm{n_O}} \text{ otherwise} \tag{4.11}$$

$$p(s_j|h_{ttbm}) \;= \frac{1}{N_{ttbm}} \;\forall\; \{s_j : s_j \subset s_{ttb}\};\; 0 \text{ otherwise} \tag{4.12}$$

$$p(y_k|s_j,\, h_{wam}) \;= (1-\varepsilon) \text{ if } \mathrm{y_k} \text{ predicted by applying WA; } \frac{\varepsilon}{\mathrm{n_O}} \text{ otherwise} \tag{4.13}$$

$$p(s_j|h_{wam}) \;= \frac{1}{N_{wam}} \;\forall\; \{s_j : s_j \subset s_{wa}\};\; 0 \text{ otherwise} \tag{4.14}$$

Traditional strategies - *Exact*

To implement the model with traditional strategies, with inference based on whether or not the choice matches the choice predicted by the strategy, and whether the cues acquired *exactly* match the cues required for implementing the strategy, and no more, we apply the below. The strategies are denoted *ttbe* and *wae* respectively. The search pattern corresponding to minimum cues required

are denoted as s_{ttb} and s_{wa} respectively. To account for choice behavior that cannot be explained for any strategy, an application error rate ε is introduced as a free parameter at an individual level.

$$p(y_k|s_j,\ h_{ttbe}) = (1-\varepsilon) \text{ if } y_k \text{ predicted by applying TTB; } \frac{\varepsilon}{n_O} \text{ otherwise} \tag{4.15}$$

$$p(s_j|h_{ttbe}) = 1 \text{ if } s_j = s_{ttb}; \quad 0 \text{ otherwise} \tag{4.16}$$

$$p(y_k|s_j,\ h_{wae}) = (1-\varepsilon) \text{ if } y_k \text{ predicted by applying WA; } \frac{\varepsilon}{n_O} \text{ otherwise} \tag{4.17}$$

$$p(s_j|h_{wae}) = 1 \text{ if } s_j = s_{wa}; \quad 0 \text{ otherwise} \tag{4.18}$$

Probabilistic modeling of strategies (multi-dimensional)

For the multidimensional case for the 2 experiments reported in this thesis, we implement the model with our novel probabilistically defined strategies. We apply the kernel based methods described in previous chapters. The strategies are denoted *ttbp* and *wap* respectively.

$$p(s_j|h_{ttbp}) = \frac{K(s_j,h_{ttbp})}{\Sigma_l K(s_l,h_{ttbp})} \tag{4.19}$$

$$p(s_j|h_{wap}) = \frac{K(s_j,h_{wap})}{\Sigma_l K(s_l,h_{wap})} \tag{4.20}$$

$$K(s_j,h_{ttbp}) = \exp\left(\frac{-\Sigma_{n=1:3}\left(F_n(s_j)-F_n(s_{ttb})\right)^2}{2\sigma^2_{ttbp}}\right) \tag{4.21}$$

$$K(s_j,h_{wap}) = \exp\left(\frac{-\Sigma_{n=1:3}\left(F_n(s_j)-F_n(s_{ttb})\right)^2}{2\sigma^2_{wap}}\right) \tag{4.22}$$

The search pattern features selected for basic implementation of the probabilistic model were:

F_1 : Proportion of cues searched

F_2 : Normalized variability in cues acquired across attributes (standard deviation of number of cues acquired for each attribute, normalized)
F_3 : Sensitivity to true validity (weighted average validity, normalized)

An application error rate is not required here, since this is a continuous probability distribution, and there is a non-zero probability for all search and choice patterns. The choice of σ governs the width of the distributions.

In addition, we define a third prototypical strategy *arbp*, where the kernel center is derived based on a sparse search pattern that picks a single, and different, attribute for each option, thus representing a highly explorative but sparse compensatory strategy, very different from prototypical rule-based WA and tallying strategies. This comes close to a guessing strategy, but is more informative than random guessing. It represents a low-effort, educated guess approach to information search.

A note on feature selection for probabilistic specification

The features selected here have meaningful interpretations for defining probabilistic versions of TTB and WA, in that the selected features for a prototypical application of these two strategies are quite different. Apart from psychological interpretation, the relative independence of these features is a useful characteristic. We calculate the selected features for all 4,096 possible cue search patterns possible within this paradigm, and run a principal component analysis on these features. Table 4.1 shows the relative independence of these features. To make a comparison, we swapped *sensitivity to validity*, for another feature, the *proportion of attributes selected*. The bottom half of the table shows that this set of features is not as independent.

Table 4.1: PCA analysis of 2 sets of selected features across all possible cue search patterns in the selected paradigm. The features selcted in the top half are relatively independent.

Feature	Latent 1	Latent 2	Latent 3
%cues	-0.004	1.0	0.0
var across attr	1.0	0.004	0.0
sensitivity to validity	0.0	0.0	1.0
Feature	Latent 1	Latent 2	Latent 3
%cues	0.35	0.68	-0.64
var across attr	-0.66	0.66	0.35
%attributes	0.67	0.30	0.68

Probabilistic modeling of strategies (single-dimensional)

For the single dimensional case for the experiments reported from Lee et al. (2014), we implement the model with our novel probabilistically defined strategies. Since on any trial, anywhere from 1 cue to 9 cues can be obtained, we define probabilistic TTB and WA strategies based on the example 1 in section 2.4.2 (example 1), using PEC as the focal measure to define kernels, where $PEC_{ttb} = 0$, and $PEC_{wap} = 1$. To accommodate cue use that is less than the minimum cues required to discern between options, that is fewer cues than what TTB requires, we define a guessing strategy with $PEC_{guessp} = -1$.

4.3.2 Alternate formulations of the learning rate

Assessment of environmental volatility and detection of environmental changes have been implicated in the modulation of learning rates (Pearson & Platt 2013; Behrens et al, 2007). We propose that the conflict between probabilities of using different strategies generated by the cognitive model can be interpreted as a proxy measure of volatility. We implement a version of the model that modulates the learning rate (η_t) on a trial-by-trial basis based on recent entropy. Higher entropy reflects greater uncertainty in the environment and hence increases learning rates. Entropy (H) is calculated

based on the strategy probabilities generated by the model:

$$H_t = -\frac{1}{\log(N_h)}\Sigma_{j=1:N_h}\big(p(s_j,t)\ log(p(s_j,t)\big) \tag{4.23}$$

$$\eta_t = \frac{1}{1+e^{-(k_H H_t+\eta_0)}} \tag{4.24}$$

4.3.3 A note on Bayesian implementation

We highlight that to the best of our knowledge, this is the first Bayesian implementation of an adaptive learning process over latent heuristic strategies. Even without considering the novel probabilistic definition of strategies, the implementation of traditional heuristics plus learning within a Bayesian model framework allows superior inferences, because the prior probability of using a particular heuristic obtained from the generative model acts as Bayesian regularization, in other words, it serves to inform the model and make better credit assignment predictions when a particular trial cannot discern between strategies. This aspect has not been addressed in previous non Bayesian applications of a learning model of heuristics. Bayesian modeling of individual heuristics and mixture models to determine proportions of heuristic strategies being used have been developed, but these models cannot incorporate a continuously changing distribution of strategy use over time, which is particularly informative in changing environments

4.4 Performance indicators

4.4.1 Overall model

We provide 2 baseline models based on traditional strategies, implemented with *minimum cue* and *exact cue* based inference methods. We highlight that these models were also implemented with the identical learning mechanism as for the probabilistic strategy based models. These baseline models

are compared against the probabilistic strategy models with fixed and entropy based learning rates. We compare the mean DIC (deviance information criteria) for each model, with a lower score being better.

4.4.2 Choice accuracy

Brier Scores

Since the models primary return a posterior predictive distribution over multiple categorical outcomes (1 out of n_O choices), we use a multi-category Brier score to assess the performance of the models. This is similar to an error measure, and in fact defaults to the mean square error in the case of a binary classification problem. We calculate the Brier score for each model at the level of each individual, as given by equation 4.25, where c is each choice option, t is each trial, and o_{tc} is the actual behavior on trial t, which is coded as 1 if option c is selected on trial t and 0 otherwise. The cognitive model provides a distribution over responses, and p_{tc} refers to the probability of choice c being selected on trial t, based on the posterior predictive distribution. A lower Brier score is better, with a perfect prediction resulting in a score of zero. For the tasks with binary choice (single dimensional, Lee et al. (2014)), the Brier score is effectively the same as a mean square error for this task.

$$BS = \frac{1}{N_t}\Sigma_{t=1:N_t}\ \Sigma_{c=1:n_O}\left(p_{tc} - o_{tc}\right)^2 \tag{4.25}$$

Modal choice accuracies

In addition, we also calculate the accuracy level of predictions by measuring the proportion of choices for which the *modal* response of the model matches the true outcome. This provides a

better intuitive sense of the model performance, but we highlight that the Brier score is a more accurate and calibrated measure. The choice accuracy is reported as Acc.

Description versus Prediction

The models for the two experiments reported in the previous chapter are implemented so that data for the first 90 trials is provided to each model, and the model provides a posterior predictive distribution for all 120 trials. The performance of the model on the first 90 trials provides an assessment of descriptive performance (or model fit), and on the last 30 trials it provides a predictive performance (or generalizability). For the models for experiments reported from Lee et al. (2014), data from the first 100 trials is provided to the models and performance on the first 100 trials serves as a measure of descriptive fit, whereas that on the last 100 trials provides an assessment of predictive performance (generalizability). The Brier scores for these two categories will be coded as BS_d (description) and BS_p (prediction) respectively. The accuracies are similarly reported as Acc_d and Acc_p.

4.4.3 Search pattern feature accuracy

The cognitive model provides a probability of each strategy being used on each trial. Each strategy has a specific search pattern associated with it (in traditional strategies), or a probability distribution over features of the search patterns (in probabilistic strategies). We assess the posterior predictive distribution of the search pattern *features* incorporated within the probabilistic heuristics. Since each feature varies on a scale of 0 to 1, we can calculate the root mean square error (Err_d and Err_p) for the model predicted versus actual search pattern features. This is similarly assessed as descriptive and predictive. Features are a reasonable way to measure search pattern accuracy in experiments with complex designs where there are over 4,000 unique search patterns. For the multi-dimensional modeling, we measure the error based on the 3 features selected in section 4.3.1.

For the single-dimensional measures, we calculate the root mean square error in number of cues used.

4.5 Model Comparison

Tables 4.2 and 4.3 show a comparison across the models tested for the 6 different experimental datasets. Note that the baseline models themselves make far superior predictions than standalone heuristic strategies, or a mixture model of strategies that does not incorporate learning. The novel approach to heuristic strategies does even better on all performance indicators.

4.5.1 Experiments 1 and 2

See table 4.2, and figures 4.3 to 4.6. The figures provide a breakdown of the improvement in both descriptive and predictive choice and search errors by individual participants. The participants are arranged in order of decreasing error based on the baseline model. The green bars show an improvement (reduction) in the error using the probabilistic models and the red bars show the deterioration (increase) in error using the probabilistic model. For evaluation of the novel probabilistic models, more green is better! The error in search features is broken down into 6 key features of search patterns, including the proportion of cues used, the proportion of cues used within selected attributes, the variability of cue use across attributes and across options, the proportion of most valid cues selected, and the overall sensitivity to validity. For experiments 1 and 2, the relative performance of the probabilistic model is more impressive in experiment 2 (Mistry et al), and where we have seen higher level of adaptivity. The difference in choice accuracy predictions when data is not supplied to the models is especially significant. Note that the reduction in error in search accuracy predictions in the range of about 0.05 to 0.1 is not trivial. To put this reduction in perspective, misclassifying a TTB strategy for WA or vice versa would yield an average error of

about 0.5 (this can vary a bit depending on the exact cue configurations), and these are 2 strategies that are considered to be at opposite ends of the spectrum. One of the key things to observe is that the traditional model can perform well for some participants, but has extremely high errors for others. The probabilistic models while maintaining a better overall error and accuracy, also provide less variability in how accurately each individual is described and predictive. This points to more robust inferences about individual differences in the cognitive process.

Table 4.2: Model comparison summary - Experiments 1 and 2. The best performing model for each measure is highlighted in bold

	Measure	Baseline *Min*	Baseline *Exact*	Probabilistic *Fixed*	Probabilistic *Entropy*
Experiment 1 (Mistry et al)	*DIC*	152	146	135	**134**
	BS_d	0.46	0.44	**0.41**	**0.41**
	BS_p	0.47	0.47	**0.42**	**0.41**
	Acc_d	69%	72%	**73**%	72%
	Acc_p	69%	70%	**73**%	72%
	Err_d	0.32	0.36	0.28	**0.27**
	Err_p	0.32	0.35	**0.28**	**0.28**
Experiment 2 (Mistry et al)	*DIC*	129	129	**106**	110
	BS_d	0.37	0.37	**0.30**	0.31
	BS_p	0.34	0.34	**0.21**	**0.21**
	Acc_d	78%	80%	**81**%	**81**%
	Acc_p	75%	82%	**90**%	**90**%
	Err_d	0.40	0.41	**0.31**	0.32
	Err_p	0.40	0.40	**0.34**	**0.34**

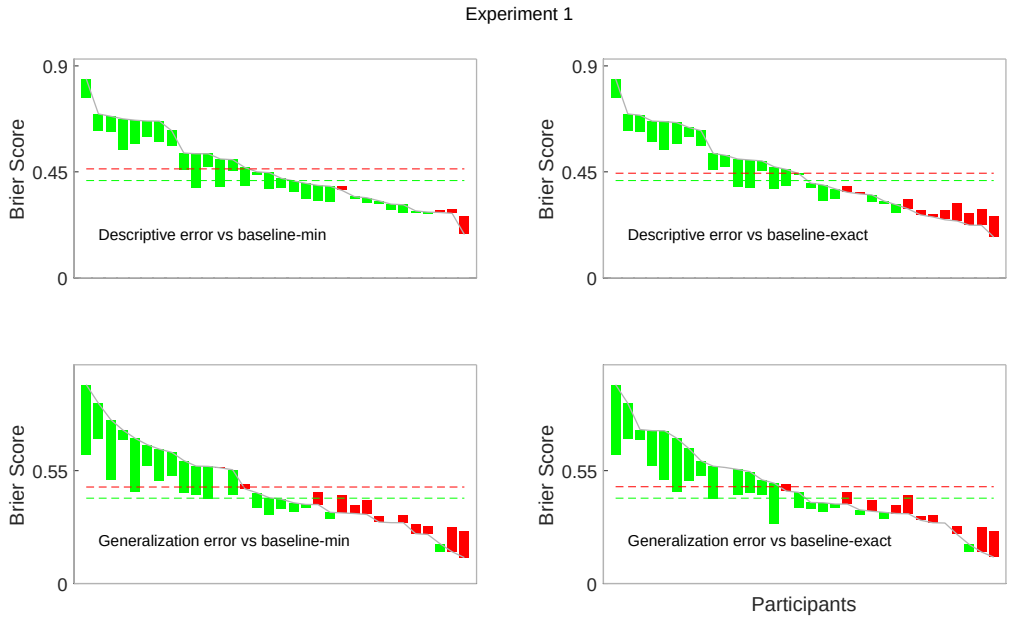

Figure 4.3: Experiment 1: Brier scores for choice with participants arranges by reducing model score (improving performance) of the baseline model (gray line). Green bars reflect the magnitude of improved performance and red bars, the reduced performance, of the probabilistic model.

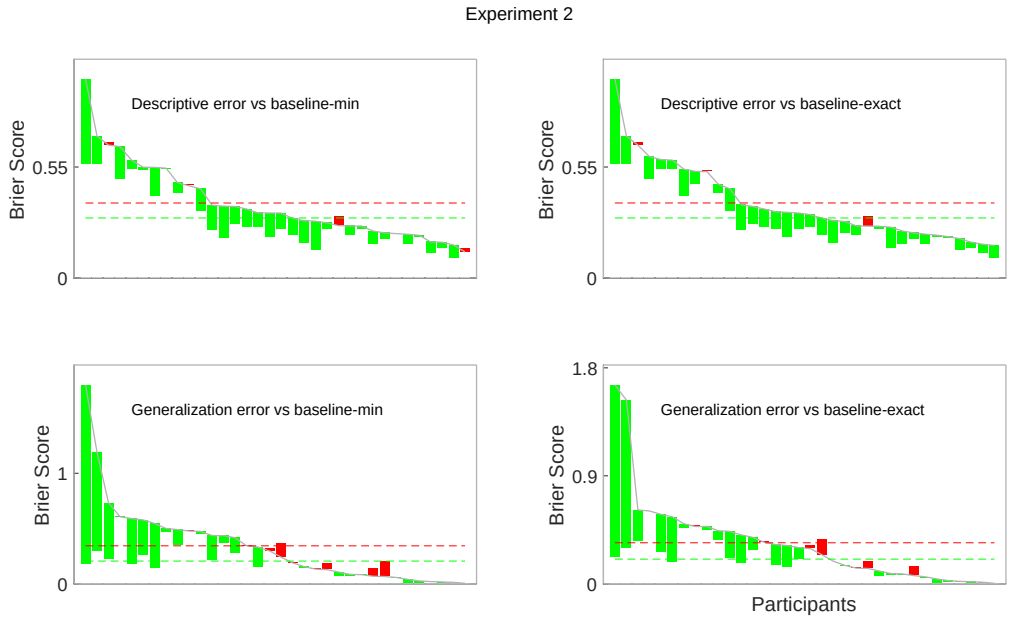

Figure 4.4: Experiment 2: Brier scores for choice with participants arranges by reducing model score (improving performance) of the baseline model (gray line). Green bars reflect the magnitude of improved performance and red bars, the reduced performance, of the probabilistic model.

Figure 4.5: Experiment 1: RMSE for search with participants arranges by reducing model score (improving performance) of the baseline model (gray line). Green bars reflect the magnitude of improved performance and red bars, the reduced performance, of the probabilistic model.

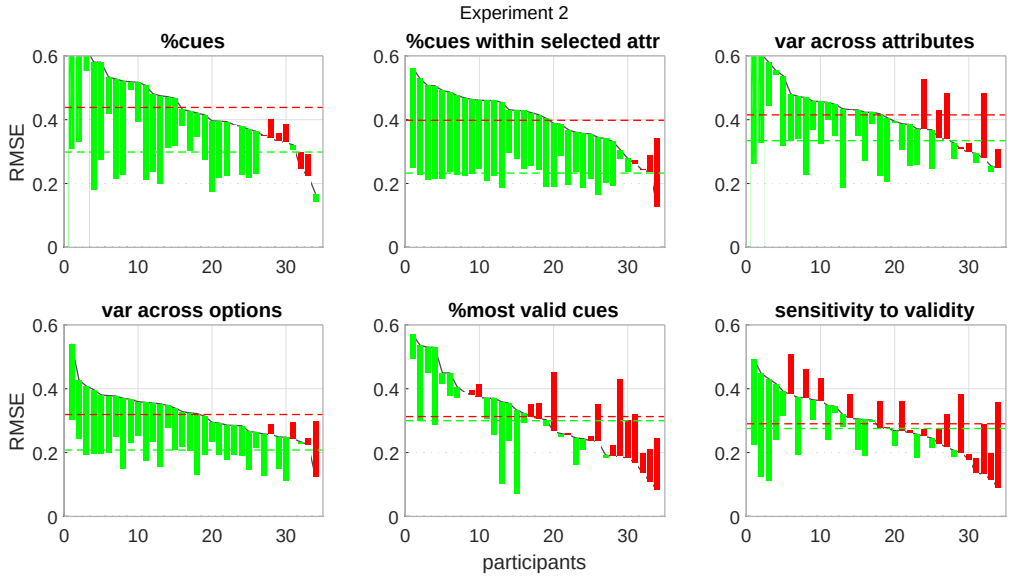

Figure 4.6: Experiment 2: RMSE for search with participants arranges by reducing model score (improving performance) of the baseline model (gray line). Green bars reflect the magnitude of improved performance and red bars, the reduced performance, of the probabilistic model.

4.5.2 Experiments from Lee et al.

See table 4.3, and figures 4.7 to 4.10. The figures provide a breakdown of the improvement in both descriptive and predictive choice and search errors by individual participants. The partici-

pants are arrange in order of decreasing error based on the baseline model. The green bars show an improvement (reduction) in the error using the probabilistic models and the red bars show the deterioration (increase) in error using the probabilistic model. The error in search features is based on the number of cues searched. For Lee 1, the probabilistic model provides higher generalizability errors than the traditional model in terms of the search errors. Note hat the search patterns in this experiment were almost always searching for all cues (84% of cues selected across al choices made). A simple WA model without any adaptivity explains behavior really well in this task. Note that the probabilistic model still provides a lower error in terms of the Brier scores obtained on choice accuracy, reflecting the higher uncertainty in the traditional model. The key observation of interest however is that when real aspects influencing decisions such as real time delays and monetary incentives come into play, people change their behavior and become more adaptive. The probabilistic model performs significantly better across almost all individuals in these conditions, as can be seen from the figures for both search and choice error.

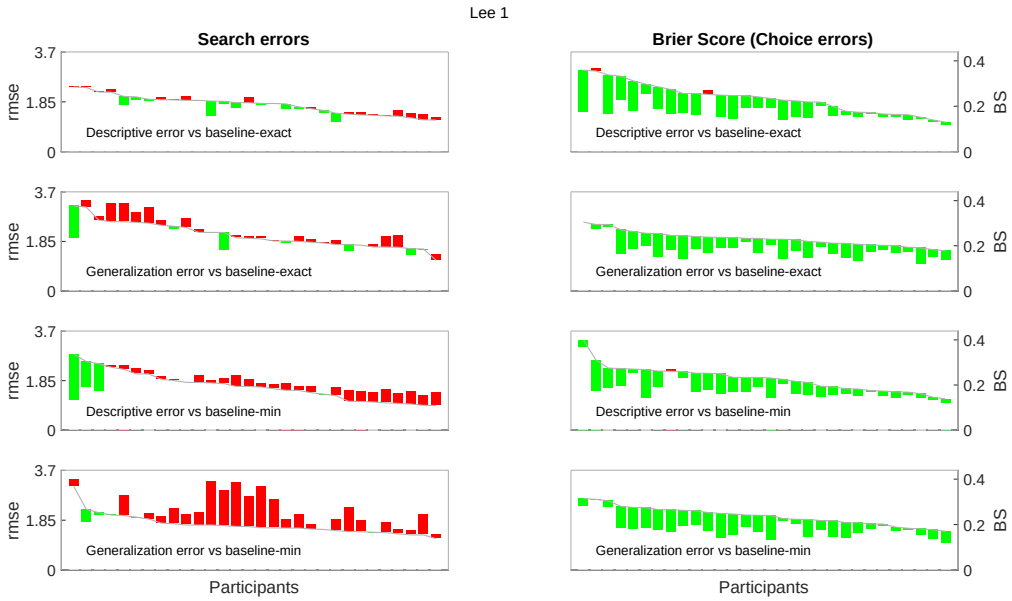

Figure 4.7: Lee 1: RMSE for search and Brier scores for choice errors with participants arranges by reducing model score (improving performance) of the baseline model (gray line). Green bars reflect the magnitude of improved performance and red bars, the reduced performance, of the probabilistic model.

Table 4.3: Model comparison summary - The best performing model for each measure is highlighted in bold.

Experiment	Measure	Baseline *Min*	Baseline *Exact*	Probabilistic *Fixed*
Lee et al 1	*DIC*	327	356	**305**
	BS_d	0.23	0.24	**0.18**
	BS_p	0.24	0.23	**0.18**
	Choice Acc_d	82%	90%	**92%**
	Choice Acc_p	78%	83%	**90%**
	Search Err_d	**1.63**	1.74	1.72
	Search Err_p	**1.68**	2.08	2.17
Lee et al 2	*DIC*	470	505	**372**
	BS_d	0.24	0.26	**0.16**
	BS_p	0.28	0.27	**0.20**
	Choice Acc_d	79%	84%	**88%**
	Choice Acc_p	76%	**89%**	85%
	Search Err_d	1.99	1.80	**1.66**
	Search Err_p	1.99	1.75	**1.71**
Lee et al 3	*DIC*	533	548	**467**
	BS_d	0.36	0.36	**0.27**
	BS_p	0.34	0.35	**0.29**
	Choice Acc_d	68%	68%	**78%**
	Choice Acc_p	71%	72%	**74%**
	Search Err_d	2.03	1.89	**1.64**
	Search Err_p	2.45	1.88	**1.66**
Lee et al 4	*DIC*	544	569	**418**
	BS_d	0.39	0.37	**0.26**
	BS_p	0.37	0.38	**0.30**
	Choice Acc_d	64%	68%	**76%**
	Choice Acc_p	68%	68%	**73%**
	Search Err_d	1.75	1.64	**1.41**
	Search Err_p	2.39	1.75	**1.59**

.

Figure 4.8: Lee 2: RMSE for search and Brier scores for choice errors with participants arranges by reducing model score (improving performance) of the baseline model (gray line). Green bars reflect the magnitude of improved performance and red bars, the reduced performance, of the probabilistic model.

Figure 4.9: Lee 3: RMSE for search and Brier scores for choice errors with participants arranges by reducing model score (improving performance) of the baseline model (gray line). Green bars reflect the magnitude of improved performance and red bars, the reduced performance, of the probabilistic model.

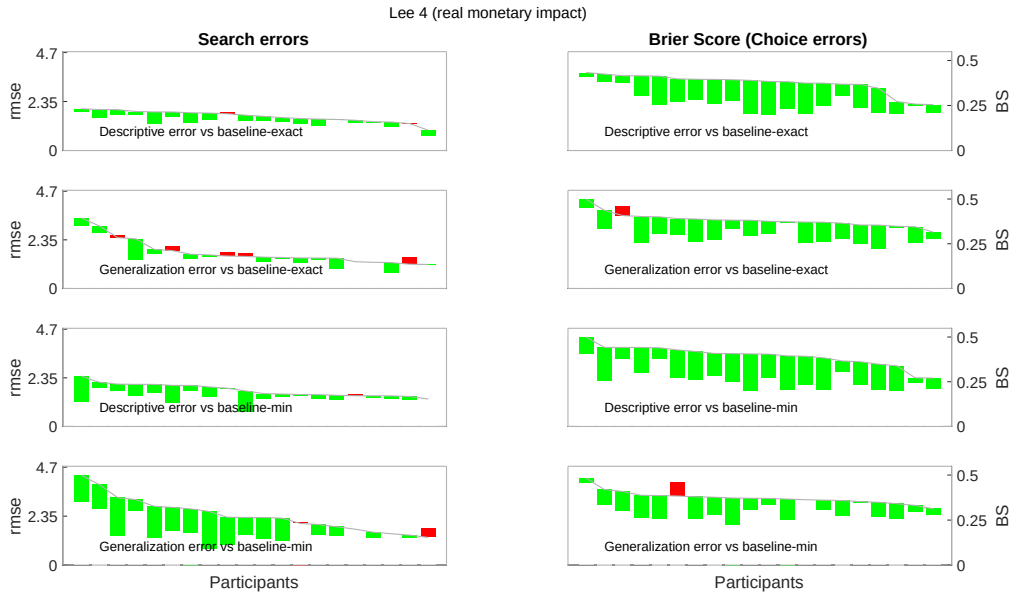

Figure 4.10: Lee 4: RMSE for search and Brier scores for choice errors with participants arranges by reducing model score (improving performance) of the baseline model (gray line). Green bars reflect the magnitude of improved performance and red bars, the reduced performance, of the probabilistic model.

4.6 Model Inferences

4.6.1 Experiment 1 and 2

Initial preferences

Figure 4.11 shows the inferred distribution of *initial* strategy preferences for each participant (data combined across experiments 1 and 2), for the 3 probabilistic strategies in the consideration set, $ttbp$, wap, and $arbp$. The last strategy represents a sparse-compensatory-exploratory strategy. This shows significant individual differences in initial preferences.

Figure 4.12 classifies participants according to the most dominant initial preference, and shows how their performance and search pattern over *all the trials* varies as a function of the initial preference. This shows that the initial preference do play a strong role even in the presence of learning and adaptive selection of strategies. Participants with strong initial preferences towards

an exploratory *arbp* strategy show lower sensitivity to validity, and lower performance, in this environment.

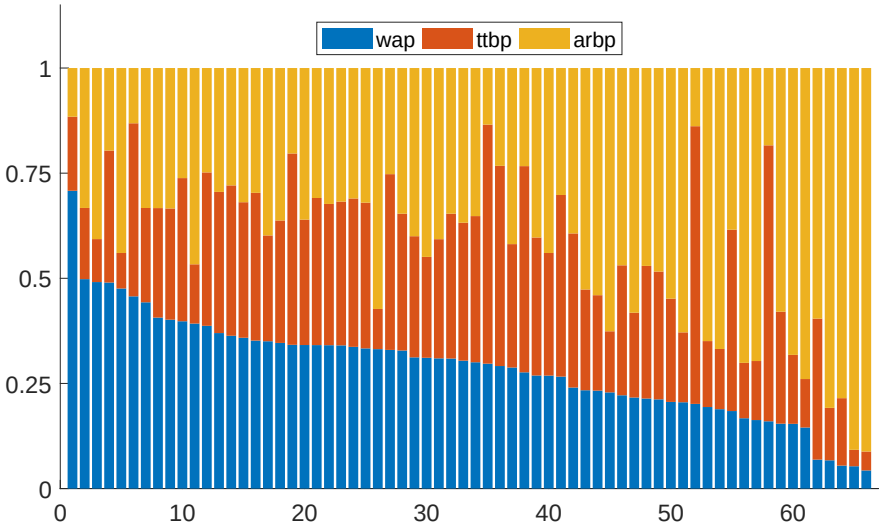

Figure 4.11: Distribution of initial strategy preference across participants

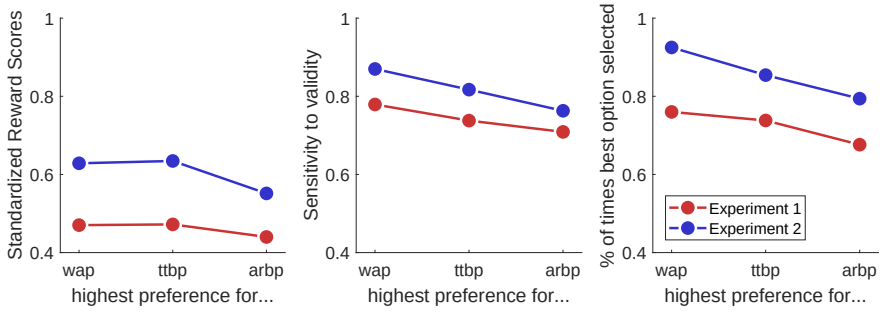

Figure 4.12: Relation between initial preference and performance

Strategy Shifts

Figures 4.13 and 4.14 show the mean inferred posterior predictive probability of each strategy being selected on each trial. The light lines show the movement for individuals and the bold lines plot the mean for individuals in each condition. The blue and yellow backdrops represent C-blocks and N-blocks respectively. The blue and red lines represent probabilistic WA and TTB heuristics. Figure 4.15 shows similar inferences made for experiment 2 using the original traditional heuristics. It can be seen that these heuristic models cannot capture continuity in the learning process.

Although the quick clean shifts may look efficient, these do not represent human learning, and these models do not provide as good a descriptive or predictive capability. ***Most importantly, the almost constant strategy use within blocks cannot explain the significant within block changes observed in behavioral measures, such as the significant drop in cue use and increasing accuracy between the first and second half of N-blocks. On the other hand, the descriptive account for such changes is quite apparent if we observe the curves underlying changes in strategy use in figure 4.14.*** The reason for the inferences appearing as they do in figure 4.15 is that for a huge majority of the trial the traditional models make no meaningful inference about the underlying process, but a small handful of trials

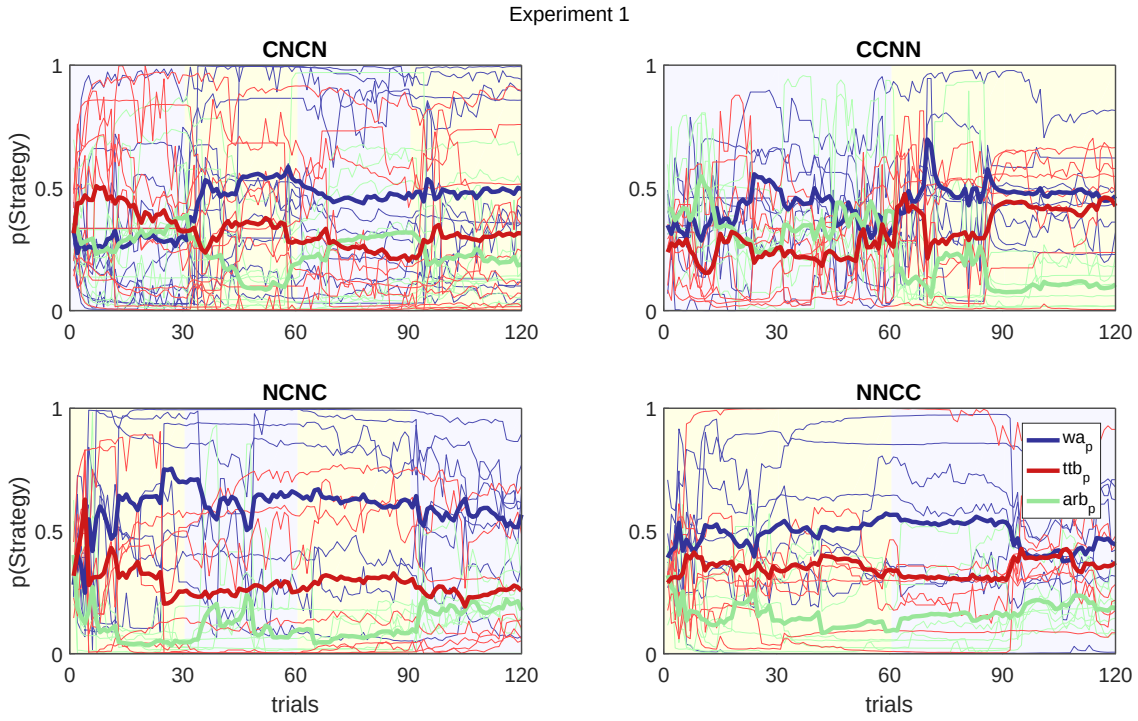

Figure 4.13: Experiment 1: Inferred probability of strategy use based on probabilistic strategy learning model

What are people learning?

In experiment 1, higher inferred learning rate does not seem to significantly affect observed behavior. In experiment 2 however, higher inferred learning rate is correlated with:

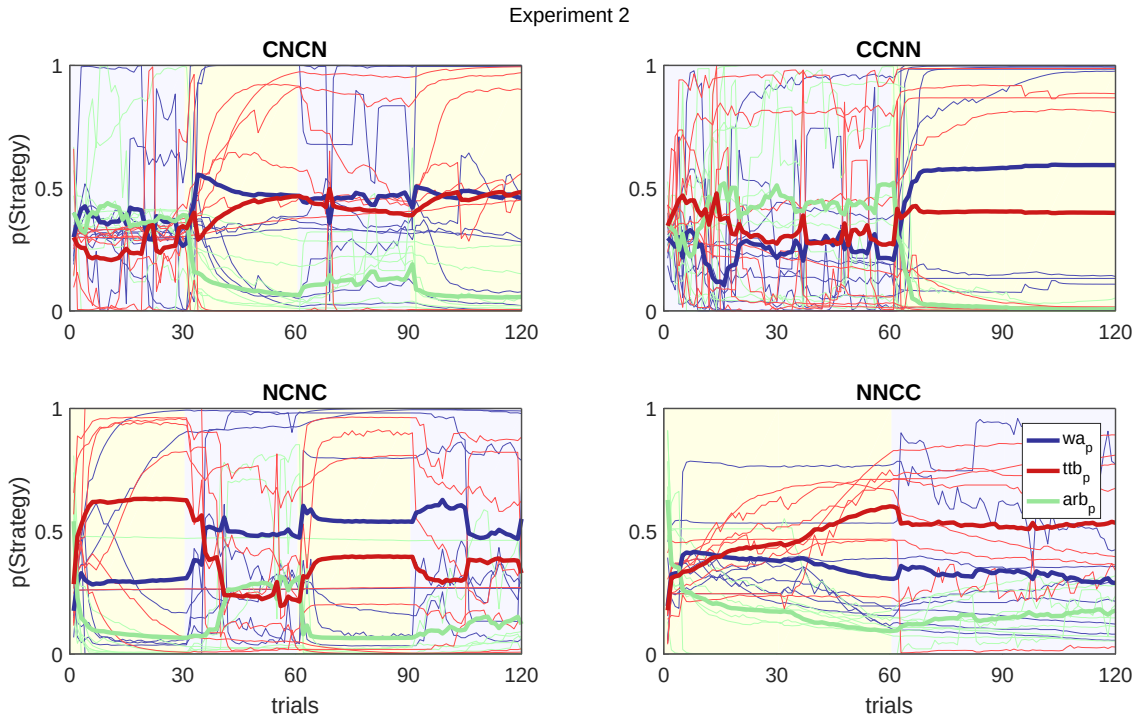

Figure 4.14: Experiment 2: Inferred probability of strategy use based on probabilistic strategy learning model

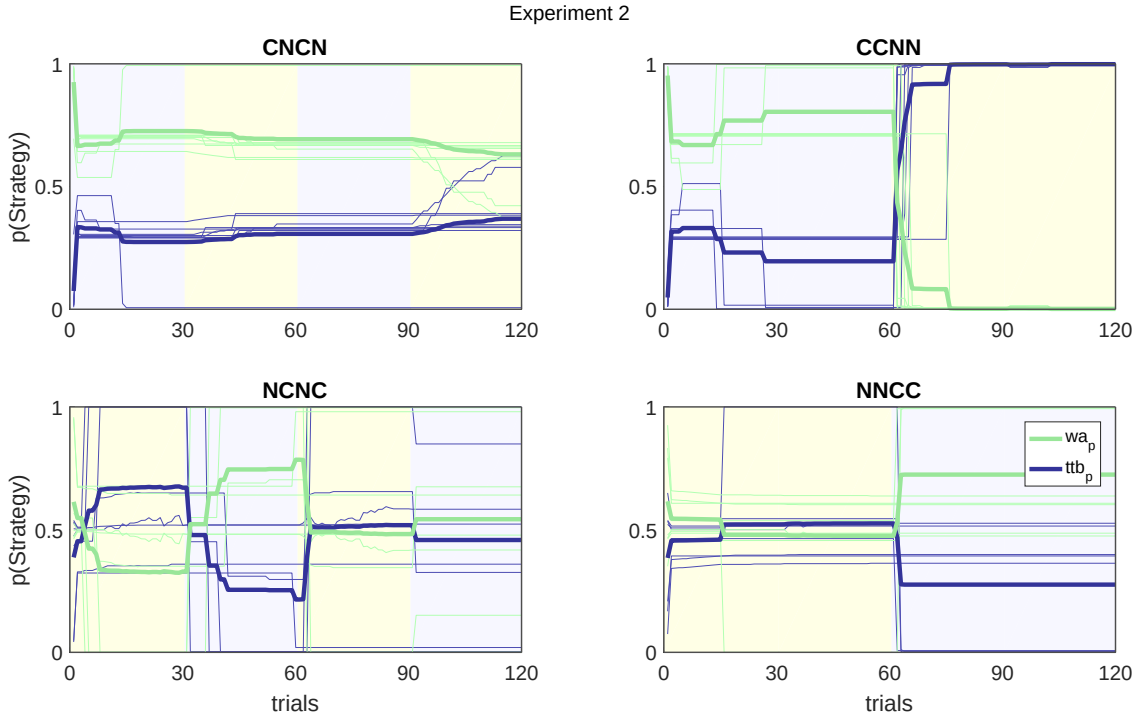

Figure 4.15: Experiment 2: Inferred probability of strategy use based on traditional strategy learning model

(a) reducing depth of information search (cues and attributes), but this effect is significant only under alternating (CNCN and NCNC) sequences, and not consecutive (CCNN and NNCC) sequences; and

(b) increasing rate of selecting the best possible option and increasing standardized reward scores, but these two effects are significant only under consecutive (CCNN and NNCC) sequences, and not alternating (CNCN and NCNC) sequences.

Thus, in alternating sequences (more frequent changes), learning seems to be focused on improving search efficiency, whereas in consecutive sequences (less frequent changes), learning seems to be focused on improving overall reward outcome.

Entropy-driven learning

We implemented the entropy driven learning rate model, for which the learning rate at any time is given by equation 4.26, where η_0 is base rate, H_t is the latent psychological entropy based on the level of uncertainty in strategy use, and k_H is the sensitivity to entropy. A high value of k_H will make learning far more sensitive to uncertainty in strategy selection.

$$\eta_t = \frac{1}{1 + e^{-(k_H H_t + \eta_0)}} \tag{4.26}$$

Figure 4.16 shows a large cluster of participants in both experiments with similar values of the base rate and sensitivity to entropy, leading to the inference that experimental factors rather than cognitive traits may be playing a greater role in the observed individual differences in entropy driven learning between the two experiments.

In figure 4.17, we see that entropy, and hence learning rate is generally higher in C-blocks compared to N-blocks. This differences is especially higher for participants who started in any C-block condition. This is not because of greater uncertainty at the start of the experiment, because we do not see the reverse effect for participants who started with an N-block. This shows how experi-

mental manipulations can be captured in the resulting cognitive process.

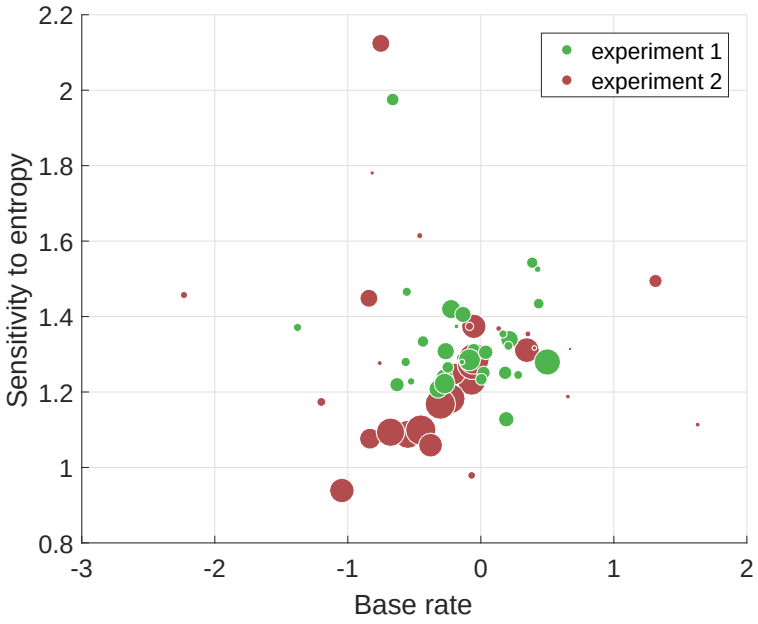

Figure 4.16: Experimental factors rather than individual differences drive entropy-driven learning

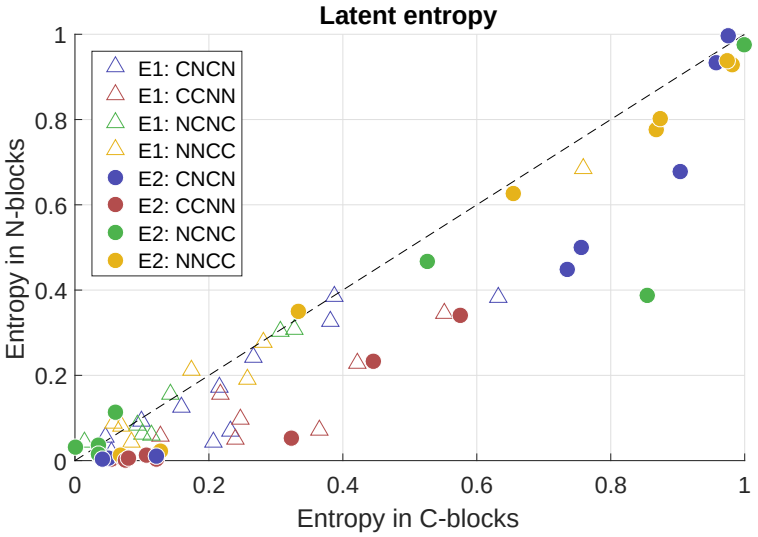

Figure 4.17: Higher entropy (uncertainty) in C-Blocks compared to N-blocks, by condition. E1 and E2 refer to experiment and experiment 2 respectively.

4.6.2 Experiments from Lee et al

Figure 4.18 shows the inferred standard deviation of the probabilistic TTB and WA strategies inferred to be used, across th 4 different experimental tasks. We can see that the introduction of

real time delays and monetary values results in a more refined (narrower) TTB prototype, but a wider WA prototype representation. This represents greater exploration of search strategies similar to WA, but focused application of TTB when that is merited.

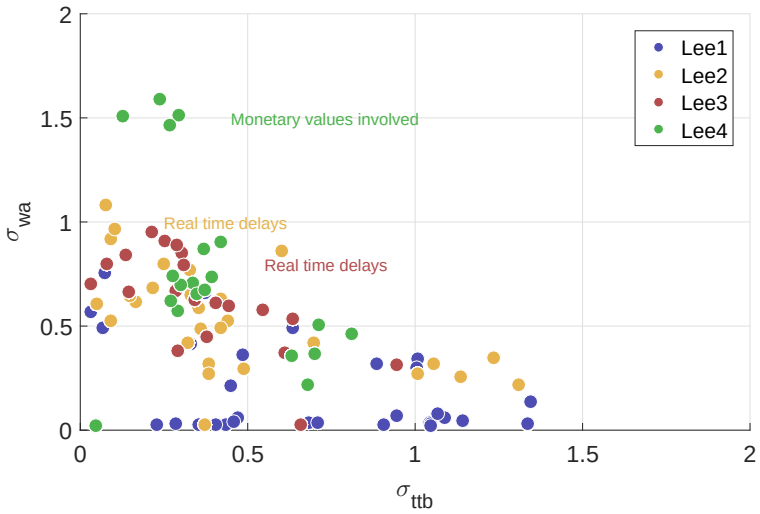

Figure 4.18: Inferred strategy precision

4.7 Modeling probabilistic heuristics to detect kernel centers

In the previous implementation, we decided the strategy kernel centers a prior, by representing a prototypical TTB and WA strategy as the kernel center. Here we test a different approach, focusing specifically on the information search patterns, and ignoring the choice selection, which represents a dependent but separate aspect of the choice process. In fact, in many cases in multi-attribute decision making, conditional on a particular information search pattern, the choice selection becomes almost trivial. Accordingly we implement a Bayesian search pattern detection model that *also incorporates a learning mechanism*. This is a very important distinction from other methods that might do a cluster analysis or some other pattern recognition algorithm that does not take into account the sequential nature of changes in the search pattern. The behavioral data in this case, are the information search features F representing the information search pattern on trial t. The search

For a particular individual, the model is specified as follows:

$$F_{it} \sim Gaussian(\bar{f}_{ih}, \sigma) \quad i \in \text{selected features} \tag{4.27}$$

Here, the subscript h refers to a particular latent strategy selected by an individual on that trial. The strategy used is modeled based on equations 4.5 and 4.6. Rather than specifying $\bar{f}_{ih}$, we infer these, allowing for individual differences across individuals. For identifiability, and to provide some reference points, we constrain the priors on these kernel feature values. We put in the constraints that strategy 1 has a prototypical representation that selects a lower proportion of cues, has higher sensitivity to validity, and has a higher variability across attributes, compared to strategy 2. In the case of traditional heuristics, TTB and WA would satisfy the roles of strategy 1 and 2, but rather than make this a strict enforcement, we explore the most likely prototypical kernel centers for each participant, and compare them to the prototypical TTB and WA values. The model works by searching for the prototypes that best explain trial to trial variability in terms of strategy switching between these two inferred strategies. The inferred prototypes are plotted in figures 4.19 and 4.20, for experiments 1 and 2. Each participant is represented by 2 dots, the blue dot representing strategy 1 and the green one representing strategy 2. The prototype representation for TTB and WA are also provided for comparison. Longer lines represent the focal dimensions that people are switching between.

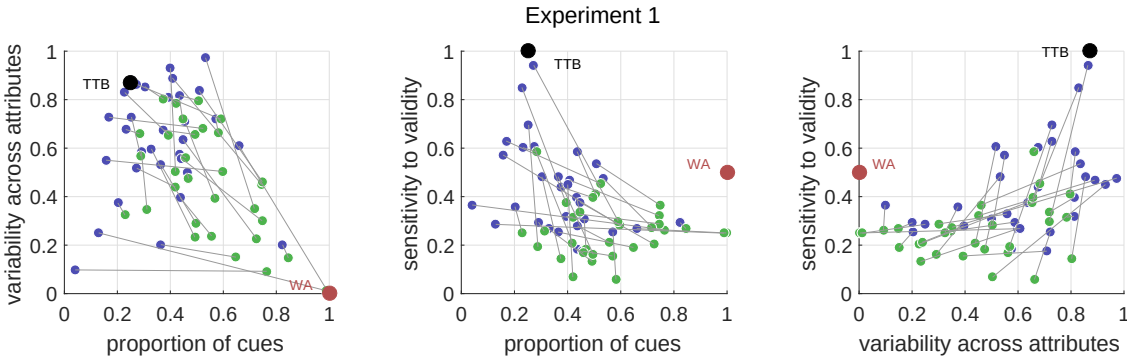

Figure 4.19: Experiment 1: Inferred kernel centers for each participant, representing two latent information search strategies identified. Individual are assumed to switch between the green and blue strategies.

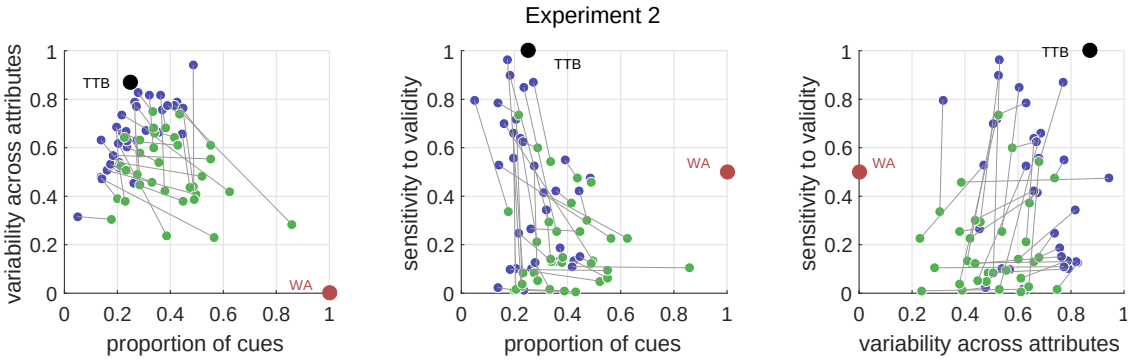

Figure 4.20: Experiment 2: Inferred kernel centers for each participant, representing two latent information search strategies identified. Individual are assumed to switch between the green and blue strategies.

4.8 Conclusion

In this chapter:

(a) We have shown the first (to the best of our knowledge) Bayesian implementation of adaptive strategy learning and switching models of multi-attribute heuristics.

(b) We have implemented and compared traditionally defined heuristics with the novel probabilistic approach to defining heuristic strategies, and shown that this novel approach generally provides more accurate descriptive and predictive capabilities.

(c) We have shown how incorporating these probabilistic heuristic strategies within a model of learning can enable improved understanding of the latent cognitive processes involved in how people adapt.

(d) We have demonstrated a model to cluster search patterns used by individuals within the framework of learning and cluster-switching, and shown how the resulting kernel centers occupy a wide and continuous range between the extreme prototypes of TTB and WA search patterns.

Chapter 5

A priori evaluation of heuristic strategies

5.1 A measurement framework for heuristic strategies

As a systematic approach to researcher decisions on what consideration set of heuristics and inference rules to use when approaching the cognitive modeling of adaptive strategy switching, we propose a measurement framework. This framework has been presented in Mistry and Trueblood (2018). This formalizes the idea of well-posedness of inference problems of strategy selection and learning based on mathematical theory. We specify the operator A (the approximate cognitive process) that maps the space of cognitive heuristics H to the space of observed behaviors X. In the most extreme case, the operator A^{-1} is strictly a well-posed inverse operator, although that is neither a realistic not a necessarily desirable requirement. Rather, we define what it would take for a set of heuristics to satisfy the strict inverse requirements. Based on this, we develop a scoring mechanism from 0 to 1, where a score of 1 essentially implies a strict inverse operation. We propose that a consideration set of heuristics should try to optimize this score in order to design the best descriptive and predictive model of behavior. One of the key considerations here is the latent identification of strategy use required to construct robust models of learning and adaptive strategy switching.

5.1.1 Notation and assumptions

In the realm of traditionally defined heuristics, we introduce an application error ε_i, which is the probability of making an error conditional on using the heuristic h_i. An error results in selecting a particular cue search and decision pattern that is not compatible with any of the heuristics. If an error is made, the probability of any incompatible behavior is uniformly distributed over all possible incompatible behaviors.

$$p(x_{kj}|h_i) = (1-\varepsilon_i I_{kj})p(y_k|s_j,h_i)p(s_j|h_i) \; + \; \frac{\varepsilon_i \, I_{kj}}{\Sigma_{k'}\Sigma_{j'} \, I_{k'j'}} \tag{5.1}$$

An alternate way to accommodate incompatible behavioral patterns using traditional heuristics is a guessing strategy (G), which can be defined as having uniform probability over all patterns instead of just over incompatible behavioral patterns:

$$p(x_{kj}|G) = \frac{1}{N_s\, n_O} \tag{5.2}$$

We denote $p(G) = g$ and assume equal prior probability of each heuristic, $p(h_i) = \frac{1-g}{N_h}$, and equal application error $(\varepsilon_i = \varepsilon)$. Note that these assumptions are not necessary, but enable more compact and intuitive expressions of the properties, and we will continue to use these assumptions in the formal specification. The overall probability of observing a particular behavioral pattern x_{kj} is given by:

$$\begin{aligned} p(x_{kj}) = \left[\frac{(1-\varepsilon I_{kj})(1-g)}{N_h}\, \Sigma_i \left\{ p(y_k|s_j,h_i)p(s_j|h_i) \right\}\right] \\ + \left[\frac{\varepsilon(1-g)I_{kj}}{\Sigma_{k'}\Sigma_{j'}\, I_{k'j'}} + \frac{g}{N_s\, n_O}\right] \end{aligned} \tag{5.3}$$

In equation 5.3, the terms in the first square brackets give the probability of observed behavior based on error-free heuristics, whereas the terms in the second brackets give the probability on account of errors or guessing. For ease of reference, we group the first $p(x'_{kj})$, and second $p(x''_{kj})$ set of terms:

$$p(x_{kj}) = p(x'_{kj}) + p(x''_{kj}) \tag{5.4}$$

5.1.2 Defining a well-posed strategy inference problem

Using Kabanikhin (2008), we define an *inverse* strategy selection problem as ill-posed, if any one of the conditions below are not met:

1. Existence of a solution: $\exists\, (h \in H)\ \forall\, (x \in X)$. This requires any possible observable behavior

(x) should be compatible and explained by at least one of the heuristics ($h \in H$).

2. Uniqueness of the solution: $\exists\, A^{-1} : X \rightarrow H$. This requires that each observable behavior (x) should be compatible with only one heuristic (h) within the set H.
3. Stability of the solution: $O(\delta_h) \approx O(\delta_x)$ *"Arbitrarily small errors in the measurement data",* δ_x should not *"lead to indefinitely large errors in the solutions"* δ_h. This requires that small changes in observed behavior (δ_x) should not result in significant changes to the inferred heuristic (h).

This is a hard set of constraints, and would be impossible, for any set of cognitive heuristics to satisfy, and A^{-1} is never expected to be a strict well-posed inverse operation. Instead, we treat these as properties that a consideration set of cognitive heuristics should try to optimize. We specify a formal but relaxed interpretation of these criteria for how a set of heuristic strategies should be evaluated:

1. *Existence property (P_1):* We propose a measure of the average probability of observing a behavioral pattern based on error-free application of heuristics, compared to the overall probability of observing it, including on account of errors or guessing, integrated over all possible behaviors and over all possible heuristics in the consideration set. This measure will vary from 0 to 1 with higher values desirable, and a value of 1 implying strict compliance with the existence property:

$$P_1 = \frac{1}{N_s\, n_O} \Sigma_k \Sigma_j \left[\frac{p(x'_{kj})}{p(x'_{kj}) + p(x''_{kj})} \right] \tag{5.5}$$

Intuitively, this measures what proportion of possible behavior can be explained by the consideration set of heuristics without using an error-based explanation. A high value of P_1 protects against the inference problem of *no matches*.

2. *Uniqueness property (P_2):* We propose using the generalized Jensen-Shannon divergence (Lin (1991)) for multiple distributions to measure uniqueness of a set of heuristics. This divergence is bounded by $[0, log(1+N_h)]$, so we adapt this to the range $[0,1]$. In equation 5.6, $\mathbb{H}_n(p_n) =$

$-\Sigma_n \left[p_n \, log(p_n)\right]$, refers to the Shannon entropy.

$$P_2 = \frac{1}{log(1+N_h)} \left\{ \mathbb{H}_{kj} \left(\Sigma_i \left[p(h_i) p(x_{kj}|h_i) + \frac{g}{N_s \, n_O} \right] \right) \right. \\ \left. -\Sigma_i \left[p(h_i) \mathbb{H}_{kj} \left(p(x_{kj}|h_i) \right) \right] - g \mathbb{H}_{kj} \left(\frac{1}{N_s \, n_O} \right) \right\} \tag{5.6}$$

The advantage over commonly used measures of divergence such as the Kullback Leibler divergence, is that $P2$ is smoothed, symmetrical, and can be simultaneously applied over multiple distributions. Further, it also linked directly to both the lower and upper bound on the Bayes probability of error (*BPE*, the lowest possible error rate of a classifier). In our context, this provides the lower and upper bounds for the lowest possible irreducible error in inferring the correct cognitive heuristic.

$$\frac{(\mathbb{H}_p - P_2 \, log(1+N_h))^2}{4 \, N_h} \leq BPE \leq \frac{\mathbb{H}_p - P_2 \, log(1+N_h)}{2} \tag{5.7}$$

where,

$$\mathbb{H}_p = -\left[(1-g) \, log\left(\frac{1-g}{N_h} \right) + g \, log(g) \right] \tag{5.8}$$

Intuitively, this property measures what proportion of behavior does not lead to overlapping inferences about the use of heuristics. A high value of P_2 protects against the inference problem of *multiple matches*.

3. *Stability property* (P_3)*:*

We define a distance metric $d_{kj_1j_2}$ between any pairs (x_{kj_1}, x_{kj_2}) of observed behavioral patterns, as the Euclidean distance based on a value of 1 if a cue is searched and 0 if a cue is not searched, measured over all $(n_A n_O)$ cues. We identify all pairs $\bar{\bar{x}}_{kj_1kj_2}$ that have the lowest possible distance.

$$\bar{\bar{x}}_{kj_1kj_2} = argmin_{(x_{kj_1}, x_{kj_2})} \, d_{kj_1, kj_2} \tag{5.9}$$

Stability is defined in terms of the mean absolute change in inferring probability of use of a heuristic, given a change in the behavioral pattern, integrated over all heuristics in the consideration set, and measured over all pairs of behavioral patterns that belong to the set $\bar{\bar{x}}_{kj_1kj_2}$, where $n(\bar{\bar{x}}_{kj_1kj_2})$ refers to the cardinality of this set. Stability is in the range [0,1], with higher values indicating higher stability of inferences made about latent heuristics.

$$P_3 = 1 - \frac{\Sigma_{\bar{\bar{x}}_{kj_1kj_2}} \Sigma_i \left| p(h_i|x_{kj_1}) - p(h_i|x_{kj_2}) \right|}{N_h \; n(\bar{\bar{x}}_{kj_1kj_2})} \tag{5.10}$$

Intuitively, this property measures how robust the strategies are in terms of whether minor changes in behavior will be attributed to significant changes in the inferred latent processes. A high value of P_3 will protect against an inference problem that tends to over-detect changes in latent switching.

Note that there may typically be trade-offs between P_1 and P_2, and it is useful to measure an aggregate value of the degree of well-posedness W of a consideration set of heuristics, measured as:

$$W = \frac{1}{3} \Sigma_{i=1:3} \; P_i \tag{5.11}$$

In addition to being well-posed, the heuristics should have a high degree of predictive capability, in that, *once it is inferred which heuristics is being used*, it should be capable of making strong inferences about what search patterns and choice options will be selected. To illustrate, the guessing strategy above has no predictive capability, since it accords equal probabilities to all observed behaviors. We thus define a predictiveness property P_4.

4. *Predictiveness property* (P_4)*:*

To measure predictiveness, we base it on the within-heuristic entropy across all possible search and decision patterns, integrated across all heuristics. This also yields a value in the range $[0, 1]$. Note that this measure has to be read in tandem with P_1 and P_2, and there will typically be some

trade-offs between P_1, P_2, and P_4.

$$P_4 = \frac{(1-g)}{N_h} \Sigma_i \left[1 + \frac{\Sigma_k \Sigma_j \left(p(x_{kj}|h_i)\ log(p(x_{kj}|h_i)) \right)}{log(N_s\ n_O))} \right] \tag{5.12}$$

Intuitively, a high value of P_4 reduces the uncertainty of predictions made by the heuristics about the range of possible behaviors. Note that P_4 is also likely to have trade-offs with P_1 and P_2, and these properties need to be viewed in tandem.

5.2 Comparing traditional and probabilistic heuristics

In table 5.1 we show the evaluation metrics for a consideration set consisting of TTB and WA, based on the experimental paradigm of experiments 1 and 2. The first half of the table shows the metrics for conventionally defined heuristics, with different values of noise ε and guessing g. The second half shows the kernel based probabilistic versions of TTB and WA prototypes, with the error model defined by different values of σ. Observe that conventional heuristics find it hard to find a balance between existence P_1 and uniqueness P_2. Probabilistic heuristics as defined here show a better balance, and overall improvement in scores. A high score on W implies better expectation about the inference process, whereas a better value of P_4 implies a better quality of predictions once we are able to make suitable inferences.

A key reason for the relatively low scores in the traditional case, even with extremely optimistic assumptions about the error rates, is because information search and decision heuristics have traditionally been defined as rule-based, without taking into account the gradual deviation of possible behavior from such deterministic rules. We now apply the same measurement framework to our novel probabilistic definition of heuristics. Note that in this case, there is no need to specify an application error ε or a guessing rate g. This is however replaced by the parameter σ.

A word of caution while using probabilistic heuristics stems from the observation that P_4 seems to drop off significantly at higher values of σ. This can be easily understood intuitively, since as the kernels become more widely distributed, the uncertainty of their predictions increases significantly. At the same time, if these distributions are too narrow, they will suffer the same fate as traditional heuristics (the W approaches that of traditional heuristics for low σ). It is important to test the values appropriately, and the measurement framework provides a tool for this.

		P_1	P_2	P_3	W= $\Sigma_n P_n/3$	P_4	$\Sigma_n P_n/4$
exact	$\varepsilon = 0.01$	0.0	0.61	0.99	0.53	0.98	*0.65*
min	$\varepsilon = 0.01$	0.02	0.55	0.99	0.52	0.74	*0.58*
choice	$\varepsilon = 0.01$	0.40	0.06	1.0	0.49	0.12	*0.40*
exact	$g = 0.01$	0.0	0.67	0.99	0.55	0.98	*0.66*
min	$g = 0.01$	0.02	0.60	0.99	0.53	0.75	*0.59*
choice	$g = 0.01$	0.40	0.09	1.0	0.50	0.12	*0.40*
prob	$\sigma = 0.01$	0.03	0.63	1.0	0.55	1.0	*0.67*
prob	$\sigma = 0.02$	0.86	0.63	0.82	0.77	1.0	*0.83*
prob	$\sigma = 0.05$	0.99	0.63	0.80	0.81	0.99	*0.86*
prob	$\sigma = 0.20$	0.99	0.59	0.80	0.79	0.35	*0.69*

Table 5.1: Evaluation metrics for a traditionally defined consideration set of WA and TTB heuristics based on 4 attributes and 3 choice options, and heuristics applied based on the true cue validities (top half of the table), compared with kernel based probabilistic versions (bottom half). The first column indicates the method of inference.

5.3 Applications and future work

The evaluation framework for heuristics should be seen as a first step towards a unified and systematic approach to defining strategy selection and learning models. The probabilistic framework is generalized enough to be applicable to a variety of experimental and empirical designs and heuristics. It can easily be incorporated with existing approaches to learning, such as rational meta-reasoning, reinforcement-learning, and cost-benefit based or cognitive effort based frameworks. Importantly, it has the potential to unify rule-based and exemplar based heuristic models. The heuristics described under our framework generally perform better than conventional heuris-

tics under the proposed evaluation measures. This framework raises a lot of possibilities in terms of future work, including experimental design based on maximizing information gain, and generating a new class of heuristics based on context specific, process driven, or exemplar measures. The evaluation framework allows us to calculate the a priori performance measures of a set of kernel based heuristics, for each particular configuration of cues. This evaluation method can be used to for experimental design to select cue configurations that allow for the strongest possible inference given a particular set of heuristics to be tested, that is, by selecting configurations that maximize measures P_1 to P_4 for the heuristics to be tested. Note that P_1 and P_2 often compete. Since these are all measured on the same scale $[0,1]$, we can use an objective function that optimizes a weighted average of the four measures.

Part II

Adaptive reference points

Chapter 6

Theoretical considerations

6.1 Introduction

Reference points have been known to play a role in a large number of psychological phenomenon, including aspects such as anchoring, evaluation of utilities (e.g. in prospect theory), serving as a criteria for signal detection, serving as a goal or aspiration, and as a base value for measuring hedonic utility, among others. While, the multiple roles of reference points in individual choice have been acknowledged (e.g. Kahneman (1992)), and the role of adaptive preferences based on previous experiences (e.g. Frederick and Loewenstein (1999)) has been acknowledged, relatively little work has been done in terms of specifying a robust model of adaptive reference point formation, since the original mean-stimulus model was proposed by Helson (1948,9). In this section, I propose a basic cognitive mechanism of adaptive reference points based on learning from previous experiences, and show how this mechanism can be ubiquitously applied across a wide range of behavioral tasks. I show that including such an adaptive reference point contributes to both: improving the quality of inferences made, as well as the predictive capabilities of the underlying cognitive models. In this chapter, I first propose a generic mechanism for reference point adapta-

tion, applicable to reference points in various contexts. I then show how such a mechanism can fit into different context-specific cognitive modeling frameworks.

6.2 Generic learning mechanism

6.2.1 Basic learning mechanism

The generic reference point adaptation mechanism is proposed as a reaction to an external variable, which, depending on the context, can be either a form of feedback, stimuli, or ancillary environmental factor. The reference point at time t is denoted as λ_t. The adaptation process at time t is captured in equations 6.1-6.3. Here, μ_t is the external influencer, which could be a stimulus or feedback value that serves to shift the reference point. The terms η_t and ρ_t refer to the immediate signal and the persistence signal, the latter being the cumulative decaying but persistent influence of all previous time periods. The immediate signal depends on a free parameter δ, characterized as *gain control*, or the sensitivity to the immediate difference. The cumulative signal ρ also depends on the free parameter α, characterized as *persistence*. High values of persistence essentially result in a stronger influence from preceding time periods. Depending on the context, the reference point may be updated at very time period, or only when an adverse event occurs, or there may be a bias where the the update is stronger in one direction that the other. These specific instances are covered in section 6.2.2.

$$\lambda_t = \lambda_{t-1} + \left(\frac{\eta_t + \rho_t}{1 + \Sigma_{i=1}^{t-1} \alpha^{t-i}} \right) \tag{6.1}$$

$$\eta_t = \delta(\mu_{t-1} - \lambda_{t-1}) \tag{6.2}$$

$$\rho_t = \delta \ \Sigma_{i=1}^{t-1} \ (\alpha^{t-i} \ (\mu_i - \lambda_i) \ ) \tag{6.3}$$

The term η_t is the contribution of the immediately preceding trial, which we call the *immediate signal*. The cumulative contribution of all previous trials is given by the ρ_t term, which we term *persistence signal*. Here, the weight given to older feedback keeps decreasing, and is a function of α. The feedback term from j trials earlier is given a weight of α^j. A weighted average of η_t and ρ_t is then computed as the final corrective feedback for the criterion level. Having large values of α result in a weighted average over a longer time window, leading to higher persistence of feedback and lower flexibility of response changes. Since α acts as a discount rate for previously acquired feedback, a value of $\alpha = 1$ means that on each trial, the effective feedback is the mean value of feedback acquired on all trials experienced so far. A low value of α, close to 0, would mean that feedback from only the most recent trial is taken into account. Depending on context, the reference point may be compared, contrasted, or assimilated within the next judgment, along with the stimulus or signal μ_t received on the next time period.

Figure 6.1 shows a simulated movement of the reference point (red lines) for the same stimulus/feedback (gray line), for different combinations of low and high values of δ and α. The different influence of these parameters is visible when viewing the difference signal, shown by the bars. This difference signal at any point in time is critical to the decision, as we will see in the next few sections. As we move from the left 2 panels (low α) to the right panels (high α), the reference point seems to oscillate around the focal stimulus or feedback signal, with more frequent swings of the difference signal above and below zero. As we move from the top 2 panels (low δ) to the bottom panels (high δ), the reference point gets closer to the focal signal, and results in smaller magnitudes of the difference. This is of course based only on one single pattern of stimulus/feedback signal, and is meant to illustrate the effect of the parameters. The actual influence will depend on the trends shown by the focal signal (e.g. whether it is monotonically changing in one direction, how volatile it is, whether it is drawn from non-stationary distributions over a period of time, etc.).

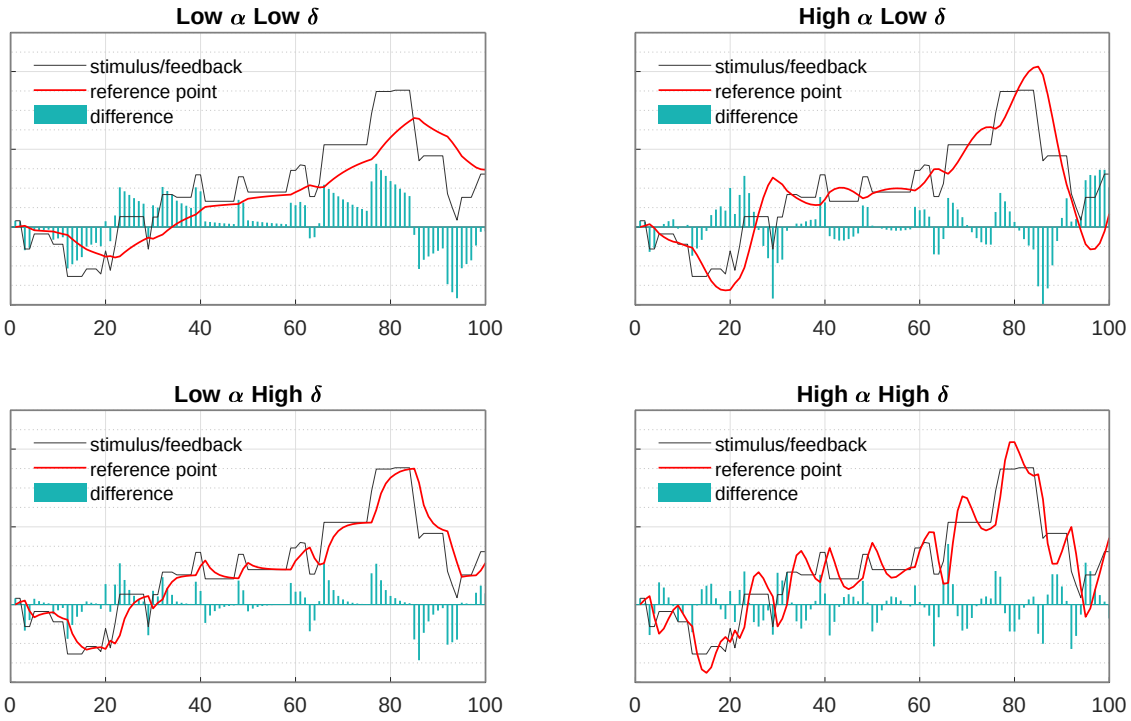

Figure 6.1: Illustration of the adaptive mechanism: Single simulation of reference point movement around the focal stimulus/feedback signal, for different combinations of low and high value of δ and α

6.2.2 Introducing asymmetry

Confirmation bias, where people place a higher weight on information that confirms rather than contradicts their beliefs, has been shown to be pervasive over a large range of cognitive processes (Nickerson (1998); Jones and Sugden (2001)). This has been shown to extend to reinforcement learning, where learning effects are higher for information that confirms rather than contradicts the current choice that individuals make (Palminteri et al. (2016)). A different line of research shows that reference point levels for evaluating gains and losses are adaptive (Gneezy (2005)), and this adaptivity has been shown to be asymmetric such that adaptation after gains is faster than adaptation after losses (Arkes et al. (2008,0)). Johnson et al. (2012) showed that in a situation with multiple possible reference points, individuals tend to use the reference point that maximizes the utility of their current behavior. These findings are directly relevant to reference point adaptation in the model presented here. We allow for asymmetry by introducing a bias term to change equation

6.1 to equation 6.4.

$$\lambda_t = \lambda_{t-1} + \left(\frac{\eta_t + \rho_t}{1 + \Sigma_{i=1}^{t-1} \alpha^{t-i}} \right) \left(I_{t-1} + B\,(1 - I_{t-1}) \right) \tag{6.4}$$

Here, I_{t-1} is an indicator function, and B is a bias term that applies depending on whether the indicator function is 0 or 1. Depending on the type of application, the indicator function may indicate either whether the reference point λ_{t-1} was higher or lower than the focal signal μ_{t-1}. Another form of indicator could be when there is an objective binary indicator, such as whether or not the previous item was correctly or incorrectly responded to. The bias B can be treated as a free parameter.

6.3 Applications

6.3.1 Application to anchoring judgments

Any quantity to be judged may be affected by anchoring effects. Formally, if the true value of the item to be judged is p_t, and the anchoring reference point is λ_t, then the judged value is given by:

$$J_t = p_t(1-k) + k\lambda \tag{6.5}$$

Here, k is the strength of the anchoring effect, with a value of 0 implying no anchoring, and a value of 1 implying completely anchoring on the reference point.

Anchor points are often inferred from data, but there is no quantitative cognitive process model to explain how these anchoring points may be formed over time. We propose the same adaptive reference point mechanism from equation 6.1. Thus, we obtain, equation 6.6, where λ_t is updated based on the feedback about the true value of the item p_t.

$$J_t = p_t(1-k) + k\lambda_t \tag{6.6}$$

This adaptive reference point based anchoring mechanism is tested with experimental data in

section 7.1, where people make repetitive judgments about different items in the same category

6.3.2 Application to expectation formation

When facing uncertainty in rewards and penalties from different sources, as might be represented by a bandit problem, people are known to use different strategies to explore and exploit the environment. In one study, it was shown that close to half the people are best represented by a win-stay-lose-shift strategy Steyvers et al. (2009). When options are not win/lose but provide different levels of rewards, this strategy can be extrapolated as a High-Stay-Low-Shift strategy, so that people shift their choices if the rewards are low but continue (stay) their existing choice if the rewards are high. This naturally requires some reference point for what counts as low or high. Accordingly, we can define the probability of a *stay* decision as in equation 6.7, where p_t is the reward obtained on trial t.

$$\mathrm{p(Stay)_t} = \frac{1}{1+e^{-(p_{t-1}-\lambda_0)}} \tag{6.7}$$

Here, for a given reward value p_t, a low reference point λ_0 increases the probability of staying, and vice versa. When the last obtained reward is equal to the reference point, the probability of a stay decision is 0.5. Rather than a static reference point λ_0, we propose an adaptive reference point based on equation 6.1, where the previous outcomes serve to adjust the reference point, λ_t. Thus we get equation 6.8.

$$\mathrm{p(Stay)_t} = \frac{1}{1+e^{-(p_{t-1}-\lambda_t)}} \tag{6.8}$$

This adaptive reference point based model of expectation formation is tested with experimental data in section 7.2.

6.3.3 Application to tracking probabilities

A natural baseline model for tracking the probability of a repetitive binary event over time is Bayesian learning, implemented using a beta counting model. Here the probability of an event occurring is drawn from a beta distribution that is updated on every trial t, where h_t is the count of times the event has occurred up to the previous trial, and $l_t = t - h_t$ is the count of times that the event has not occurred. This model assumes an uninformative prior and perfect Bayesian learning.

$$\text{Event}_\text{t} \sim \text{Beta}(1 + h_t, 1 + l_t) \tag{6.9}$$

As an alternative, we propose the adaptive reference point mechanism of equation 6.1 as a tracking mechanism for the probability of an event, given as in equation 6.10, that depends only on λ_t, which is increased based on equation 6.1 every time the event occurs, and is decreased every time it does not. Since the rate of change are driven by individual level parameters α and δ, this allows for sub-optimal learning as well as strong persistence of prior signals. As the reference point λ_t increases, the judged probability increases, and vice versa. The probability is 0.5 when $\lambda_t = 0$.

$$\text{p(event)}_\text{t} = \frac{1}{1 + e^{-\lambda_t}} \tag{6.10}$$

This adaptive reference point based probability tracking mechanism is tested with experimental data in section 7.3

6.3.4 Application to signal detection theory

Classical signal detection theory can account for behavioral patterns extremely well within a fixed environment. It does not, however, provide a descriptive account of how people adapt their criterion in response to environmental manipulations, such as changing base rates, changing discriminability, or changing utilities, for different types of correct decisions and errors (although it prescribes what the normative change in criteria *should* be). There are some existing theories of how people

may adopt a flexible rather than static criterion across trials Treisman and Williams (1984); Erev (1998); Turner et al. (2011). Here we introduce an alternative adaptive account of how people set criteria for categorizing stimuli within the SDT framework, that is based on the basic reference point adjustment mechanism described in the previous section.

We make a basic assumption that people are sensitive to feedback, and thus adjust their criterion, only on incorrect or punitive trials. We thus specify a criteria adjustment mechanism based on equation 6.4, with $B = 0$, where I_{t-1} is an indicator function that is 1 if the $(t-1)$th trial was incorrect and 0 otherwise. Under this assumption, we can rewrite 6.4 as equation 6.11.

$$\lambda_t = \lambda_{t-1} + I_{t-1}\left(\frac{\eta_t + \rho_t}{1+\Sigma_{i=1}^{t-1}\alpha^{t-i}}\right) \tag{6.11}$$

In the case of adaptive SDT, the term μ_t is the signal on trial t and λ_t is the criterion on trial t. The term $(\mu_{t-1} - \lambda_{t-1})$ is the underlying difference signal between the stimulus and criterion, and represents the error signal on incorrect trials. If the $(t-1)$th trial was a miss because the criterion was too high, this term will be negative and serve to lower the criterion. If the previous trial was a false alarm because the criterion was too low, this term will be positive and serve to increase the criterion. This difference signal is modulated by the gain control parameter δ. Higher values of δ imply a larger corrective feedback given a particular level of sensory feedback. Note that if the previous trial is a hit or correct rejection, the criterion will not change. However, if the previous trial is incorrect, the change made includes feedback based not only on the immediately preceding trial, but also feedback weighted and averaged from previously experienced trials, including correct trials, so that learning from history is not heavily biased. On correct trials, the difference signal term is positive for hits and negative for correct rejections.

This mechanism is tested on experimental data in section 7.4.

Interpretation in terms of ROC characteristics

Classical SDT often uses receiver operating characteristic (ROC) analysis which plots performance in terms of hit and false alarm rates. Figure 6.2 characterizes ASDT in ROC terms, showing the results of simulations that systematically varied α and δ in 4 different environments, that in turn varied in terms of base rate and discriminability. This represents a range of environments typically used in signal detection problems. We show that the adaptive SDT has a strong relationship to the dynamics of the ROC. The resulting performance is plotted along the ROC in figure 6.2. The left panel shows the sensitivity to α. As α increases from 0.01 to 1, it results in a smooth change *along the ROC*, and *away* from the optimal criterion. For very low values of α, behavior still deviates from optimal performance, but to a smaller extent. Very high values of α close to 1 show maximum deviation away from optimality. The right panel of Figure 6.2 shows the sensitivity to δ, with increasing values of δ showing *a movement away from the ROC*, with reducing hit rates in low BR and increasing false alarm rates in the high BR conditions. Thus α and δ capture two separable behavioral deviations: along the ROC or away from the ROC. Changes in δ capture what in traditional SDT analysis, is captured as a difference in sensitivity. We note that the simulations show that the mean criterion level is extremely sensitive to values of α, with higher values of α leading to less extreme average criterion values. Higher values of δ on the other hand result in higher variability in the criterion across trials.

Adding payoff sensitivity to the reference point adjustment

We posit that if the task rewards are symmetric, that is, there is no difference in the rewards for correctly identifying a hit or a correct rejection, or between the costs incurred for a miss or a false alarm, the adaptive mechanism above is sufficient. However, if the tasks are asymmetrical, a payoff sensitivity needs to be incorporated into the criterion adjustment mechanism. First, I propose that a prospect theory-like value function (equation 6.12) is used to transform the objective rewards

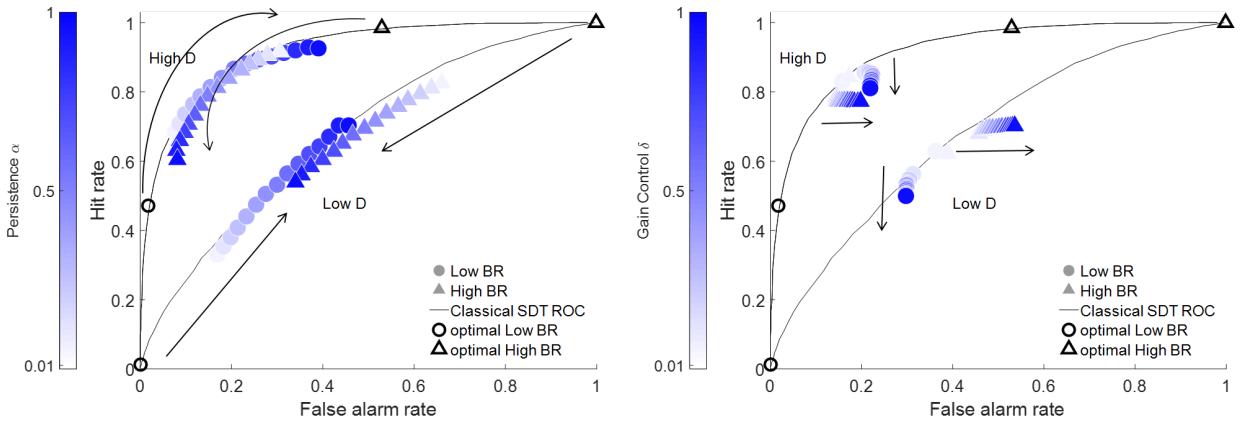

Figure 6.2: ROC curves for high-D and low-D environments (i.e. different experimental level of discriminability), and the hit rate and false alarm rate based on optimal criterion placement for both low BR and high BR conditions. The colored plots show how a change in α (left panel) and δ (right panel) affect how individual behavior moves away from optimality. Increasing α results in movement along the ROC (does not affect sensitivity or discriminability), but changes in δ shift performance to a lower ROC (impact sensitivity).

and penalties, and the transformed rewards and penalties are used to calculate the optimal bias R (Stüttgen et al. (2011)), as shown in equation 6.13. However, people may show only partial sensitivity to the optimal bias, and this is denoted by θ. Finally, the actual bias, $\bar{B}$, as shown in equation 6.14, is applied as a multiplier to the difference signal only when the stimulus on the last trial was greater than the criteria (and a multiplier of 1 otherwise). Here, V_F, V_C,V_M, and V_H refer to the payoffs in case of false alarms, correct rejections, misses and hits respectively. Here, L is a loss aversion parameter.

$$.f(x) = (x >= 0)\, x^{\gamma} - (x < 0)\, L\, (-x)^{\gamma} \tag{6.12}$$

$$R = \frac{f(V_F) - f(V_C)}{f(V_M) - f(V_H)} \tag{6.13}$$

$$\bar{B} = \psi((1-\theta) + \theta\, R) \tag{6.14}$$

Thus, under the assumption of asymmetric payoffs, we can rewrite equation 6.11 as equation 6.15, where J_t is 1 if $\mu_{t-1} > \lambda_{t-1}$ and 0 otherwise.

$$\lambda_t = \lambda_{t-1} + I_{t-1}\left(\frac{\eta_t + \rho_t}{1+\Sigma_{i=1}^{t-1}\alpha^{t-i}}\right)\left(\bar{B}J_{t-1} + (1-J_{t-1})\right) \tag{6.15}$$

6.3.5 Application to utility functions

Economic methods often model consumption decisions (how many units of x do people purchase at time t) in terms of the psychological utility $U(x)$ of x and the unit cost p_t of x. The consumption decision is modeled as the quantity x that maximizes $V(x) = U(x) - px$. For instance, a commonly used utility function is of the form:

$$U(x) = \frac{ax^{1-b}}{1-b} \tag{6.16}$$

Maximizing $V(x) = U(x) - px$ yields a solution:

$$x = \left(\frac{p}{a}\right)^{-1/b} \tag{6.17}$$

The parameters a and b are traditionally estimated by converting this to a regression equation of the form below, where $\beta = -1/b$.

$$log(x_t) = A + \beta\, log(p_t) \tag{6.18}$$

This equation 6.18 is assumed to capture the change in consumption units x_t as the corresponding price per unit p_t changes, where A and β are parameters at the individual or population level. In typical applications, $\beta < 0$, so that consumptions reduces as the price increases.

Thaler (1999, 2008) proposed that consumption choices were driven by a combination of acquisition utility and transaction utility. Acquisition utility is defined as above, in terms of rational economic analysis, as the the utility of consumption less the value of the price paid. Transaction utility typically reflects the *value of the deal*, and is the difference between the price paid and a *reference price*. The transaction utility can be positive if the actual price is lower, and negative if it is higher, than the reference point, λ. The mental accounting theory proposes that the acquisition

and transaction utilities are *separately* evaluated, and may be weighted differently. Thus, we now propose equation 6.19, where k is a relative attention weight, and the value of the transaction is higher than rationally observed if the reference pointλ is higher than the price, and vice versa.

$$V(x) = U(x) - px - k(p - \lambda)x \tag{6.19}$$

We propose that this reference point λ is not static, but changes over time based on previous experiences with prices or other similar indicators. Thus, we propose the identical reference point adjustment mechanism as in equation 6.1, and replace λ in equation 6.19 with an adaptive reference point λ_t on every time period t. Maximizing this yields a solution:

$$x_t = \left(\frac{p_t + k(p_t - \lambda_y)}{a} \right)^{-1/b} \tag{6.20}$$

The regression equation 6.18 then changes to equation 6.21.

$$log(x_t) = A + \beta \, log(p_t + k(p_t - \lambda_t)) \tag{6.21}$$

Most econometric methods treat β as a form of demand elasticity to price, and assume short term and long term constancy and symmetry. However, there is empirical evidence for instability of demand elasticities, asymmetries in upward and downward responses to price changes, and other behavioral aspects that cannot be explained without additional assumptions about the cognitive process. This reference point mechanisms shows how a constant β can still lead to inferring unstable and asymmetric price elasticity inferred because the adaptive reference point mechanism is not taken into account. Note that on each trial or time period, the consumption quantity x is assumed to be decided based on maximizing the utility model. The reference point is an *inducing* reference point, since a higher reference level makes any price seem relatively more attractive, thus increasing utility. Extremely high values of the reference point mathematically, would result in infinite utility, inducing people to spend all resources and maximize the units of consumption. Such reference levels are however psychologically implausible. Mathematically, the upper bound on psy-

chological reference point depends on the current price levels and is given by $\lambda_t <= \frac{p_t\ (1+k)}{k}$, beyond which, utility becomes infinite.

This adaptive reference point based utility function is tested with experimental data in section 7.5, and also adapted to work with real world data on consumption taxes in chapter 9.

6.4 Conclusions

This chapter proposes that a basic adaptation mechanism, well specified in quantitative terms, can account for and assimilate into multiple modeling frameworks that account for different types of judgment and decision making tasks. In the next chapter, each of these propositions is successfully tested using experimental data collected from secondary sources.

Chapter 7

Ubiquity of the adaptive reference point

All of the cognitive models within this chapter were implemented as Bayesian graphical models using JAGS.

7.1 Price Judgments

We implement the adaptive reference point based model of anchoring judgments, as specified in section 6.3.1, to analyze behavior in an experimental price judgment task.

7.1.1 Data

This is a secondary dataset, obtained from work reported in Matthews and Stewart (2009). This data includes data from 28 subjects who were asked to judge the price of 102 real items that differed on a single physical dimension. The item set consisted of women's shoes. Participants had to judge the price of each item, and after judging the price were provided feedback on the true price of the item. The original study found that providing feedback resulted in anchoring judgments towards

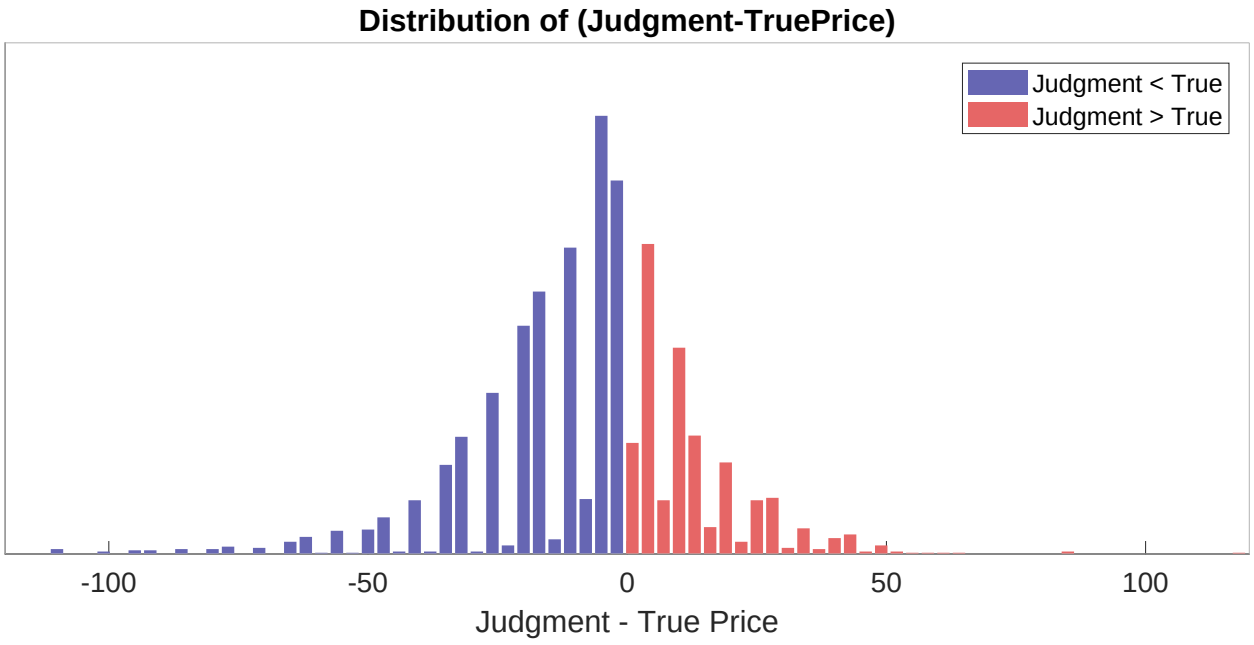

Figure 7.1: Distribution of judgment errors

the most recent item, but also that there was an overall bias in judgment towards the center of the true price range. They present a discriminative regression model to measure the effect of previous item prices, but do not offer any formal cognitive model of the anchoring process. Figure **??** shows the distribution of difference between the judged and true price of all items across all participants.

7.1.2 Modeling

For the baseline model, we propose that the trial-by-trial judgments J_t are drawn from a normal distribution centered around the true price, but with an anchoring bias based on a *fixed* reference point λ, where k is the strength of the anchoring effect, and σ is a form of expertise, the precision with which judgments are made. The strength of anchoring k lies in the range $[0,1]$ and effectively serves as a weight for the reference point. A value of zero implies no anchoring, and a value of 1 implies complete anchoring and ignoring the true price effect.

$$\mathrm{J_t} \sim \mathrm{Gaussian}\big(p_t - k(p_t - \lambda), \sigma\big) \tag{7.1}$$

Along with that, we build an adaptive reference point model, that infers the latent adaptive reference point λ_t based on equation 6.1, with α and δ as individual level parameters. Here, p_t is the true price, revealed after each judgment, and serves to update the reference point before the next trials as in equation 6.1.

$$J_t \sim \text{Gaussian}\big(p_t - k(p_t - \lambda_t), \sigma\big) \tag{7.2}$$

7.1.3 Modeling results

Figure 7.2 shows show the RMSE error for each participant, with the gray lines showing the error based on the baseline model, the green bars showing an improvement based on the adaptive reference point model and the red bars showing a deterioration based on the adaptive model. The descriptive error was based on providing the full data to the models, whereas the generalizability error shows the results where 68 of the 102 judgments were hidden from the model. The adaptive model improves descriptive and predictive error compared to a baseline fixed reference point model.

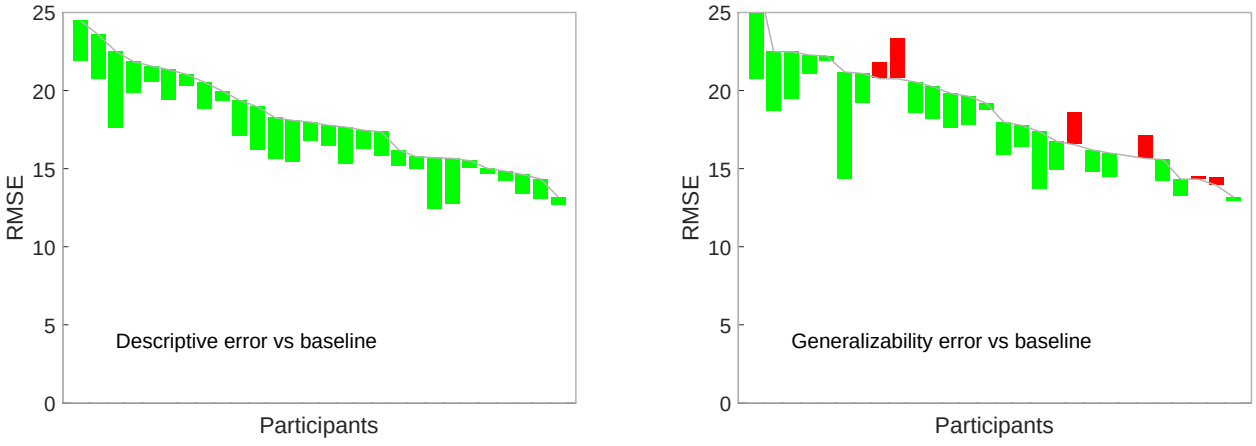

Figure 7.2: Price judgments modeling results: The plots show the RMSE error, with the gray lines showing the error based on the baseline model, the green bars showing an improvement based on the reference point model and the red bars showing a deterioration based on the adaptive model. Each point / bar represents a single participants.

7.1.4 Inferences

Figure 7.3 shows the robustness of the psychological interpretation of the strength of anchoring k, with the latent inferred strength being strongly positively correlated with the mean absolute error in judgments of participants. In figure 7.4, the plot on the right shows that because of random

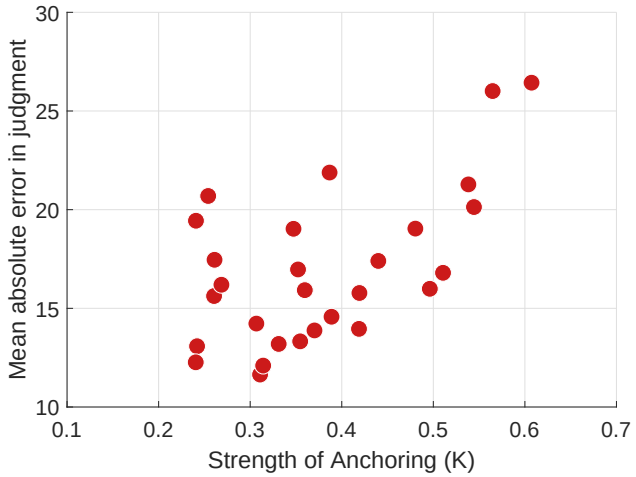

Figure 7.3: Increased latent strength of anchoring k (an inferred parameter), greater the observed mean error in judgments by participants, on an average.

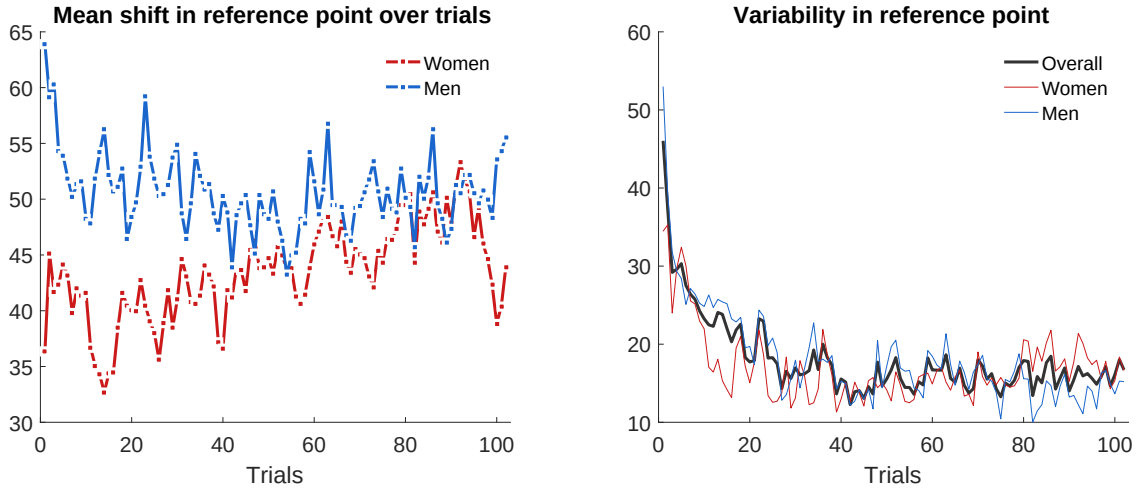

Figure 7.4: Mean shift in reference points. The plot on the right shows that because of random ordering of stimulus, the variability in reference drops over time, explaining the regression to the mean effect recognized in the original study. The plot on the left shows the mean reference point average from men and women separately. It shows that men begin with a higher reference point than women, but after about half the trials, the reference point for men drops and that for women increases, till they are approximately equal.

ordering of stimulus, the variability in reference drops over time, explaining the regression to the mean effect recognized in the original study. The plot on the left shows the mean reference point average from men and women separately. It shows that men begin with a higher reference point than women, but after about half the trials, the reference point for men drops and that for women increases, till they are approximately equal. It is worth remembering at this point that the items being judged for price were women's shoes!

7.2 Bandit problems

We implement the adaptive reference point based expectation formation model, as specified in section 6.3.2 to analyze behavior in an experimental bandit task.

7.2.1 Data

This is a secondary dataset, obtained from work reported in Yechiam and Busemeyer (2008). This data includes 88 subjects from study 1 in the paper, each of whom played three distinct bandit problems. In each of the three types of problems, each subject made 200 sequential decisions, selecting one of two alternatives from which to receive payoffs. The payoff distribution for each alternative was fixed throughout each individual task, but unknown to the participants. The types of tasks were independent from each other. The first task (payoff-sensitivity) included 2 alternatives with different expected values and equal variance. The second (small-probability) task had one constant reward and one low probability reward alternative. The third (high-variance) task had one alternative with high expected value and variance, and the other with low expected value and variance. One of the primary questions raised in this paper was whether parameter estimates for choice models inferred from one of these tasks are transferable to other tasks, and the results applied with varying degrees of success. In the task, one key measure was the degree to which

people stay (exploit) or shift (explore). Figure 7.5 summarizes the behavioral results.

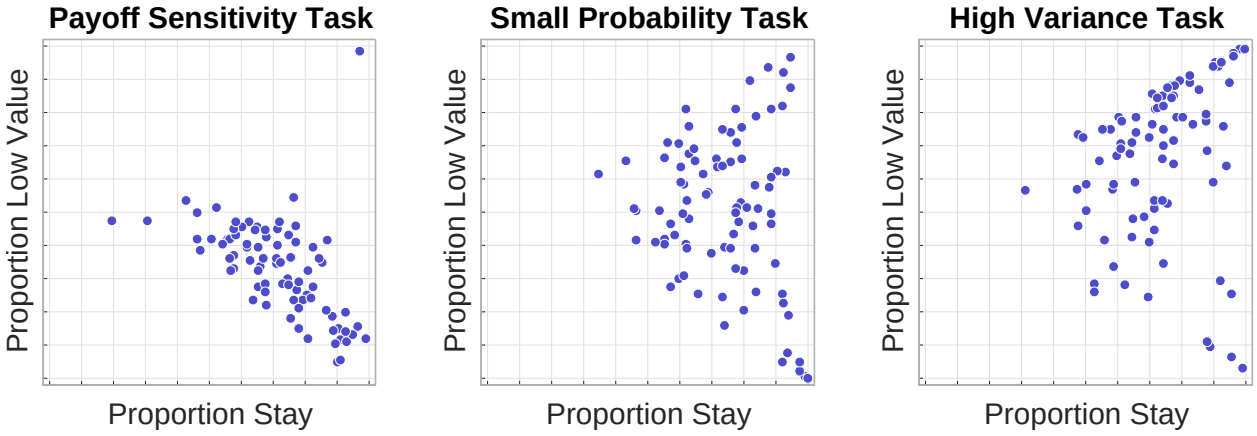

Figure 7.5: Behavioral results for the three tasks, showing the proportion of stay decision for each individual, and the proportion of low value choices

7.2.2 Modeling

The third (high variance) task was used to model participant behavior and infer individual level parameters, based on the first 100 trials. These same parameters were used as a generalizability test for predicting choices on the remaining 100 trials, as well as on all the trials in the other two tasks. As a baseline model, we implement a High-Stay-Low-Shift model with a fixed reference point λ_0. On each trial, the reward from the previous trials, denoted as p_{t-1} is compared to this reference point, and whether the decision is a stay decision is based on a draw from a Bernoulli distribution with a logistic function determining the Bernoulli parameter. Thus, a high reference point λ_0 serves to make the denominator in equation 7.3 very high, resulting in a low probability of staying, and vice versa. When the reward is exactly equal to the reference point, the probability of staying is drawn from a Bernoulli(0.5).

$$\text{Stay}_\text{t} \sim \text{Bernoulli}\left(\frac{1}{1+e^{-(p_{t-1}-\lambda_0)}}\right) \tag{7.3}$$

We also implement an identical model but with an adaptive reference point λ_t, which is updated on every trial based on the learning mechanism proposed in chapter 6. The reward p_t obtained on

the previous trial serves to create the difference signal $(p_{t-1} - \lambda_{t-1})$, which is used to update λ_t on the next trial based on the underlying parameters δ and α.

$$\text{Stay}_t \sim \text{Bernoulli}\left(\frac{1}{1+e^{-(p_{t-1}-\lambda_t)}}\right) \tag{7.4}$$

Figure 7.6 illustrates how the model works by providing the mean posterior predictive for 4 of the

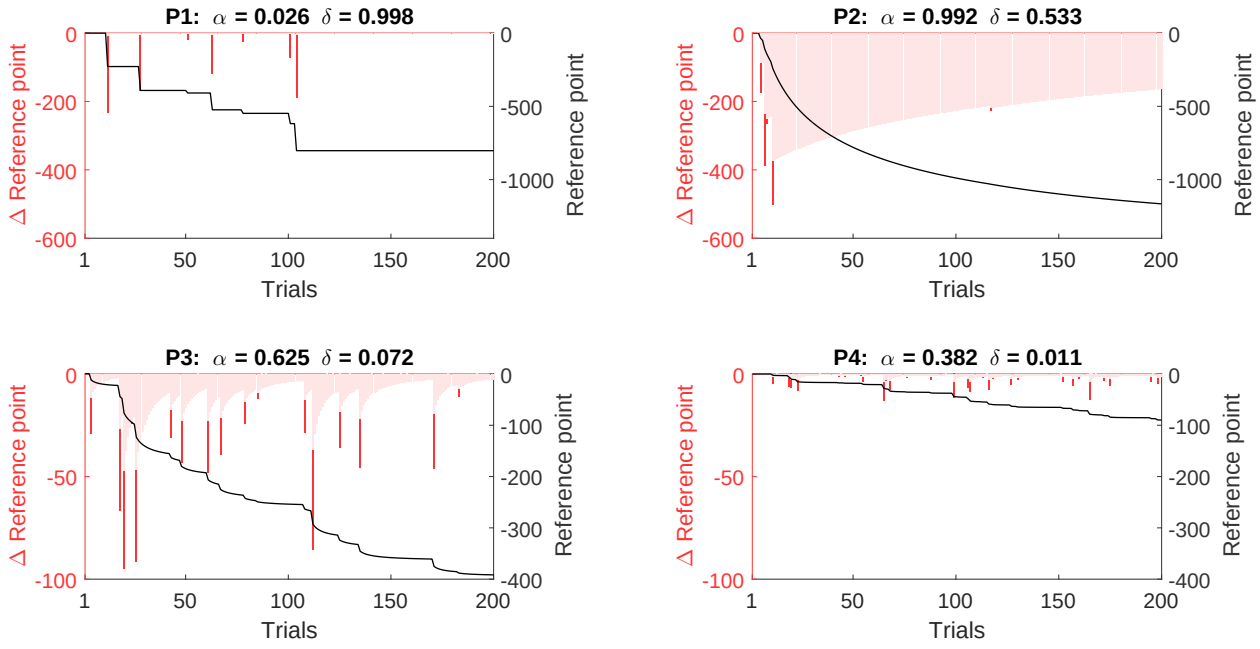

Figure 7.6: Illustration of feedback and persistence signal in 4 participants. The plots show the latent inferred reference point movement (black lines) and the trial by trial immediate feedback (bold red) and decaying persistence signals (light red). The 4 participants are selected with varying inferred values of δ and α, to represent different behavioral characteristics.

participants who have relatively different values of α and δ. The black lines in the figures shows the latent reference point over 200 trials. The bold red lines show the immediate feedback signals and the light red lines show how the influence of that signal decays over time in terms of future influences on the reference point. P1 is an example of a participant with very high inferred gain control δ and very low persistence α. The reference point moves in steps, every time a feedback signal is received, but there is no persistent influence because of the low α. P2 is an example of a participant with extremely high inferred persistence α. The reference point moves smoothly, and the impact of the initial negative feedback signals takes a very long time to decay. P3 and P4

show a much smaller movement of the reference point(see the different values on the right hand scale), because of very low values of the gain control δ. P4 barely registers any feedback, whereas P3 gives a classical representation of feedback signals with short decay streaks. The mechanism in this case is almost self-regulating, because a large initial drop in the reference point (P1 and P2) results in fewer feedback signals, whereas for P3 and P4 with lower sensitivity, the relatively higher reference point gives rise to more frequent negative feedback signals.

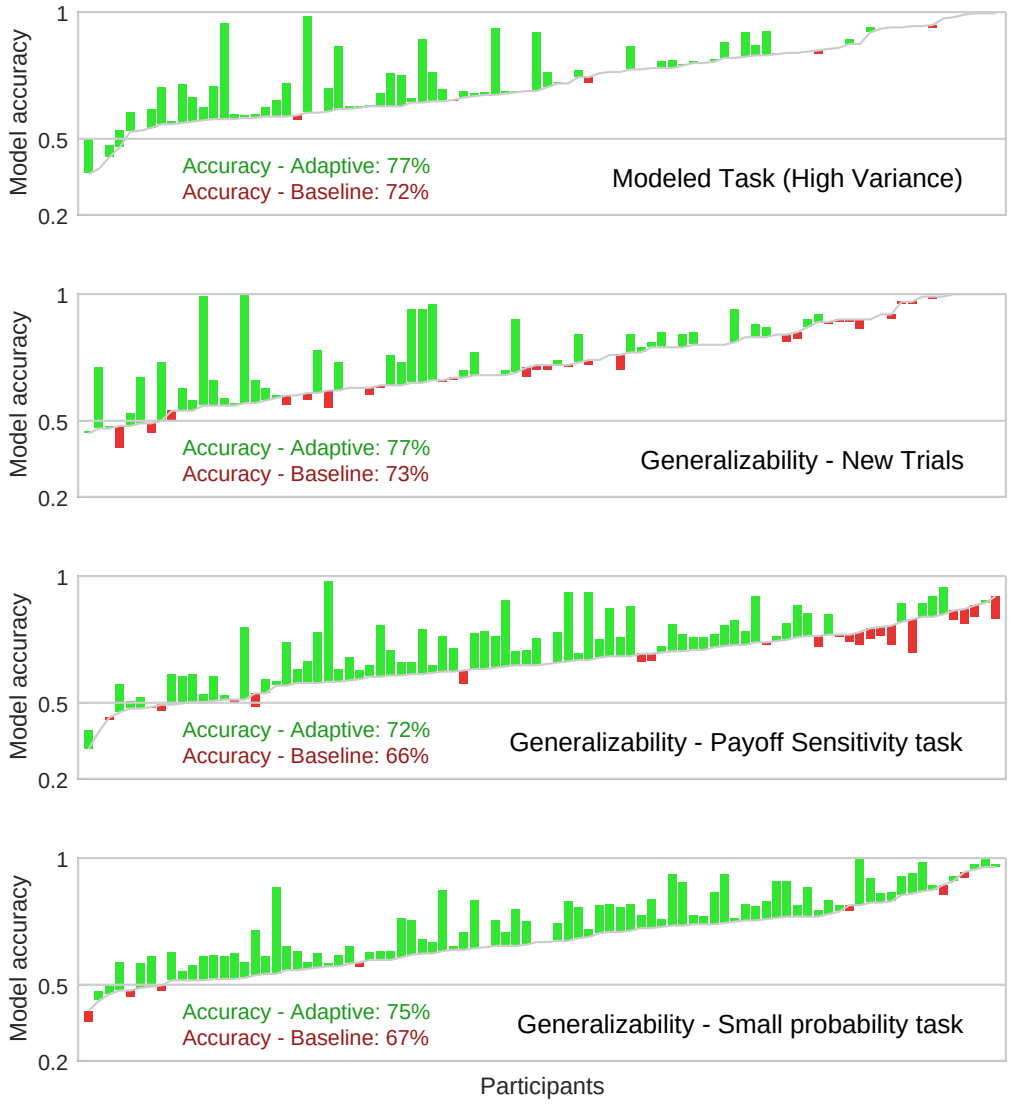

Figure 7.7: Modeling Bandit problems: Performance of the adaptive model against the fixed reference point model. The performance is coded in terms of whether the model correctly predicts a *Stay* decision on each trial. The participants are arranged in terms of increasing model accuracy for the base line model (gray line), with the green bars showing a relatively higher accuracy of the adaptive model and the red bars showing a drop in accuracy from the adaptive model.

7.2.3 Modeling results

Figure 7.7 shows the performance of the adaptive model against the fixed reference point model. The performance is coded in terms of whether the model correctly predicts a *Stay* decision on each trial. The participants are arranged in terms of increasing model accuracy for the base line model (gray line), with the green bars showing a relatively higher accuracy of the adaptive model and the red bars showing a drop in accuracy from the adaptive model. The model prediction is based on the modal choice made by the models. The adaptive reference point model is especially better than a fixed model when it comes to generalizing to new tasks based on the same parameters, because it allows the cognitive parameters to adjust the reference point. The adaptive model shows consistently higher accuracy, increasing 3 percentage point from from 72% to 75% in descriptive fit, and between 4 to 8 percentage points in the generalizability task. See figure 7.7 for details.

7.2.4 Inferences

Figure 7.8 shows the distribution of the mean posterior inferred values of gain control δ and persistence α for the participants. The 4 participants illustrated in figure 7.6 are highlighted. The colors represent participants falling in the four different quadrants, representing different high-low combinations of persistence and gain control. This should be seen in tandem with figure 7.9, which shows the mean trial by trial movement of the reference point across all the participants in each quadrant. The reference point is in the units of the reward. The blue horizontal bars in figure 7.9 show the distribution of the rewards. It is clear that most of the rewards are small, and the distribution is skewed towards the center, with a small probability of very large negative rewards. Participants with high *delta* have much lower adapted reference points, and thus react subsequently only when facing very large negative rewards. This allows them accrue higher average rewards over all the trials, as shown in the circles in figure 7.9.

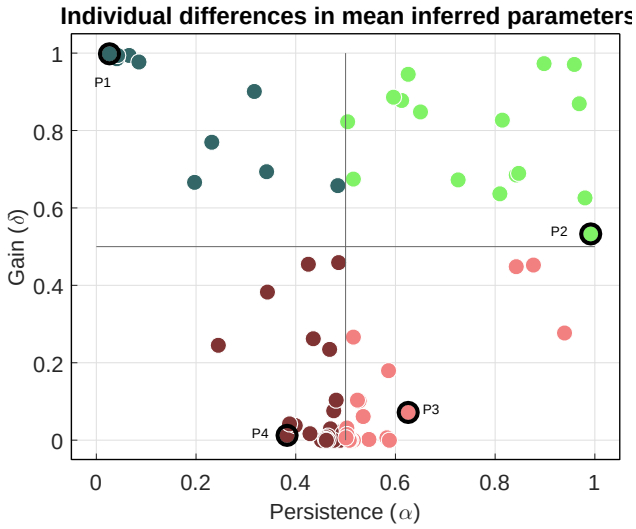

Figure 7.8: Inference on Bandit problems: Distribution of inferred parameters for all participants. The color coded dots should be read in tandem with figure 7.9

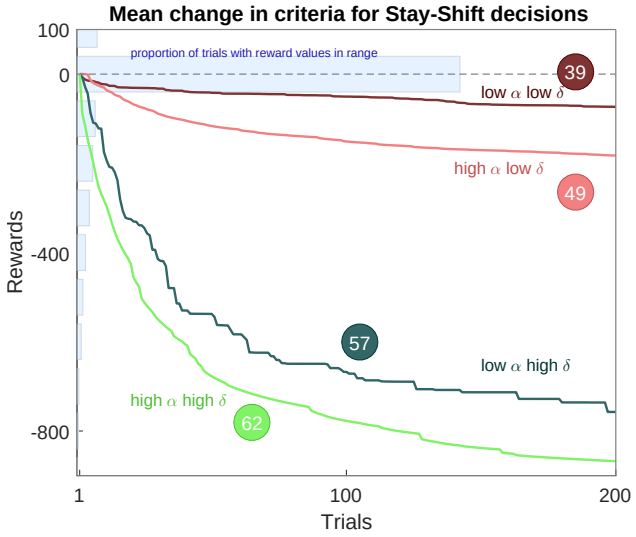

Figure 7.9: Inference on Bandit problems: Movement of the average reference point across participants in each quadrant of figure 7.8. This shows the clear difference in behavior based on the cognitive parameters, especially between participants with high and low values of δ. The color coded lines should be read in tandem with figure 7.8

7.3 Probability tracking and extrapolation judgments

We implement the adaptive reference point based probability tracking model, as specified in section 6.3.3 to analyze behavior in an experimental task.

7.3.1 Data

This is a secondary dataset, obtained from work reported in Frydman and Nave (2016). This data includes 38 subjects that made a series of 400 sequential judgments (judgments were elicited using a willingness to pay judgment) on whether a hypothetical financial instrument would increase or decrease, given the previous history. The average reported beliefs show that the probability of the current streak (increase or decrease) continuing is monotonically increasing with the length of the current streak (that is, number of consecutive increases or decreases). For example, the reported probability of a streak continuing is about 65% after a streak of 3 compared to about 60% after a streak of 2, and 50% after a single increase. The average reported probability moves closer to 70% after a streak of 6.

7.3.2 Modeling

We implement a Bayesian learning model based on a beta distribution, where the probability of an increase is drawn from beta distribution that is updated on every trial t, where h_t is the count of increases up to the previous trial, and l_t is the count of decreases up to the previous trial.

$$\text{Increase}_\text{t} \sim \text{Beta}(1+h_t, 1+l_t) \tag{7.5}$$

This is contrasted with an adaptive reference point based tracking model, where the probability of an increase on any trial is drawn from a Bernoulli distribution with the parameter being a logistic function of the adaptive reference point λ_t. This reference point is assumed to be updated on each trial based on the actual outcome observed, as in equation 6.1, with δ and α as individual level parameters.

$$\text{Increase}_\text{t} \sim \text{Bernoulli}\left(\frac{1}{1+e^{-\lambda_t}}\right) \tag{7.6}$$

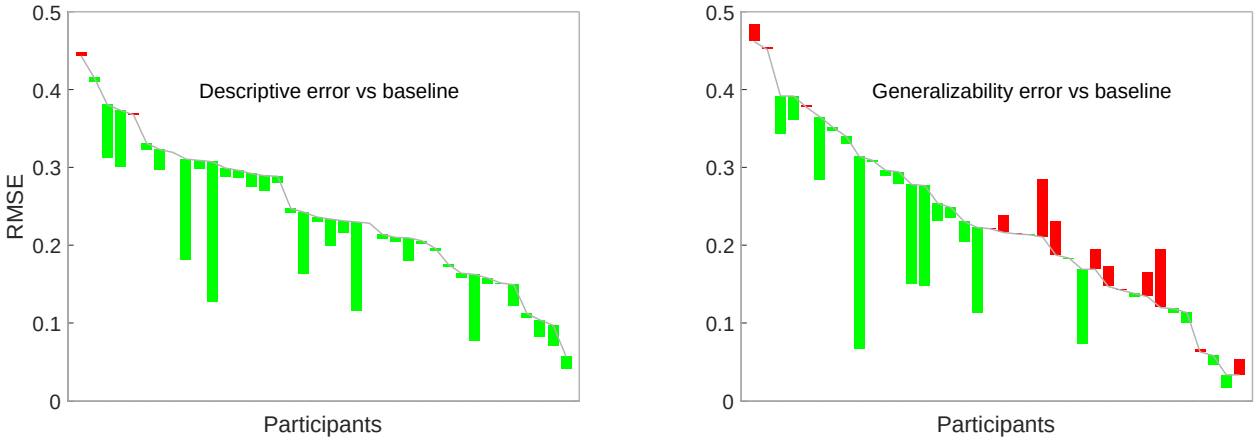

Figure 7.10: RMSE for probability tracking extrapolation judgments. The gray line gives the error for the baseline Bayesian learning model and the green and red bars show the improvement or deterioration based on the adaptive reference based model. Generalizability error is based on a model where the latter 200 trials are hidden from the model.

Figure 7.10 shows the RMSE error. The gray line gives the error for the baseline Bayesian learning model and the green and red bars show the improvement or deterioration based on the adaptive reference based model. Generalizability error is based on a model where the latter 200 trials are hidden from the model. There are some participants for whom a Bayesian learning model gives better predictive performance, but on an overall level, the adaptive model seems superior.

7.3.3 Inferences

Figure 7.11 shows the average movement of the reference point for 4 groups of participants, clustered based on high and low levels of α and δ. Participants with higher reference points report a higher probability of the increase occurring. In this task, higher persistence α plays a greater role in biasing the participants above and below a mean true probability of 0.5, with higher gain control δ leading to more extreme judgments in both directions. Once again, the parameters show a clear differentiation in terms of how they affect behavior. Figure 7.12 attempts to explain the finding in the original paper that people tended to provide a higher probability of increase after longer streaks, even thought the true underlying distribution was independent of the streaks. The figure shows the average (across all participants) current η_t and persistence ρ_t signals applied towards reference

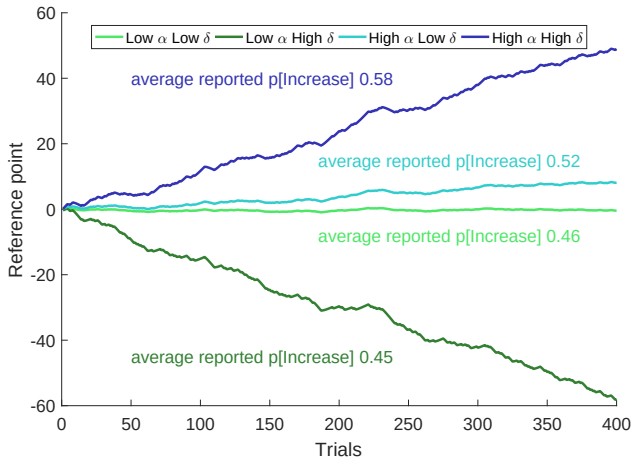

Figure 7.11: Probability tracking and extrapolation: Average movement of the reference point for 4 groups of participants, clustered based on high and low levels of α and δ

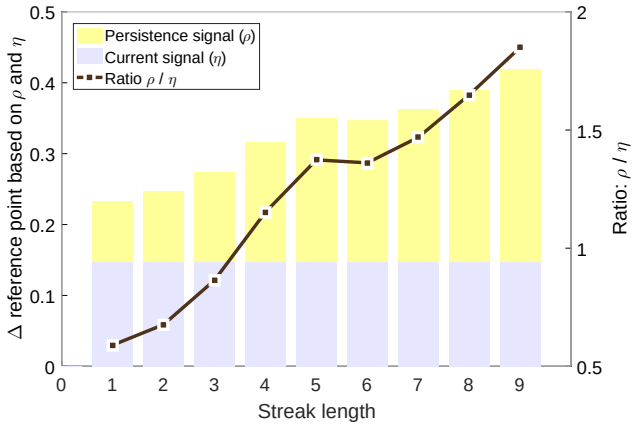

Figure 7.12: Probability tracking and extrapolation: The latent inferred average immediate and persistence signal clustered by streak length of the current trial. Increasing persistence signal capture the streak length effect, directly influencing the reference point.

point on every trial, grouped by the current streak, up to the current trial. This aggregates absolute values of both positive and negative signals, since the idea to show how streak length contributes to changes in magnitude. An upward or down movement contributing to the streak would tend to increase or decrease the reference point for the next trial. The bars in the figure show that the average current feedback signal (blue bars) are similar, independent of the streak history. The yellow bars, representing the mean persistence signal are continuously increasing as the streak length increases, showing that the decaying persistence signal captures the cumulative streak history. The black line shows the resulting ratio of persistence to current signal keeps increasing with increased

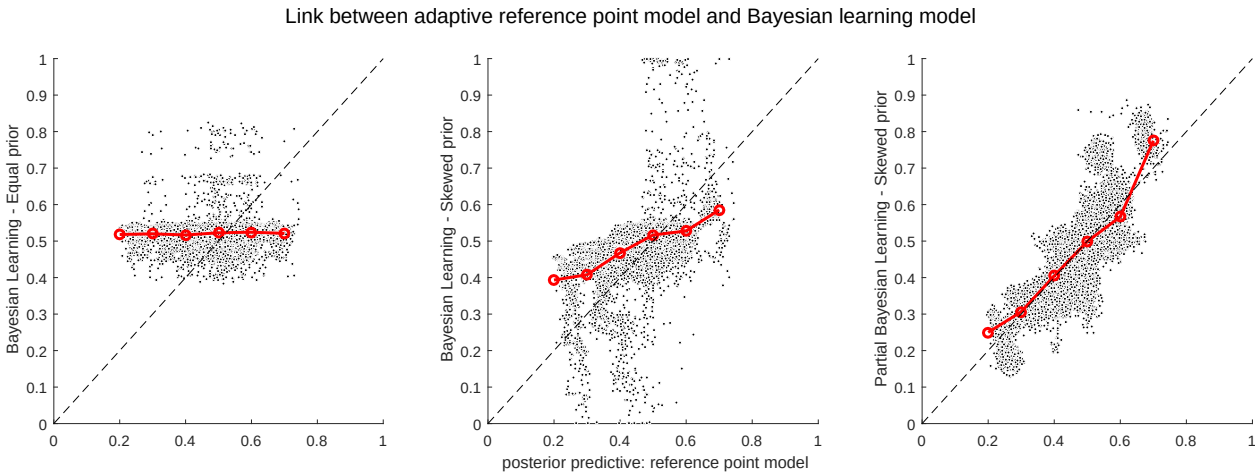

Figure 7.13: Probability tracking and extrapolation: Comparison of the posterior predictive of the adaptive reference point model with the posterior predictives resulting from fitting 3 versions of the Bayesian learning model. The baseline Bayesian learning model described, a Bayesian learning model with skewed prior expectations, anda sub-optimal Bayesian learning model (where belief updating is sub-optimal) with skewed priors. The adaptive reference pont comes closest to the third.

streak length, contributing to the observations in the original paper.

Figure 7.13 compares the posterior predictive of the adaptive reference point model with the posterior predictive resulting from fitting 3 versions of the Bayesian learning model - the baseline Bayesian learning model described earlier, a Bayesian learning model with skewed prior expectations (free parameters), and a sub-optimal Bayesian learning model (where belief updating is sub-optimal) with skewed priors. The posterior predictive from the last model and the adaptive reference point models are very similar, showing that there may be a strong representational link between these models.

7.4 Perceptual learning

We implement the adaptive SDT model specified in section 6.3.4 to analyze behavior in an experimental task.

7.4.1 Experimental Data

This data come from collaborations, with experiments reported by Skewes and Gebauer (2016), looking at differences in perceptual learning between neurotypical and autistic participants. This analysis has been presented in Mistry et al. (2018). Autism spectrum disorder (ASD) is a highly prevalent condition with about 1 in 68 affected globally. Sensory symptoms are common in ASD, and include hypo- and hyper-sensitivity to stimulus, and sub-optimalities in perceptual inference (Turi et al. (2015)). One area in which perceptual differences are particularly common in autism is auditory perception (OConnor (2012)), including auditory localization (Teder-Sälejärvi et al. (2005)). Skewes and Gebauer (2016) examined the potential cause of sub-optimality in perceptual judgments for the spatial sources of sounds in adults with ASD. In the task, on each trial, participants had to categorize an auditory stimulus into one of two categories. The categorization was based on a cover story of classifying different species of crickets, with the territory of one species being distributed to the left and the other to the right. Based on the spatial location of the sound stimulus, participants had to categorize which species the sound on each trial originated from. Each trial was followed by corrective feedback. The stimuli for the two species were spatially overlapping to some extent to introduce uncertainty into the task. Each participant completed 960 trials split into 4 randomized blocks. The 4 blocks consisted of a 2 X 2 factorial design, with each block having either a low (25%) or high (75%) base rate (BR) of one species, and a low or high discriminability (SD, standard deviation). The blocks were presented in randomized order. In the low discriminability environment there was greater spatial overlap of the auditory stimulus from the 2 species. For this task, the criterion is defined as the spatial boundary such that any stimulus perceived to come from the right of this criterion is categorized as species 1 and from the left as species 2. As a matter of convention, objective spatial locations to the right are given positive values and to the left are given negative values, so that the species on the right is considered the "signal" and on the left, the "noise".

The key results from a classical SDT analysis were that both ASD (n=19)and NT (n=23) partic-

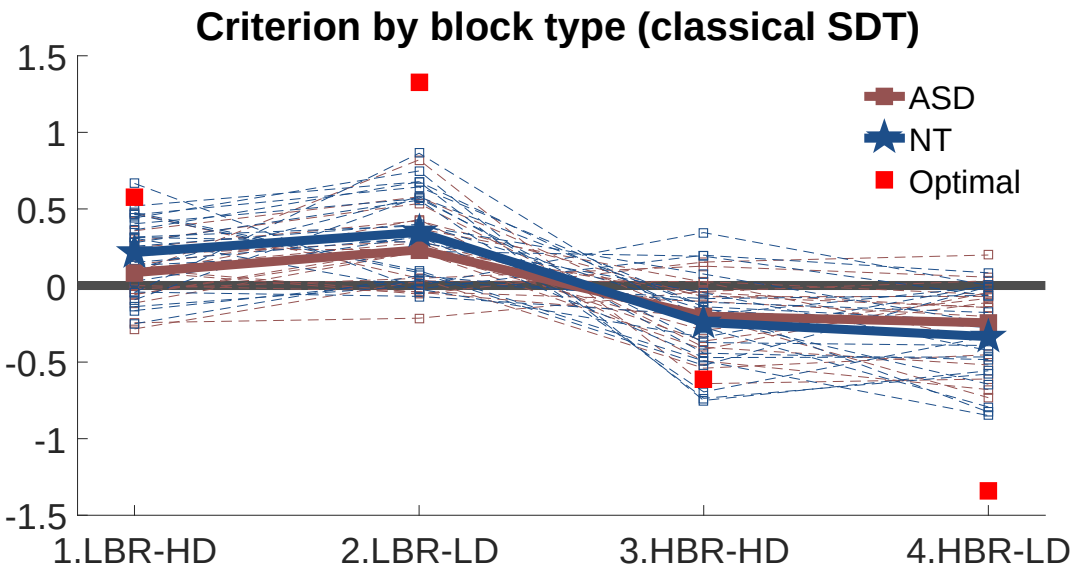

Figure 7.14: Inferred criterion based on classical SDT for individuals (dotted lines) and group means (thick lines) in the 4 blocks that vary in base rate (LBR=low; HBR=high) and discriminability (LD=low; HD=high). The red squares show the optimal criterion placement.

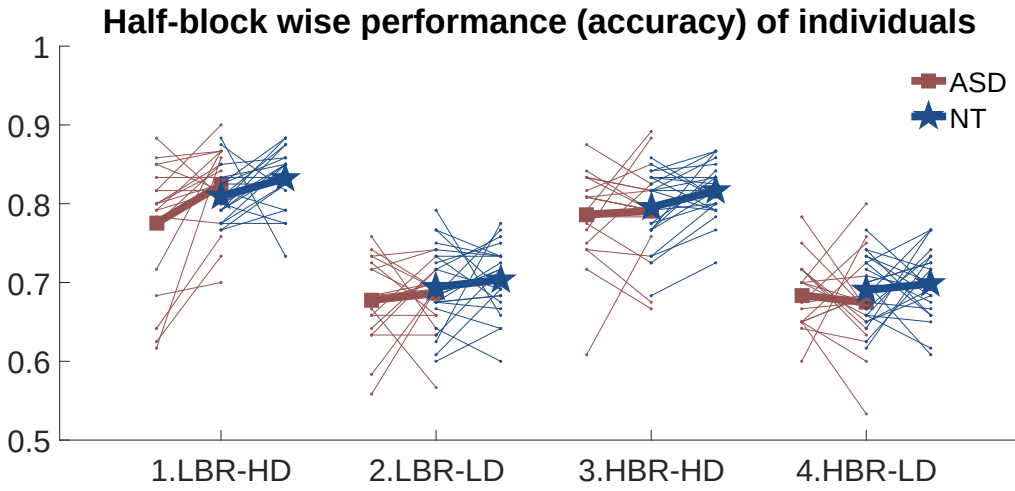

Figure 7.15: Accuracy of categorization for individuals (dotted lines) and group means (thick lines) in the 4 types of blocks. The blocks are split into two halves of 120 trials each, so the slope of the lines shows within block changes. The NT and ASD plots are displaced adjacent to each other to improve the clarity of the figure. ASD participants show greater variability and some show lower levels of performance, but the differences at a group level are very small.

ipants showed sensitivity to base rate as well as discriminability manipulations. This sensitivity, however, was suboptimal, and both groups demonstrated significant deviation from the optimal response criterion, as shown in Figure 7.14. This deviation was larger for the ASD group than for

the NT group for all 4 conditions. As a result ASD participants also demonstrated lower accuracy as shown in Figure 7.15. A one-sided Bayesian t-test JASP-Team (2016) produced a Bayes factor (BF) of 4.0 in favor of the accuracy for ASD participants (mean 73.7%) being lower than NT participants (75.5%). ASD participants demonstrated less extreme criterion values in response to base rate manipulations, but the BF for this was not conclusive. Figure 7.15 also shows the performance of individual participants divided into the first 120 and second 120 trials, for each of the 4 types of blocks. There does not seem to be a significant improvement within blocks for either group. In general, accuracy was lower in conditions with lower discriminability. The participants in the ASD group show greater variability, especially in the lower accuracy range.

7.4.2 Adaptive criterion setting

We implement the adaptive SDT model detailed in section 6.3 on this dataset. On the t^{th} trial, participants are assumed to adapt a criterion λ_t, such that if their perceived stimulus μ_t is higher than λ_t, they identify the stimulus as species 1, and otherwise, as species 2. Our adaptive SDT model (ASDT) assumes that people do not adapt a fixed criterion across all trials, but keep changing the criterion in response to feedback. Such changes would be responsive to differences in rewards, the perceived size of the error, or the past history of correct and incorrect feedback. In this task rewards are symmetric, that is, there is no difference in the rewards for correctly identifying species 1 (hit) or species 2 (correct rejection). Similarly there is no difference in the penalty depending on whether species 1 (miss) or species 2 (false alarm) was incorrectly identified. Since the two categories were fictional, there is no reason to believe that participants have an inherent bias towards either. Accordingly, ASDT groups all correct and all incorrect decisions together. It is assumed that people shift their criterion only after receiving feedback about errors.

7.4.3 Model description

Figure 7.16 shows the stimulus and criterion based on classical SDT as well as our adaptive SDT model for a single participant. The effectiveness of adaptive SDT is especially visible in the predictions in the low discriminability (LD) blocks. Figure 7.17 shows the application of our adaptive SDT model to 2 NT and 2 ASD participants. The achieved accuracy of the 4 participants and the mean values of the inferred α and δ parameters for these participants are shown on the left side. The 4 participants were selected to show behavior where both parameter are low (participant 1), low α but high δ (participant 2), high α and low δ (participant 3), and both parameters high (participant 4). The first column shows the immediate sensory error signature η_t across all 960 trials. It can be seen that participants 2 and 4, with higher values of δ, show higher η values. The second column shows the persistence related error signature ρ_t, and here, participants 3 and 4, with high values of α, show higher ρ_t values. The third column shows the sum of these two, which is what contributes to the total criterion correction on each trial. Of interest is the fact that across most trials, the persistence based feedback signature seems to be the inverse of the immediate sensory error feedback, thus leading to muted corrections when α is large. The last column shows the resulting criterion movement from trial to trial. All four participants show some sensitivity to base rate (BR) and the variability - standard deviation (SD) , but this is much higher in participants 1 and 2, who accordingly show higher accuracy rates.

7.4.4 Inference about individual parameters

We then use the complete data set to infer individual level α and δ parameters for the 19 ASD and 23 NT participants. Figure 7.18 shows the joint posterior density of the parameters for the two groups. The size of the squares is the joint probability density. The overall densities look quite similar for ASD and NT participants. The ASD group shows slightly higher values of α (mean 0.26, SD 0.26) versus the NT group (mean 0.19, SD 0.20), and similar values for δ (mean 0.4,

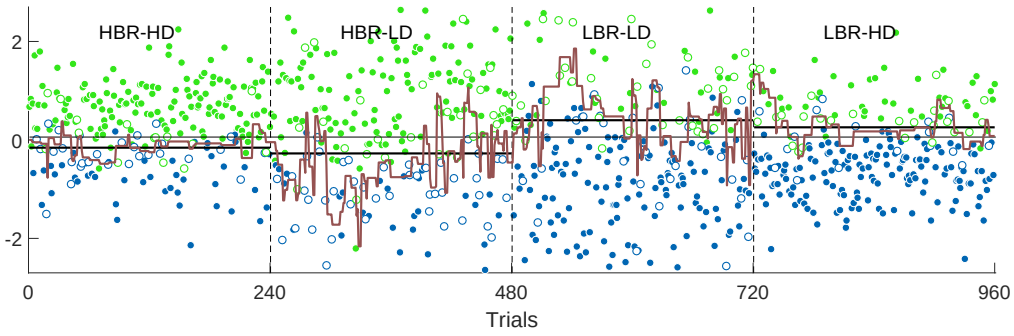

Figure 7.16: Criterion dynamics for participant 2: The brown line shows the inferred adaptive criterion inferred by ASDT. The thin black line shows the criterion based classical SDT, combining all 960 trials together. The thick black lines in each block lines show the criterion based on classical SDT computed for each block separately. The dots show the standardized stimulus values. Filled green dots show hits, filled blue dots show correct rejections. Empty green dots show false alarms and empty blue ones show misses. A model that predicts well should show green dots above the criterion, and blue dots below the criterion.

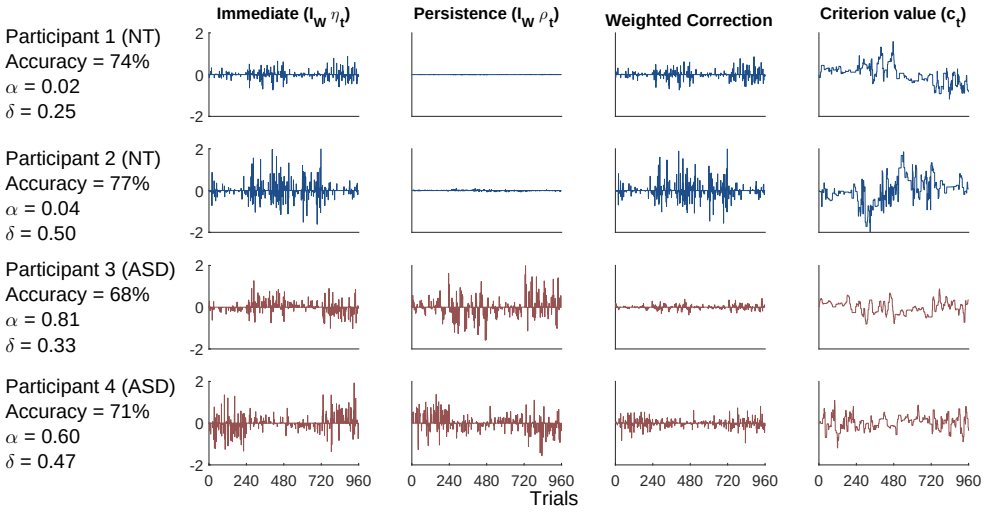

Figure 7.17: A process perspective inferred from the model for 4 of the 42 participants, to show how the adaptive model infers distinct forms of behavior. The four columns show η_t, ρ_t, $\eta_t + \rho_t$, and c_t.

SD 0.09 for ASD versus mean 0.41, SD 0.07 for NT) but neither is significant. A Bayesian t-test suggests no main effect of diagnosis (ASD vs NT) on either parameter with BFs of 0.47 and 0.31 respectively, testing for a difference between the two groups for α and δ. A Bayesian ANOVA analysis however reveals a significant main effect of the Autistic traits questionnaire (AQ) score, with a Bayes factor of 4.4. Higher AQ scores demonstrate higher values of α. In Figure 7.18 the

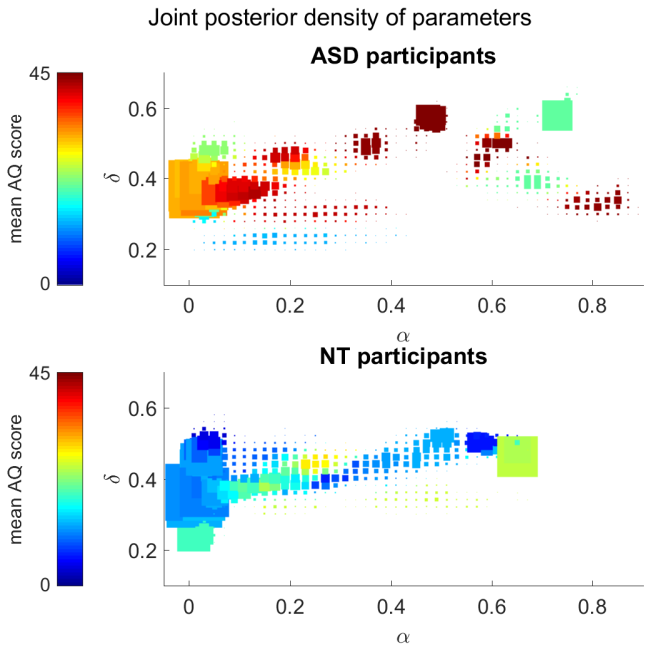

Figure 7.18: The joint posterior probability densities for the 2 model parameters for the ASD and NT participants. The size of the squares shows the probability density and the color shows the mean AQ (autistic traits questionnaire) score for the particular combination of α and δ values.

color represents the weighted AQ scores. For NT participants, this score is almost uniformly low as expected, except for the highest level so fα within NT participants. With a few exceptions, α seems to increase with an increasing mean AQ score, shown by the density clusters in dark red towards the right. A Bayesian test of correlation yields strong evidence for a negative correlation between α and the actual accuracy of participants in the task ($r = -0.54, BF154$), and mild evidence for a positive correlation between δ and accuracy ($r = 0.39, BF = 4$). This supports the notion that any suboptimality is driven primarily by higher persistence signal (α), than by the gain (δ).

7.4.5 Model performance

We implemented ASDT within a Bayesian inference framework for statistical inference (Plummer et al. (2003)). To test the model, we infer the parameters using only data from one of the blocks at a time and calculate the accuracy of the out-of-sample predictions for the remaining 3 blocks based

on the mean posterior predictives. A floor benchmark is the accuracy with which the classical SDT based criterion calculated using the hit rate and false alarm rate from a single block is able to predict the responses for the remaining blocks. Table 7.1 shows a comparison of the predictions based on using data from each of the 4 blocks for the classical SDT and ASDT models. The ASDT model provides superior predictions, and provides a psychological process perspective to explain how the criterion adapts over time.

Table 7.1: Accuracy of out-of-sample predictions using the difference blocks (LB=Low base rate; HB=High base rate; LD=Low discriminability; HD=High discriminability). SDT is based on classical SDT analysis, and ASDT is based on our proposed model of adaptive criterion setting.

	Out of sample prediction using block				
	LB-HD	LB-LD	HB-HD	HB-LD	All
Autism Spectrum Disorder (ASD)					
SDT	81.2%	82.0%	80.8%	79.9%	81.0%
ASDT	85.8%	86.0%	85.9%	86.1%	86.0%
Neurotypical (NT)					
SDT	82.1%	78.8%	78.9%	79.7%	79.9%
ASDT	87.0%	87.1%	87.3%	86.6%	87.0%

7.4.6 Aberrant precision interpretation

Suboptimality in sensory (and other) tasks by adults with ASD has been proposed to be a disorder of metacognition (Friston et al. (2013); Van de Cruys et al. (2014)). Within this framework, Lawson et al. (2014) propose two mechanisms that constitute an aberrant precision account of autism. The first is enhanced neuromodulatory gain for how prediction errors are encoded in individuals with autism. Adaptive gain control in neurotypical individuals is expected to adjust to environmental volatility so that there is higher gain in more volatile environments. It has been proposed that in individuals with autism however, gain control might be excessively enhanced because of the expectation of highly precise sensory inputs. This in turn would lead to a lack of context sensitivity, as reported by Palmer et al. (2015). Thus we conclude that the gain control processes controlled by δ in our model corresponds to this mechanism. We would thus expect to see higher values of

δ for ASD participants under this framework. We do not however observe this and δ values for both groups are strikingly similar. We propose that excess neuromodulatory gain control is not a key driver of suboptimality for the Skewes and Gebauer (2016) task. This result supports the conclusion that autism is not characterized by uniform differences in the weighting of prediction error (Manning et al. (2016)).

The second mechanism under the predictive coding framework constitutes a lack of sensory attenuation, sometimes manifested as a failure to suppress prediction errors generated by repetitive stimuli over time (e.g. Kleinhans et al. (2009)), or in failing to notice changes in the predictive value of specific information (Van de Cruys et al. (2014)). The key aspect is that individuals with autism can form accurate representations of low-level prediction errors, but the translation of these into higher level signals differs when compared to NT individuals. Specifically, the higher level signals might be influenced to drive repetitive behavior and perceive prediction errors over time in a consistent manner. This may thus lead to behavior that is more resistant to change. In our model, we propose that α captures this mechanism. High values of α would indicate persistence of sensory feedback over time, leading to increased consistency of actions and longer time frames to respond to environmental changes. We would expect to see higher values of α for ASD participants under this framework. We see some indication of this, as values of α do show a small but significant increase with increasing AQ scores. We propose that increased persistence and thus a lack of response flexibility is the key driver for any increased suboptimality observed in this pool of ASD participants. Relating this to classical SDT analysis, increased lack of response flexibility would result in an increase in deviation *along* the ROC, not necessarily demonstrating reduced sensitivity.

There is general consensus that lower level sensory error signals can be more precise, but are transformed into attenuated or less precise higher level prediction error signals in people with ASD. A perspective for explaining this has been using Bayesian updating (Pellicano and Burr (2012)). The basic idea is that individuals with ASD may demonstrate inefficient Bayesian updating since

they may have diffused priors, called hypo-priors, but strong sensory signals. We propose a related but slightly different explanation. Even if individuals with ASD start with diffused priors, updating with a strong sensory signal on a trial by trial basis would result in sharp posteriors. Since the posterior on one trial would form the basis for the prior on the next, a diffused prior would not be sustainable over trials. A sustained diffused prior might however be maintained from trial to trial if apart from a strong sensory signal, there was a second signal that also influenced these priors. On any trial, if all previous error information has been accounted for efficiently in the updated prior, Bayesian updating would require that only new information is taken into account for further changes to be made to the criterion. This is represented by the term η_t. Hence any significant contribution from ρ_t leads to interference and ineffective updating. Even if η is a sharp sensory signal, if ρ is partly in opposition to η, the result would result in sustained diffused beliefs, as have been proposed in theory. Slightly higher levels of α and the resulting higher values of the ratio of absolute magnitudes of ρ_t to η_t (mean ratio of 5.4 for ASD versus 2.9 for NT) though not statistically significant, directionally align with Pellicano and Burr (2012), who suggest that autistic perception might suffer from hypo-priors.

7.5 Price-High-Low: A consumption preference task

We implement the adaptive reference point based consumption utility model, as specified in section 6.3.5 to analyze behavior in an experimental task.

7.5.1 Data

This is a secondary dataset, obtained from the work reported in Sitzia and Zizzo (2012). 384 participants were required to make a series of 20 sequential decisions. On each trial, they had to select how many units of a particular lottery to buy. The lottery remained fixed across all trials

within a subject. However, the purchase price per lottery was varied sequentially within-subject, with the objective to test how different pricing schedules (increasing, decreasing, etc.) influenced purchasing behavior. Participants were initially endowed with experimental units of currency, and could spend as much of it as they wanted on the lotteries. At the end of the task, the unspent currency, as well as any winnings based on the lotteries, were added and converted to real monetary payouts. The leftmost panel in figure 7.19 shows the between-subject conditions. Each colored line is the price stimulus for each of the 5 conditions. Participants in each of the 5 conditions start with different, extremely high (EH), high (H), moderate (M), low (L), and extremely low (EL) levels of prices fo the first 10 trials (shape block), but all of them observe the same identical moderate price in the last 10 trials (compare block). The second subplot in the figure shows the behavior, in terms of average units bought over trials for each condition. The key observations made in the original paper are that participants with higher initial price observation purchase more units at the constant latter half price than those that has observed a lower price. The distribution of individual participants based on prices seen in the compare block for each condition is shown in the third subplot. It can be seen that the average values are higher for th ered and blue (EH and H) conditions compared to the yellow and purple (EL and L) conditions.

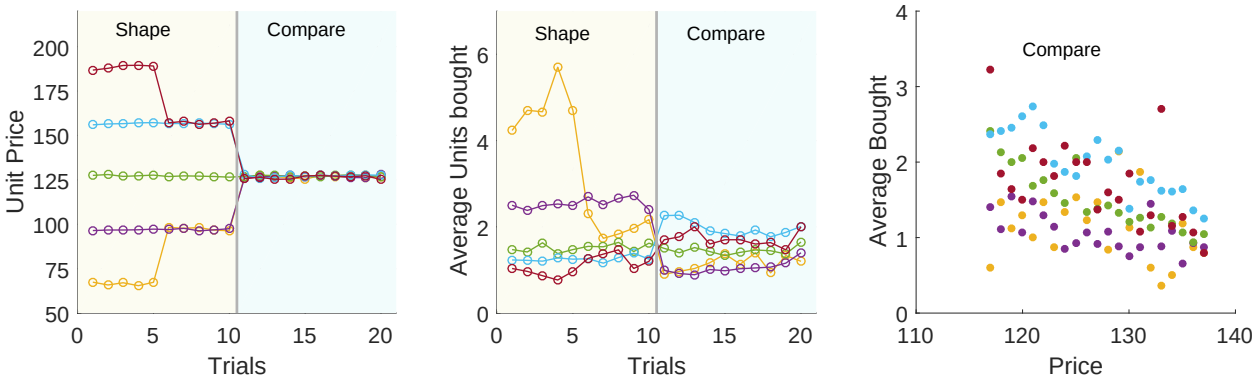

Figure 7.19: Stimulus and responses: Price-High-Low task

7.5.2 Modeling

We implement the baseline regression model for utility function based price elasticity, as specified in section 6.7 as the baseline model. This allows demand to vary based on the current value of the price.

$$log(x_t) = A + \beta\ log(p_t) \tag{7.7}$$

In addition, we implement a fixed reference point (λ_0)model of utility.

$$log(x_t) = \begin{cases} A + \beta\ log(p_t + k(p_t - \lambda_0)) & \text{if } \lambda_0 \leq p_t(1+k)/k \\ \infty & \text{otherwise.} \end{cases} \tag{7.8}$$

Finally, we also implement the adaptive reference point model of utility based on section 6.7. The adaptive reference point moves based on the last price observed, in accordance with individual level α and δ parameters. In addition, we allow a payoff bias. In case of consumption, people we propose that people have an innate consumption bias, and will be quicker to adjust their reference points when the prices increase than when they decrease. We infer the bias as a free parameter in the model. This bias (between 0 and 1) acts as a multiplier term to the difference signal in equations 6.2 and 6.3, only when the price is lower than the reference point, thus allowing for any potential consumption bias. The individual level parameter k measures the strength of anchoring on the reference point. Figure 7.20 illustrates the working of the model with one actual example. Figure 7.21 shows the model comparison, with the adaptive model showing lower rates of error against both the standard utility and fixed reference point based utility model.

$$log(x_t) = \begin{cases} A + \beta\ log(p_t + k(p_t - \lambda_t)) & \text{if } \lambda_t \leq p_t(1+k)/k \\ \infty & \text{otherwise.} \end{cases} \tag{7.9}$$

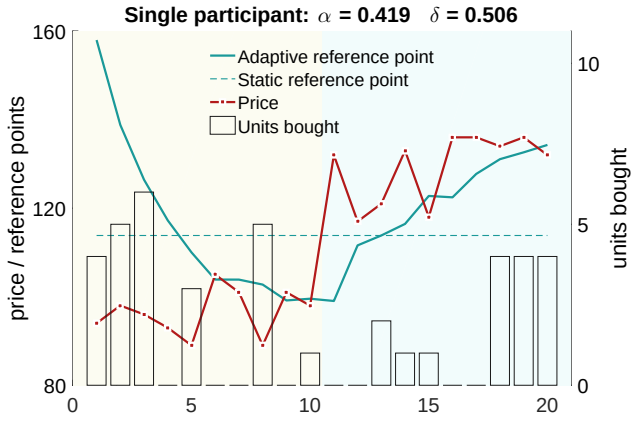

Figure 7.20: Illustration of the model working: Units bought are observed and the price is the stimulus. The inferred reference points based on the static and adaptive reference point model are also shown.

7.5.3 Inferences

Figure 7.22 shows the correlation between the latent inferred strength of anchoring and the total units bought. The strength of this correlation keeps increasing as we move from the extremely low initial price to the extremely high initial price conditions. This reflects the fact that in high price conditions, higher or lower price elasticity cannot explain higher consumption. Rather, higher consumption is only explained by a higher fixation on comparison of the actual price against a reference, and having a high reference point.

Figure 7.23 shows the distribution of inferred parameters of the adaptive reference point process. Whilst most participants are clustered near the center, some individual differences are seen, especially with higher values of δ and α in the High (H) and extra high (EH) conditions, suggesting that there may be some influence of task condition on the learning parameters.

Figure 7.24 shows that most people show a bias, with a multiplier in the range of 0.1 to 0.6 when the price is lower, reflecting that they are quicker to update their reference point after a price increase than after a price decrease. Figure 7.25 shows the average reference point movement by condition. The higher reference points in the H and EH condition also reflect a *consumption bias*,

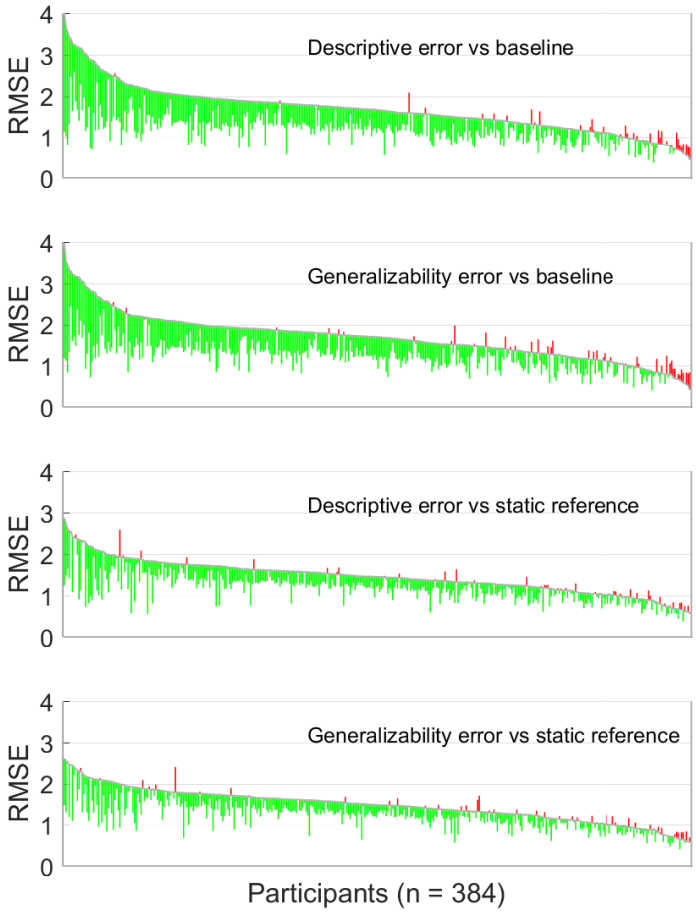

Figure 7.21: RMSE comparison for the price-high-low task: Gray lines show the RMSE of baseline models and the green and red bars show the improvement and deterioration in error using the adaptive reference point model.

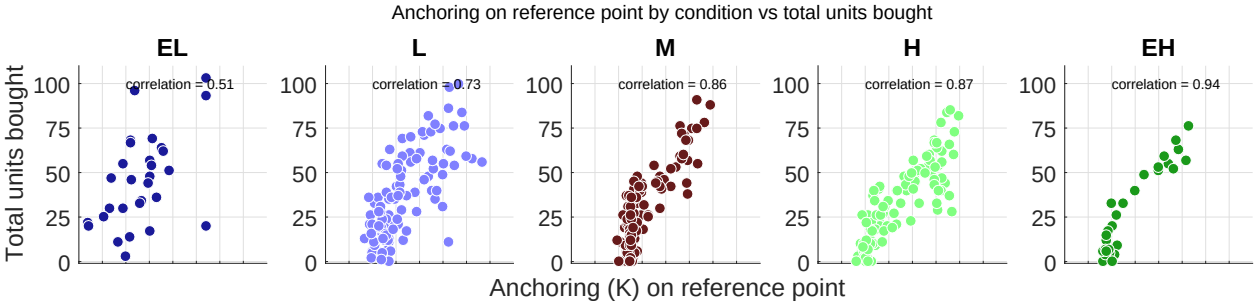

Figure 7.22: Correlation between the latent inferred strength of anchoring and the total units bought. The strength of this correlation keeps increasing as we move from the extremely low initial price to the extremely high initial price conditions.

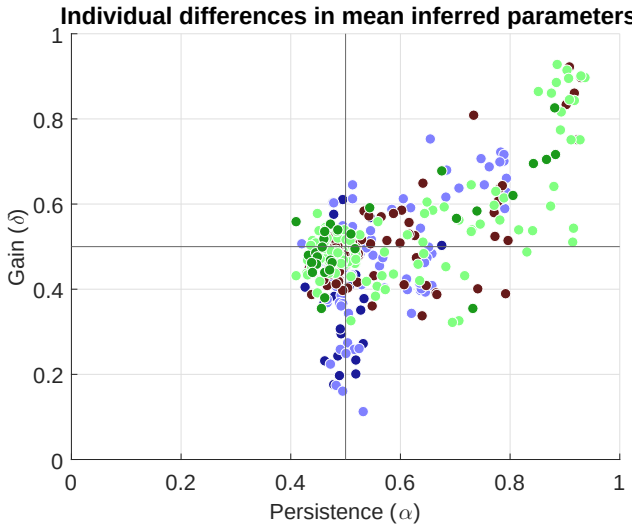

Figure 7.23: Distribution of inferred parameters of the adaptive reference point process. Whilst most participants are clustered near the center, some individual differences are seen, especially with higher values of δ and α in the High (H) and extra high (EH) conditions, suggesting that there may be some influence of task condition on the learning parameters.

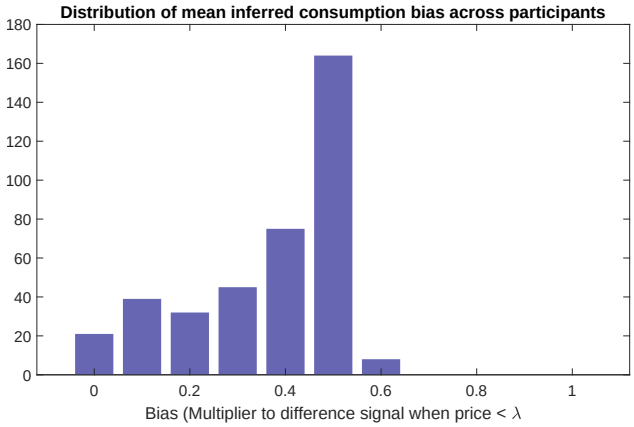

Figure 7.24: Distribution of the consumption bias parameter across participants. Most people show a bias, with a multiplier in the range of 0.1 to 0.6 when the price is lower, reflecting that they are quicker to update their reference point after a price increase than after a price decrease.

away from the mean value, since the low values of the bias parameter (0.1 to 0.6 above), make the movement from high to low reference points slower than the other way round, resulting in the higher end point reference values and a higher consumption for these conditions.

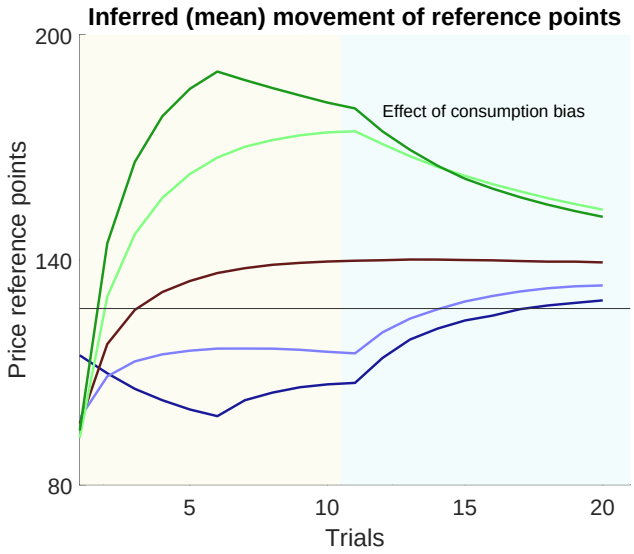

Figure 7.25: Average reference point movement by condition. The higher reference points in the H and EH condition also reflect a *bias*, away from the mean value.

7.6 Conclusions

In this chapter we have shown the ubiquity of the adaptive reference point mechanism by applying this to several different experimental tasks and incorporating them within different cognitive modeling or response frameworks.

Part III

Adaptive populations

PART III - Adaptive populations

Population level behavior, behavior that is typically measured at the macro level, is usually treated as an econometric rather than psychological or cognitive data. Such data takes the form of time series overs days, months or years, and is typically characterized by various levels of population adaptivity, often stemming from cognitive primitives.

In this part of the thesis I show how robust cognitive models can be incorporated into econometric analysis of population level time series, to account for adaptivity, provide good descriptive and predictive capabilities, and make suitable inferences about the cognitive process that drives population level behavior.

Chapter 8

Cognitive modeling of adaptive behavior in real world populations - Intifada violence

8.1 Introduction

Jeliazkov et al. (2008) analyze the daily incidence of violence during the Second Intifada using analytical Bayesian implementation of a second order discrete Markov process. They find that the data are *characterized by weak dynamics and strong instability across sub-periods, showing distinct violence patterns within each political regime*. In this chapter, we extend this work in several ways. First, we propose an agent-based cognitive model of violence, where the propensity for violence is governed by a non-homogeneous, zero-inflated Poisson process. This approach allows for a greater influence of past events than a Markov process, but imposes greater structure compared to vector autoregressive models commonly used to analyze time series. The cognitive model proposes repetition and retaliatory functions to measure the latent build up for violent behavior. These functions are similar to the reaction functions typically used in VARs, but are constrained and simplified based on assumptions drawn from psychological theory. These functions act as a

common cause, and influence both the propensity for violent behavior to occur, and the intensity of violent behavior, conditional on violence occurring. The propensity for violence based on the latent build-up is non-linear, and the effects of violence in previous periods are not linearly separable, as often assumed in linear VAR models. The model is implemented as a fully Bayesian computational model, and is not limited by any assumptions necessary for analytical approaches, including whether the time series is stationary, whether the Poisson increments are independent, or whether model residuals are normally distributed etc. The key model parameters capture the base rates of violence, the weights placed on repetitive versus retaliatory violent behavior, the recency impact for repetitive and retaliatory behavior (which measures the lag dynamics), and scale and shift parameters for the zero-inflation process. First, we implement this model to infer structural instability between sub-periods, and whether such instability is characterized by changes in the relative weights on repetition and retaliation, by the recency effects, the base rate of violence independent of previous dynamics, or the zero-threshold for initiation of violent behavior. These inferences give a deeper insight into the mechanics of violent behavior depending on the political regime shifts. The model is a generative model, and we apply the model sequentially, by providing it with last 30 days of data, and letting the model make predictions about the violence level in the next 7 days, and find that this model provides superior predictive estimates even when there are significant shifts in the underlying violence patterns.

8.2 Data

Jeliazkov et al. (2008) consider a problem and data set, involving the pattern of violence during the Second Intifada between Israelis and Palestinians, that has received considerable attention in political science and statistics. The data set (obtained from B'Tselem), basically measures the numbers of fatalities for both the Israeli and Palestinian sides over the course of over 2400 days, and divides the days into 10 meaningful periods delineated by significant political or military

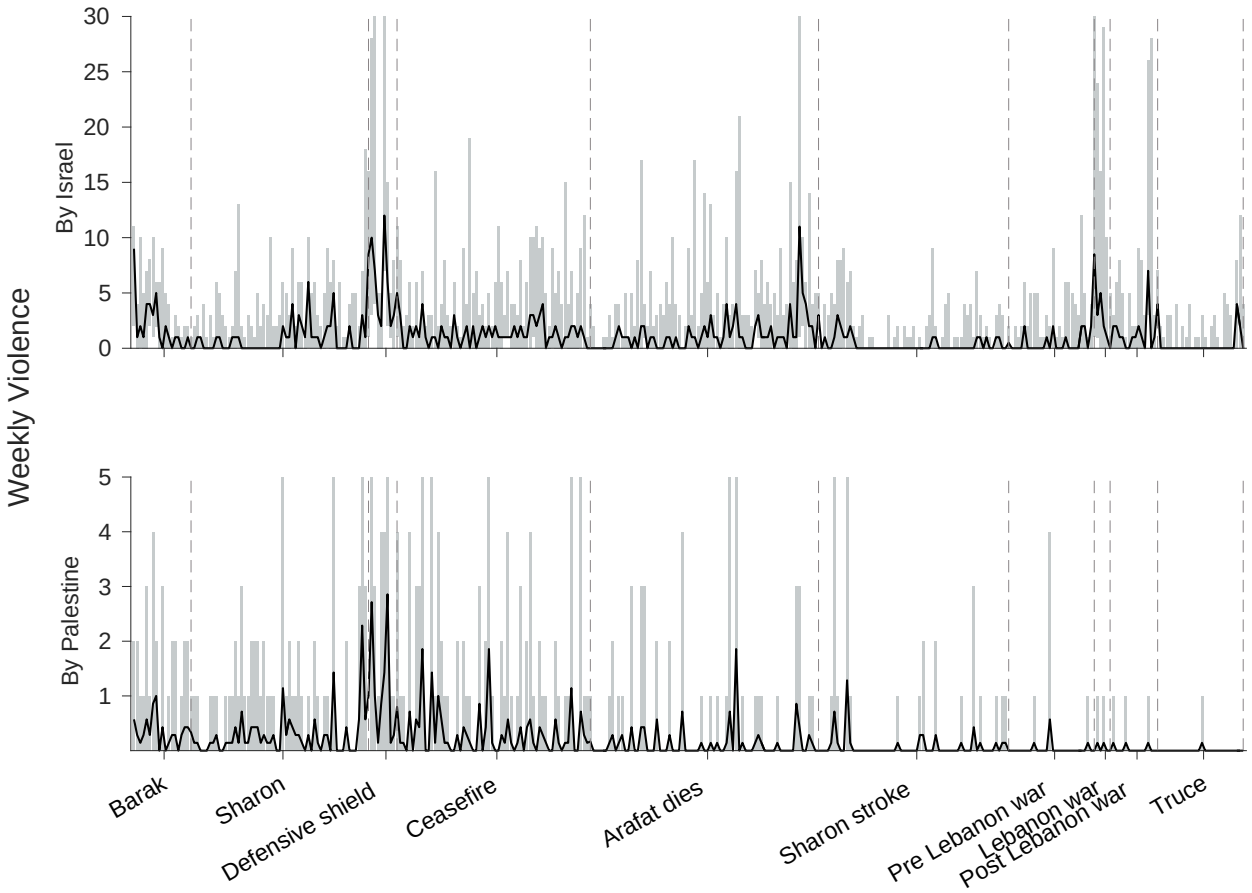

Figure 8.1: The pattern of change in weekly violence over time. The top panel shows violence by Israel against Palestine, and the bottom panel shows violence by Palestine against Israel. Bars show the range in violence over the individual days in each week, and the solid line shows the average violence for each week. The weeks are divided into ten named periods, divided by broken lines.

events.

The data take the form of counts of fatalities, $v_{s,t}$ and $v'_{s,t}$, for violence on the tth day committed, respectively, by the sth side and against the sth side. Because there are two sides, the violence committed by the first side is the violence received by the other side, so that $v_{1,t} = v'_{2,t}$. Similarly, the violence committed by the second side is the violence received by the first side, so that $v_{2,t} = v'_{1,t}$. The data also include a partition of each day into one of 10 periods, with the period for the tth day denoted by p_t.

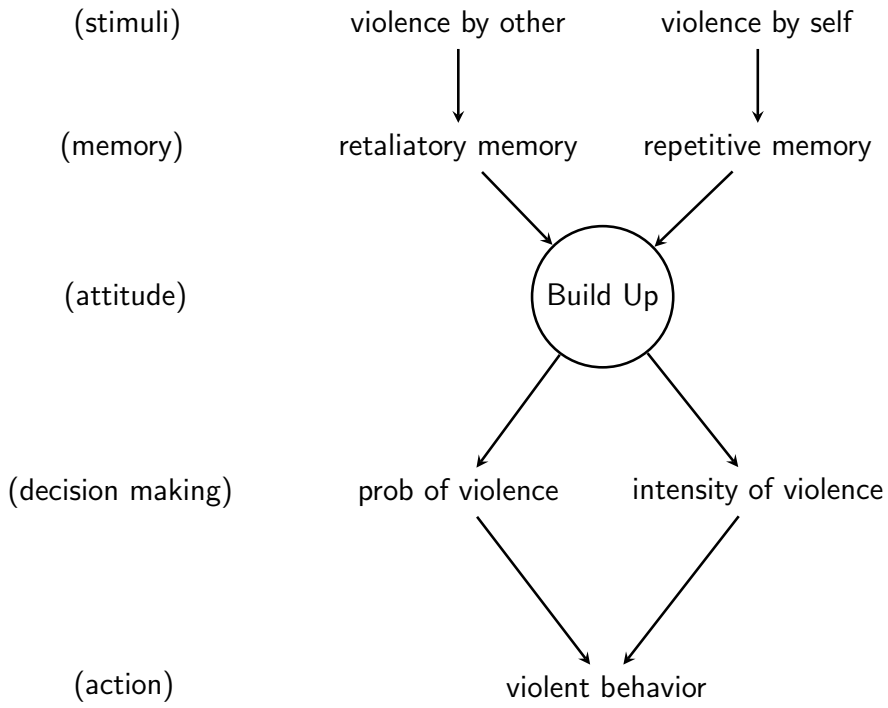

Figure 8.2: Conceptual model

8.3 Building a Psychological Model

The cornerstone of our model is a latent psychological construct, which we call *build up*, that measures the strength of the tendency toward violence for each side over time. We conceive of build up as something like a social memory, and assume it has similar properties to a human episodic memory (Norman et al. (2008)). We also assume that build up plays a role both in determining the probability that any given day will involve fatalities, and, given that it does, the magnitude of those fatalities. This sort of common-cause assumption is a hallmark of psychological modeling, in which the same latent variable is assumed to affect multiple sorts of behavioral observations (Lee (ress)). The mechanism for determining the probability of fatalities also borrows from psychological theory, taking a non-linear form often used to map latent strengths to behavioral probabilities in psychophysics (Wichmann and Hill (2001)). To describe the model, we first discuss how build up is defined and interpreted, before specifying how it determines the probability of violent attacks and the intensity of those attacks. Figure 8.2 provides a conceptual representation of this model.

8.3.1 Build Up

Build up depends on previous violence by both sides. The influence of previous violence may be positive, which aggregates further violence, or negative, which deters further violence. Further violence aggregated by previous violence by the same side is naturally interpreted as repetitive violence, while further violence aggregated by previous violence by the other side is naturally interpreted as retaliatory violence. When the influence of previous violence from the same side is negative, one interpretation is that the side is satisfied with the impact of their previous actions, and another interpretation is that the capacity for additional violence by that side has been exhausted. When the influence of previous violence from the other side is negative, one interpretation is that the capacity of the side subjected to the previous violence has been significantly diminished, and another possible interpretation is that the previous violence has acted as an effective deterrent.

Formally, the psychological build-up $\psi_{s,t}$ for an agent from side s at time t is given by a weighted combination of the retaliatory $\alpha_{s,t}$ and repetitive $\beta_{s,t}$ components:

$$\psi_{s,t} = \omega^{\alpha}_{s,p_t}\alpha_{s,t} + \omega^{\beta}_{s,p_t}\beta_{s,t}. \tag{8.1}$$

The weights ω^{α}_{s,p_t} and ω^{β}_{s,p_t} correspond to the relative importance of the repetitive and retaliatory components for the sth side during the time period p_t. The repetitive and retaliatory build ups are recency-weighted tallies of previous observed violence.

$$\begin{aligned}
\alpha_{s,t} &= \gamma^{\alpha}_{s,p_t} v_{s,t} + \left(1-\gamma^{\alpha}_{s,p_t}\right)\alpha_{s,t-1}; \quad \alpha_{s,1} = v_{s,1} \\
\beta_{s,t} &= \gamma^{\beta}_{s,p_t} v'_{s,t} + \left(1-\gamma^{\beta}_{s,p_t}\right)\beta_{s,t-1}; \quad \beta_{s,1} = v'_{s,1}.
\end{aligned} \tag{8.2}$$

The recency weights γ^{α}_{s,p_t} and γ^{β}_{s,p_t} quantifies to how persistent the influence of previous violence is on current build up. If a recency weight is close to one , only the most recent events will be

significant. As a recency weight decreases towards zero, the influence of previous events increases.

8.3.2 Probability of Violent Attacks

The probability of violent attacks π_t at any time t is based on the latent build up ψ_{t-1} and is modeled with a logistic function

$$\pi_{s,t} = \nu_{s,p_t} / \left(1 + \exp\left(\frac{-\left(\psi_{s,t-1} - \tau_{s,p_t)}\right)}{\lambda_{s,p_t}}\right)\right), \tag{8.3}$$

that includes a threshold value τ_{s,p_t}, an upper bound ν_{s,p_t}, and a scale λ_{s,p_t} for the logistic function. As the build up $\psi_{s,t-1}$ crosses the threshold τ_{s,p_t}, the probability of violence will be half the maximum possible probability ν_{s,p_t}. How quickly the probability increases or decreases as the build up increases or decreases is controlled by the scale λ_{s,p_t}. The probability $\pi_{s,t}$ is used to generate a latent indicator $\delta_{s,t} \sim \text{Bernoulli}\left(\pi_{s,t}\right)$, with $\delta_{s,t} = 1$ indicating the sth side perpetrated violence on the tth day, and $\delta_{s,t} = 0$ indicating that they did not.

8.3.3 Intensity of Violent Attacks

The intensity of observed violence is also modeled as depending on the combination of the latent build up $\psi_{s,t-1}$, and a base rate ϕ_{p_t} of violence in the period p that is independent of the latent build up. The intensity is calculated as

$$\theta_{s,t} = \psi_{s,t-1} + \phi_{s,p_{t-1}}, \tag{8.4}$$

and is the basis of modeling the observed violence as

$$v_{s,t} \sim \begin{cases} \text{Poisson}\big(\theta_{s,t}\big) & \text{if } \delta_{s,t} = 1 \\ 0 & \text{otherwise.} \end{cases} \tag{8.5}$$

8.3.4 Graphical Model

Figure 8.3 shows a graphical model representation of the model specified above.

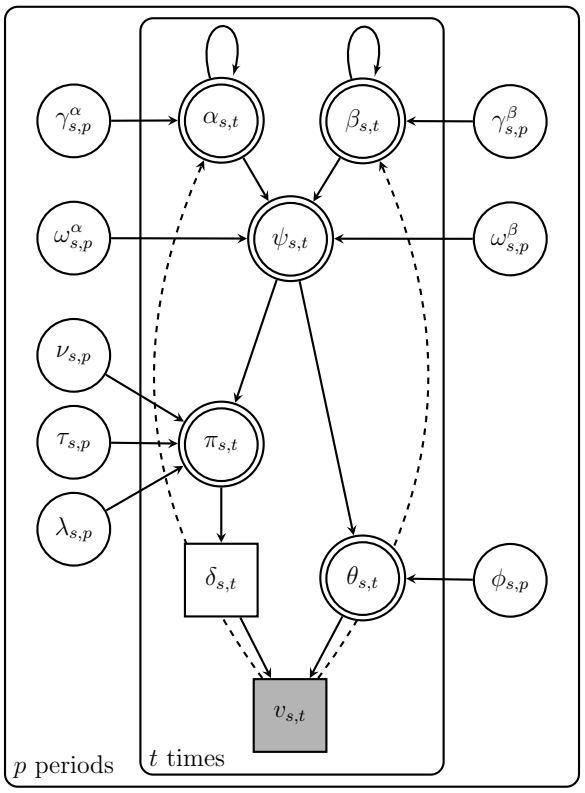

$$
\begin{array}{rcl}
\mu_s^{\gamma^\alpha}, \mu_s^{\gamma^\beta} & \sim & \text{Uniform}\big(0,1\big) \\
\xi_{s,p_t}^{\gamma^\alpha} & \sim & \text{Uniform}\big(-\mu_s^{\gamma^\alpha}, 1-\mu_s^{\gamma^\alpha}\big); \quad \xi_{s,1}^{\gamma^\alpha} = 0 \\
\xi_{s,p_t}^{\gamma^\beta} & \sim & \text{Uniform}\big(-\mu_s^{\gamma^\beta}, 1-\mu_s^{\gamma^\beta}\big); \quad \xi_{s,1}^{\gamma^\beta} = 0 \\
\gamma_{s,p_t}^{\alpha} & = & \mu_s^{\gamma^\alpha} + \xi_{s,p_t}^{\gamma^\alpha} \\
\gamma_{s,p_t}^{\beta} & = & \mu_s^{\gamma^\beta} + \xi_{s,p_t}^{\gamma^\beta} \\
\alpha_{s,t} & = & \gamma_{s,p_t}^{\alpha} v_{s,t} + \big(1-\gamma_{s,p_t}^{\alpha}\big)\,\alpha_{s,t-1}; \quad \alpha_{s,1} = v_{s,1} \\
\beta_{s,t} & = & \gamma_{s,p_t}^{\beta} v'_{s,t} + \big(1-\gamma_{s,p_t}^{\beta}\big)\,\beta_{s,t-1}; \quad \beta_{s,1} = v'_{s,1} \\
\hline
\mu_s^{\omega^\alpha}, \mu_s^{\omega^\beta} & \sim & \text{Gaussian}\big(0,1\big) \\
\xi_{s,p_t}^{\omega^\alpha} & \sim & \text{Gaussian}\big(0,1\big); \quad \xi_{s,1}^{\omega^\alpha} = 0 \\
\xi_{s,p_t}^{\omega^\beta} & \sim & \text{Gaussian}\big(0,1\big); \quad \xi_{s,1}^{\omega^\beta} = 0 \\
\omega_{s,p_t}^{\alpha} & = & \mu_s^{\omega^\alpha} + \xi_{s,p_t}^{\omega^\alpha} \\
\omega_{s,p_t}^{\beta} & = & \mu_s^{\omega^\beta} + \xi_{s,p_t}^{\omega^\beta} \\
\psi_{s,t} & = & \omega_{s,p_t}^{\alpha}\alpha_{s,t} + \omega_{s,p_t}^{\beta}\beta_{s,t} \\
\hline
\mu_s^{\lambda} & \sim & \text{Gaussian}\big(1,1\big)_{T(0.01,\infty)} \\
\xi_{s,p_t}^{\lambda} & \sim & \text{Gaussian}\big(0,1\big)_{T\left(0.01-\mu_s^{\lambda},\infty\right)}; \quad \xi_{s,1}^{\lambda} = 0 \\
\lambda_{s,p_t} & = & 1/\left(\mu_s^{\lambda} + \xi_{s,p_t}^{\lambda}\right) \\
\mu_s^{\tau} & \sim & \text{Gaussian}\big(0,1\big) \\
\xi_{s,p_t}^{\tau} & \sim & \text{Gaussian}\big(0,1\big); \qquad \xi_{s,1}^{\tau} = 0 \\
\tau_{s,p_t} & = & \mu_s^{\tau} + \xi_{s,p_t}^{\tau} \\
\pi_{s,t} & = & \nu_{s,p_t}/\left(1+\exp\left(-\left(\psi_{s,t-1}-\tau_{s,p_t)}\right)/\lambda_{s,p_t}\right)\right) \\
\delta_{s,t} & \sim & \text{Bernoulli}\big(\pi_{s,t}\big) \\
\hline
\mu_s^{\phi} & \sim & \text{Gaussian}\big(0,1\big) \\
\xi_{s,p_t}^{\phi} & \sim & \text{Gaussian}\big(0,1\big); \quad \xi_{s,1}^{\phi} = 0 \\
\phi_{s,p_t} & = & \mu_s^{\phi} + \xi_{s,p_t}^{\phi} \\
\hline
\theta_{s,t} & = & \psi_{s,t} + \phi_{s,p_t} \\
v_{s,t} & \sim & \begin{cases} \text{Poisson}\big(\theta_{s,t}\big) & \text{if } \delta_{s,t}=1 \text{ and } \theta_{s,t}>0 \\ 0 & \text{if } \delta_{s,t}=0 \text{ or } \quad \theta_{s,t}\le 0 \end{cases}
\end{array}
$$

Figure 8.3: Graphical model implementation of the model of intifada violence.

8.4 Modeling Results

8.4.1 Descriptive Adequacy

Here, we provide all the data to the model, and use this to test the descriptive adequacy of the model and make suitable inferences about the population level cognitive process. Figure 8.4 shows the violence by either side. For ease of comparison, the average weekly mean and weekly range of posterior predictive from the models (brown) have been plotted on the negative y-axis, with the actuals on the positive y-axis (gray). The descriptive adequacy of the model can be observed from figure.

8.5 Evaluating Model Predictions

We implement an $n+1$ to $n+7$ prediction, that is, we provide the model with the last 30 days of data and obtain the predictions for the *next 7 days*, the data for which is not provided to the model. This is done consecutively for all days. We thus obtain model generalizability predictions for lags from $n+1$ to $n+7$. The figures 8.5 and 8.6 show the mean posterior predictive values of the violence based on $n+1$ and $n+7$ values respectively. The RMSE error of the predictions increase from 2.3 to 2.9 as we move from an $n+1$ to $n+7$ prediction.

One important aspect of the modeling predictions are that in the context of the data, a miss (predicting no violence when there is going to be violence) is costlier than a false alarm (predicting violence but none occurs). We test the model against a simple *repeat yesterday* model, where the assumption is that if violence occurred on the previous day, it will occur today, and vice versa. Note that the data is heavily zero-inflated, and this simple heuristic is expected to do quite well on an overall level. Our model is slightly worse in terms of false alarms (30% versus 20% for the *repeat yesterday* assumption), but significantly better in terms of the misses (20% versus 47% for

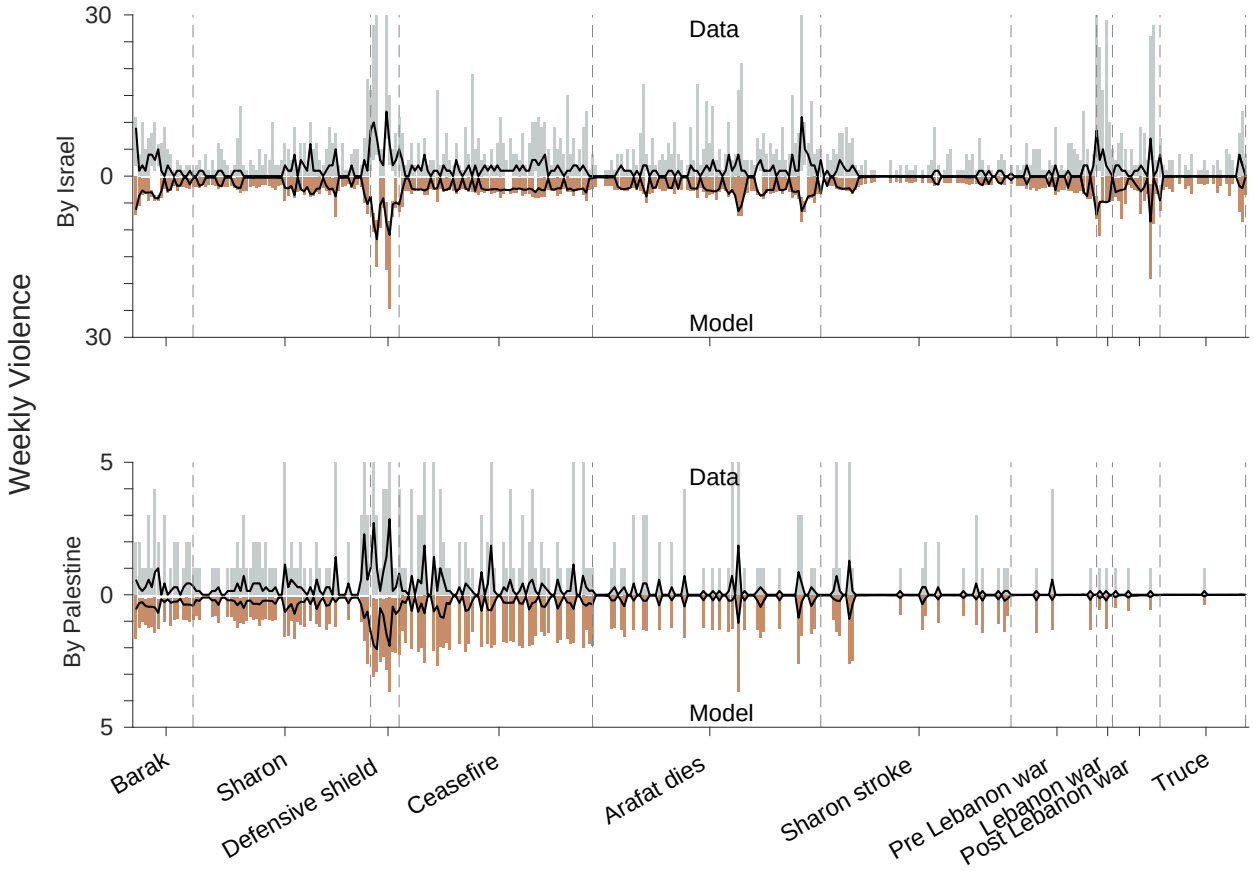

Figure 8.4: Descriptive adequacy of the model. The top panel shows violence by Israel against Palestine, and the bottom panel shows violence by Palestine against Israel. The weeks are divided into ten named periods, divided by broken lines. Within each panel, the upper half shows the data, with bars showing the range in violence over the individual days in each week, and the solid line showing the average violence for each week. The bottom half shows the range and mean of the corresponding posterior predictive distribution of the model.

the *repeat yesterday* assumption), which we argue, are the critical points for prediction.

8.5.1 Inferences

Figure 8.7 shows how the repetitive and retaliatory components of the latent build up decay over time. For most periods, the influence of past violence, as far as the dynamic component of violence is concerned, is wiped out within a week. There are some notable exceptions, for instance in the case of Israel, during the period when Arafat dies, and during the pre Lebanon war. The figure

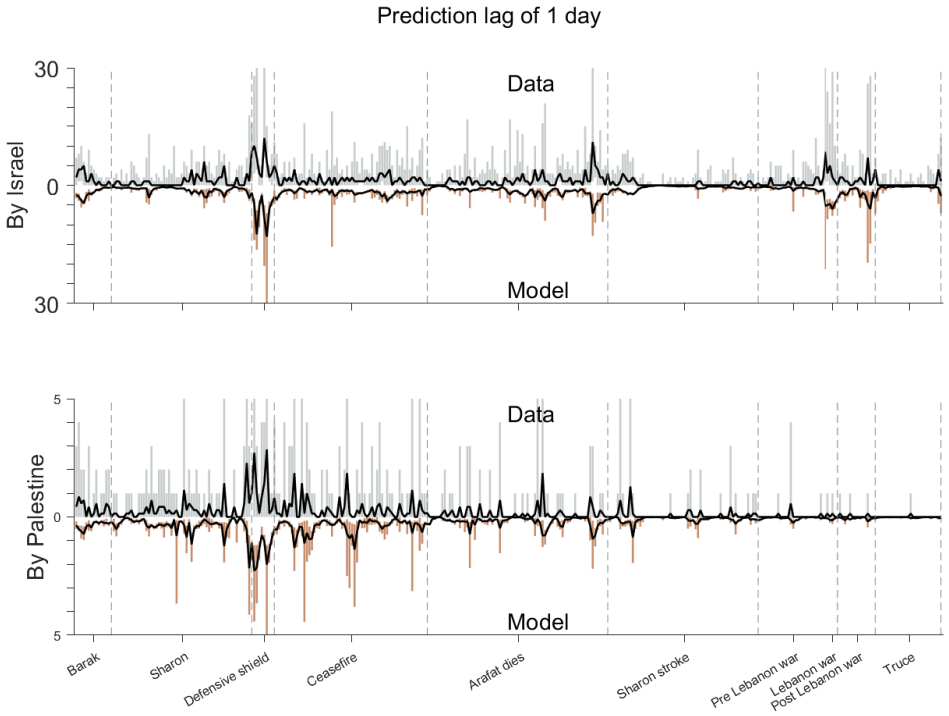

Figure 8.5: $n+1$ prediction

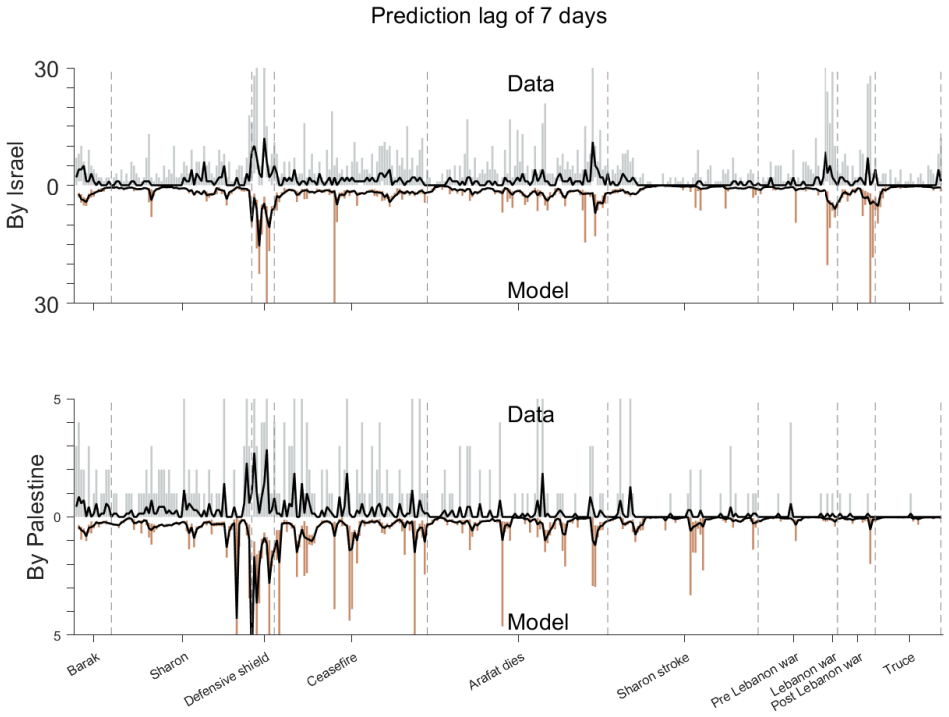

Figure 8.6: $n+7$ prediction

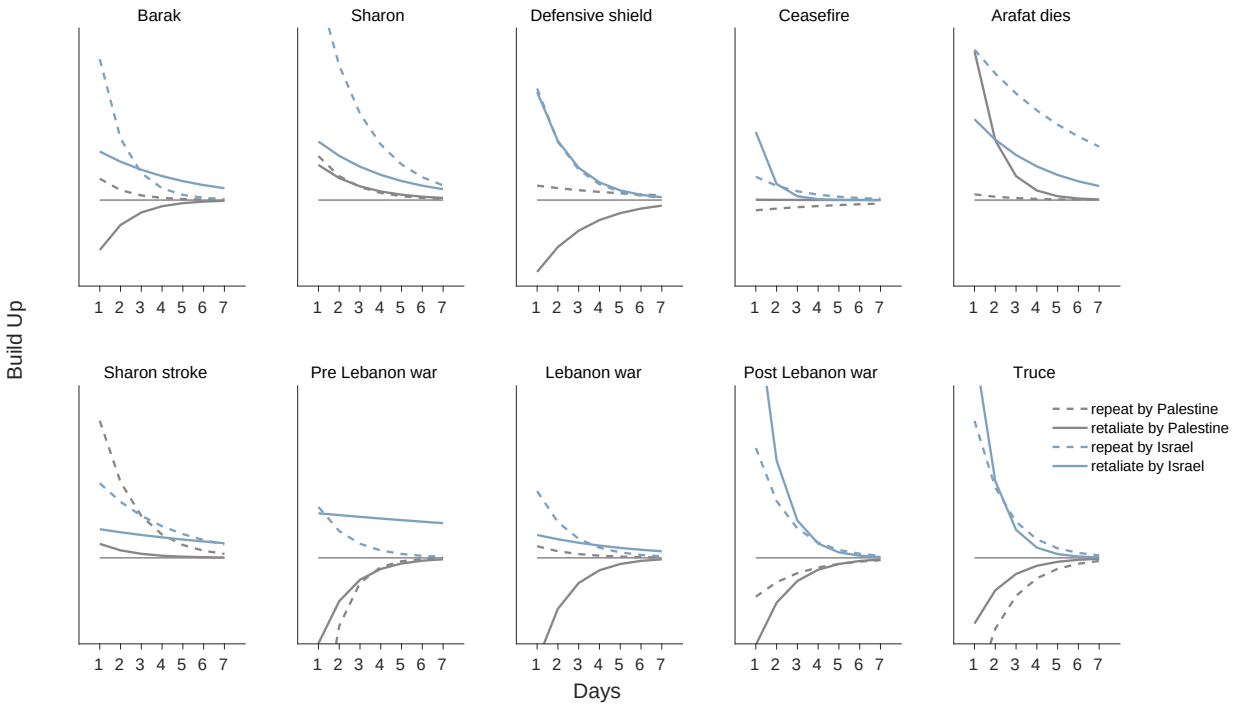

Figure 8.7: Shows how the repetitive and retaliatory components of the latent build up decay over time. The influence of these components is shown with increasing lag as we move from the left (1 day later) to the right (7 days later).

also shows how retaliatory influence switches between positive to negative, or vice versa, between certain periods, for Palestine. For instance, the retaliatory influence is positive during the Sharon period, but becomes significantly negative during the Defensive shield period, which was a period with a goal of capacity destruction by Israel. Conversely, there is a huge increase in the retaliatory component in the period where Arafat dies. In general, the ceasefire (4th time period) demonstrated stronger dilution of the dynamic repetitive and retaliatory responses than the truce period (10th time period).

Figure 8.8 shows the *mean* build up of the retaliatory and repetitive components in each period (note this varies on a day-to-day basis in the model). The key observations are that the retaliatory buildup seems to be more influential on violence by Palestine, and the repetitive buildup for Israel. The negative build in latter periods for Palestine may be indicative of capacity reduction based on attacks from Israel. Figure 8.9 shows the characteristic curves inferred for probability of violence

versus latent build up during each period. Some key observations are the generally higher sensitivity for Israel, the sudden shifts in the curves, for example a drop in the period form Defensive shield to Ceasefire, and an increase in the period from the pre Lebanon war to the Lebanon war, are captured by the model from the data, without any prior knowledge about these periods.

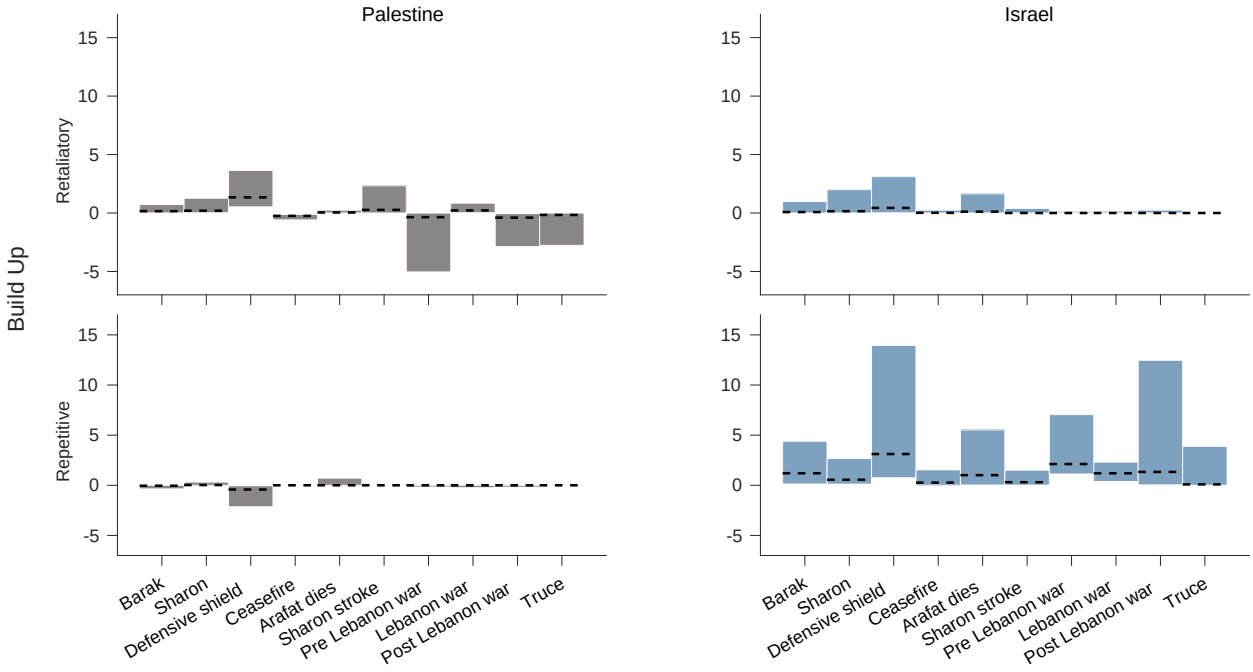

Figure 8.8: Average build up of the retaliatory and repetitive components in each period.

The cognitive modeling approach reveals relatively weak dynamics, that is, the role of day to day violence compared to the long term internalization of violence. Figure 8.10 shows the distribution of the ratio of latent build-up to the base rate in determining the intensity of violence. High values imply that short term dynamics play a large role, whereas low values imply a long term build up that is internalized, and not affected by day to day violence. In general, retaliatory violence by Palestine shows a higher degree of dynamics, whereas repetitive violence by Israel shows a higher degree of dynamics. The impact of specific events such as a jump in the repetitive dynamics for Israel during the war and the during Operation Defensive Shield, and the sudden fall during the truce period are once again captured by the model, just from the data, without any prior information about the nature of these periods.

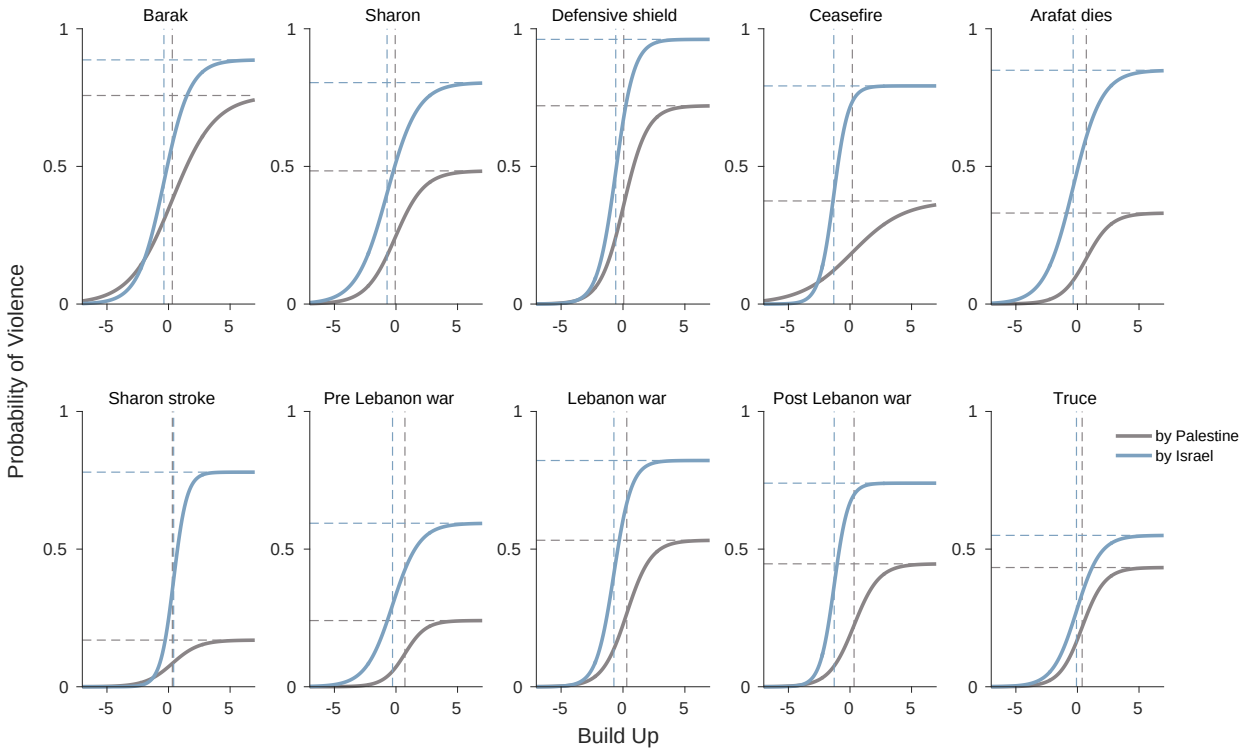

Figure 8.9: The probability of violence as a function of latent buildup on either side.

8.6 Conclusions

Vector autoregressive (VAR) models are commonly used to model interdependent time series in econometric studies. These models allow for interpretation of the time series based on impulse response functions, which measure the reaction of target variables to impulse shocks in one or more of the input variables, and historical and predictive decomposition of error variances based on these impulse response functions. These functions can be analytically determined because the change in the target variable as a result of a change in one of the input variables is linearly separable in VAR models. Unlike the VAR models, the model specified here does not need to make any assumptions about the time series (such as stationarity). The repetition and retaliation functions here are similar to the reaction functions typically used in VARs (Jaeger and Paserman (2006); Haushofer et al. (2010)), but are constrained and simplified based on assumptions drawn from psychological theory. This psychological perspective means that, while our model ultimately can be conceived as a set of

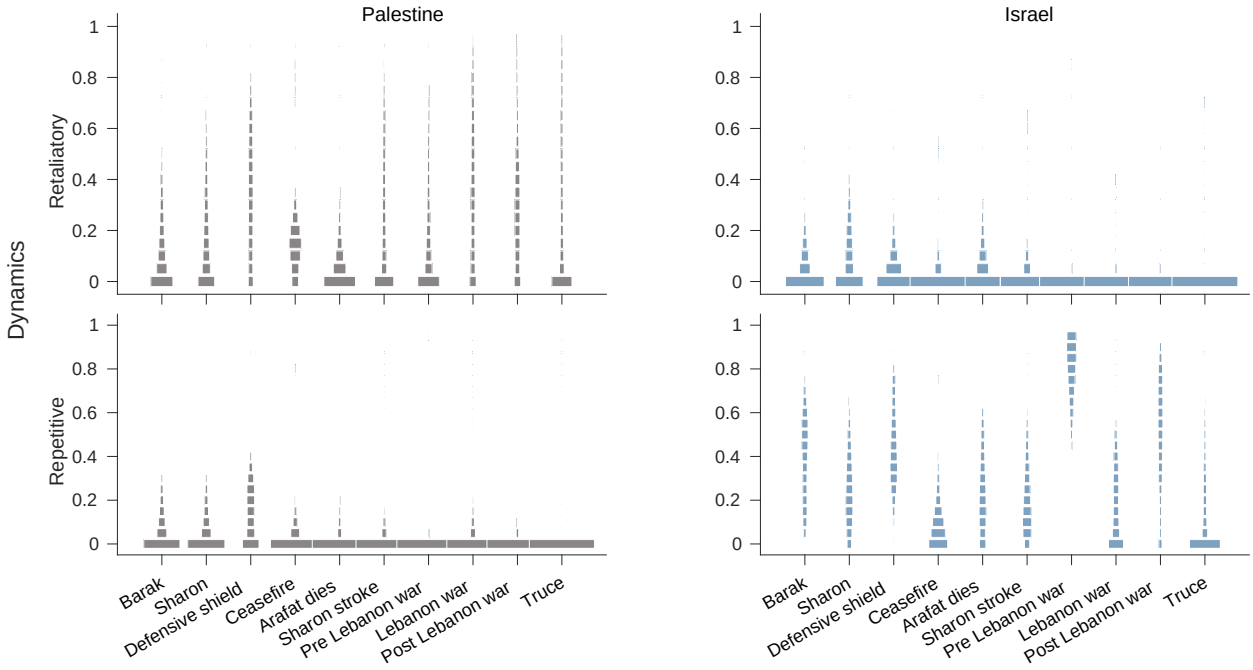

Figure 8.10: Distribution of the ratio of latent build-up to the base rate in determining the intensity of violence. High values imply that short term dynamics play a large role, whereas low values imply a long term build up that is internalized, and not affected by day to day violence.

statistical assumptions, it makes theoretical commitments and has parameters with interpretations that are more grounded in psychological theory. We argue that this approach leads to a model that is both statistically simpler, and psychologically more structured, than previous statistical approaches, and especially vector autoregressive (VAR) modeling approaches. For example, VAR modeling potentially allows a free parameter quantifying the weight given to each previous time, although previous authors have imposed various statistical constraints, often by choosing a finite window that limits the time-lag over which previous violence can exert an influence. Our approach, in contrast, makes a theoretical assumption about constraints on social memory that lead to a single interpretable recency-weighted parameter for each period controlling the influence of past violence on current build up. Overall, our model provides suitable descriptive and predictive capability, and allows us to make strong inferences about the underlying process.

Chapter 9

Cognitive modeling of adaptive behavior in real world populations - Alcohol consumption patterns in response to tax changes

9.1 Introduction

Consumption behavior is often analyzed by identifying sufficient statistics based on the observed elasticities of consumption behavior with respect to underlying drivers of change, such as prices or tax rates. The underlying assumption in rational economic analysis is often that these elasticities, for instance, the rate of change of consumption in response to tax changes, are stable at a population level over time. If instead, these elasticities are dependent on evolving choice preferences, the resulting analysis may be highly biased or inefficient. For instance, assume that the change in consumption ∂W in response to a change in policy ∂t (e.g. a change in the tax rate) is assumed to be determined by an empirically observable elasticity ε, so that $\partial W/\partial t = f(\varepsilon)$. Once ε is estimated, it can be used to measure and predict the change in consumption resulting from all future changes in t. This is the typical assumption in many econometric models.

Now assume that ε is in fact a function of some constant cognitive preference parameters Ψ, and environmental factors E. If instead we can empirically measure Ψ, then $\partial W/\partial t = f(\varepsilon) = f(g(\Psi, E))$. This has the advantage that any changes in ε due to changing preferences can be captured through a relatively constant cognitive parameter set Ψ. It also has the advantage that environmental factors E, that might affect elasticities because of the way cognitive preferences react to specific environmental changes can be captured. However this is contingent on specifying a structural model $g(\Psi, E)$. Since elasticities are an emergent property of consumption functions, I propose that a generic model of consumption that captures changing preferences through a set of constant cognitive parameters may be a viable solution.

In this chapter, I present a generic cognitive framework for modeling consumption decisions by individuals. This framework is based on segregated mental accounting and the existence of a psychological transaction utility that works to complement or offset economic utility. I model transaction utility as being reference point sensitive. Reference points in turn are modeled as adaptive psychological constructs that evolve over time based on hedonic adaptation, show susceptibility to

confirmation bias, and in the domain of public policy response, are influenced by perceptions of fairness and trust in the government. In essence, this is an adapted version of the adaptive reference point based utility models introduced in section 6.3.5.

Note that this chapter is primarily theoretical, and provides an analytical approach to integrating cognitive models within econometric analysis. However, a small computational modeling exercise with real world beer consumption data is included towards the end as a proof of concept.

9.2 Developing a cognitive-econometric utility framework

I propose a framework to evaluate behavioral responses to public policies based on embedding a cognitive process perspective within econometric analysis. I construct any response to public policy as a consumption problem faced by a representative agent (Chetty (2015)). This formulation can be quite flexible, for instance, as shown later in the paper, tax evasion behavior can also be modeled within this framework, by treating the proportion of income reported as the numeraire. This section describes a stepwise development of the framework.

9.2.1 Mental accounting and transaction utility

Thaler (1999, 2008) proposed that consumption choices were driven by a combination of acquisition utility and transaction utility. Acquisition utility is the utility of consumption less that value of the price paid. Transaction utility typically reflects the value of the deal, and is the difference between the price paid and a reference price. The transaction utility can be positive if the actual price is lower, and negative if it is higher, than the reference point. The mental accounting theory proposes that the acquisition and transaction utilities are *separately* evaluated.

In the generic consumption problem, the decision maker selects a consumption quantity, x, to maximize the total utility $T(x)$, which comprises an acquisition utility and a transaction utility. In the context of responses to public policy, transaction utility is specified as a segregated economic evaluation of the relevant policy tool t, that impacts the decision maker (such as taxes, subsidies, defaults etc.), against an internal reference point t_θ. Often, the exact impact of the policy may not be transparent or salient to the decision maker, hence the appraisal of transaction utility is made based on the *perceived* value of t, denoted as $t_s = st$, a multiple of the true value. Here, $s = 1$ indicates an accurate perception of the true value t. The reference point t_θ reflects the personal beliefs of the individual about the appropriate level of t. For example, it can reflect beliefs about the fair level of taxes, or about social insurance entitlements, etc. The total utility is defined in equation 9.1. Here, $V(z)$ is a money metric utility function, p is the pre-policy cost per unit of x, and δ is the weight placed on the transaction utility. The transaction utility is based on applying the policy t (say taxes), on some $f(x)$. The function $f(x)$ depends on how the policy tool is applied. For instance, in the case of ad valorem sales taxes, the tax rate t is applied on the total purchase cost, and $f(x) = px$. Note that applications of the mental accounting framework typically assume $\delta > 0$. This implies that individuals act in self-interest to maximizes their own economic position, an approach that is also the norm in neoclassical economics. Thus, if the perceived value of t_s is less than the internal reference point t_θ, the individual experience a positive utility, reflecting the value of a 'good deal'. On the other hand, if $t_s > t_\theta$, this is evaluated as paying too much, and induces a negative transaction utility. The money-metric utility $V(z)$ may be concave (e.g. $V(z) = z^\gamma, 0 < \gamma < 1$), the transaction utility may be under or over weighted ($\delta \neq 1$), or may depreciate over time ($\delta(n) < \delta(n-1)$, where transaction utilities may be repeatedly evaluated over multiple time periods). These, amongst many others aspects of mental accounting, can constitute a violation of fungibility, and can thus influence choices.

$$T(x) = [V(U(x)) - V(px)]_{acquisition} + \delta\,[V(t_\theta f(x) - t_s f(x))]_{transaction} \tag{9.1}$$

For the rest of the paper, we define $V(z) = z^{\gamma}$ if $z \geq 0$ and $V(z) = -(-z)^{\gamma}$ if $z < 0$. This is similar to the utility function as per cumulative prospect theory Tversky and Kahneman (1992), but without loss aversion, that is $\lambda = 1$ under prospect theory. I highlight that we are dealing with purchase and transaction costs rather than losses. It has been proposed that psychological coding of negative values as costs rather than losses implies that these quantities are not treated in a loss aversive manner Kahneman and Tversky (1984); Thaler (2008, 1999), to ensure hedonic efficiency. Thus, $\lambda_{costs} << \lambda_{loss}$, and it is reasonable to assume that $\lambda_{costs} \approx 1$.

9.2.2 Deontological and Utilitarian transaction utilities

The previous section defines transaction utilities in a strictly self-serving sense. However, many consumption decisions involve conflicting considerations. For instance, the decision to purchase goods that damage the environment, the decision to evade taxes, or the decision to purchase mandatory health insurance, may result in a conflict between economic utility on one hand, and a moral obligation on the other. Dual process theories of moral judgment (Greene (2007,0)) propose that when facing a moral conflict, people may make decisions based either on utilitarian or deontological premises. Utilitarian responses are based on evaluating the overall consequences to society, whereas deontological responses are based on following a rule or obligation, without rationalizing the decision based on potential consequences. Conway and Gawronski (2013) were able to dissociate not only mechanisms for utilitarian and deontological responses, but also a third propensity where neither drive behavioral responses. However, most decisions giving rise to conflict, that are studied under this paradigm, involve no direct costs or benefits to one's own self, unlike the examples of consumption based conflicts described above. In all of these examples, the self-serving utility may differ significantly from a utilitarian perspective that is based on overall social welfare. I propose that for such consumption problems, there are three possible cognitive processing modes that an individuals may adopt.

First, a self-interest based response considers transaction utilities defined in the traditional sense. This implies that the psychological benefit of paying less than the reference price is always positive, or $\delta > 0$ (see equation 9.1).

Second, it is possible that decisions are based on a deontological perspective, that is, dominated by a rule, obligation or ideology, without consideration to the transaction utility. In this case, $\delta = 0$. Psychologically, this means that considerations over whether or not something is fair, beneficial, or harmful do not enter decision considerations. These decisions tend to default to normatively expected behavior.

Third, decisions could be made based on a utilitarian perspective, with the objective of maximizing social welfare, even if this is partly in conflict with maximizing one's own utility. This corresponds to $\delta < 0$. For example, purchasing goods that are harmful to the environment may invoke a utilitarian process. In this case, it is possible that paying more than the reference point actually contributes to positive psychological utility, and helps justify the purchase. Similarly, getting a 'good deal' on such a harmful product may actually reduce transaction utility, thus inducing guilt and reducing consumption. Transaction utility under a utilitarian perspective operates in the opposite direction to a self-interest based transaction utility.

Transaction utility thus depends on whether a self-interest based ($\delta > 0$), deontological ($\delta = 0$), or utilitarian ($\delta < 0$) cognitive process underlies the decision. This of course raises the question as to what aspects determine what underlying process dominates a particular consumption decision. The dual process theory of moral judgments proposes that deontological processes are more likely to be driven by emotion and intuitive responses, whereas utilitarian processes are driven by greater reflection and cognitive control (Paxton et al. (2014)). Increased cognitive load has been found to significantly alter utilitarian judgments whereas manipulation of empathic concern influenced deontological judgments (Conway and Gawronski (2013)). From a public policy perspective, an important premise here is that deontological processes do not consciously consider the fairness or benefits of a policy while responding to it. Deontological processes might be more susceptible

to influence by default options, and peer group influences. Utilitarian processes would be specifically concerned with the fairness of policies and redistribution goals, and trust in the governing parties. Utilitarian reference points thus reflect *'the bare minimum I should be paying'*, whereas self-interest reference-points reflect *'the maximum they should be charging me'*. Reference points become irrelevant under deontological processes.

9.2.3 Hedonic and Trust based adaptation of reference points

Transaction utilities may be evaluated positive or negatively against a reference point based on what underlying processes drive responses. However this reference point is not typically constant. I propose that the reference point evolves over time, based on hedonic adaptation (Frederick and Loewenstein (1999)). A similar principle has been used to describe how people may adapt a reference point for the rate of taxation (Bernasconi et al. (2014)). Essentially, this mechanism works by moving the current reference point (t_θ) closer to the current perceived value (t_s), at a certain rate L_t. Additionally, I propose that the reference point for policy costs (benefits) moves higher (lower) when there is a perception of fairness or trust in the government, and lower (higher) when there isn't. This perception of trust is coded as $F = 1$ or $F = -1$. The rate of adaptation in response to these perceptions is governed by L_f. This characterizes equation 9.2. Here, the reference point is updated at every time period n. The updating is assumed to happen every time there is a consumption decision or a policy change. For instance, the reference point for the purchase of goods may be updated every time the goods are purchased, or the reference point for income tax may be updated every time a tax return is filed.

$$t_\theta^{(n)} = t_\theta^{(n-1)} + L_t(st - t_\theta^{(n-1)}) + L_f F^{(n)} \tag{9.2}$$

This can be rewritten as a recurrence equation 9.3, where t_0 is the initial reference point at $n = 0$. We define $\alpha^{(n)}$ as the proportion of time periods up to n where $F(n) = 1$, as opposed to $F(n) = -1$.

Then $\alpha^{(n)}$ measures the proportion of time periods during which the individual had a positive perception of fairness and trust in the system. The sensitivity to trust, $L_f > 0$ when t is a cost (e.g. taxes) and $L_f < 0$ when t is a benefit (e.g. subsidies).

$$t_\theta^{(n)} = st + (1 - L_t)^n (t_0 - st) + L_f (2\alpha^{(n)} - 1)n \tag{9.3}$$

The hedonic adaptation serves to *condition* people towards the current levels of t. The rate at which this belief changes may vary significantly by individual. Note that a hedonic learning rate $L_t = 1$ implies that an individual will update their belief about the reference point (t_θ) instantly after a single time period (e.g. a single purchase), to their current perceived level st. A low value of L_t implies that people stick to their initial beliefs. If people learn fast (high L_t), and have an accurate perception $(s \approx 1)$, this implies that their transaction utility will be close to zero, and consumption behavior will indicate that individuals tend to ignore the taxes or subsidies in question. If L_t is low, and people believe in a low fair reference point $(t_\theta << st)$, the transaction utility will have a high negative value under self-interest based processing ($\delta > 0$), resulting in a lower consumption demand. If this is the only mechanism affecting the reference point (i.e. if $L_f = 0$), as long as there is some learning ($L_t > 0$), people are expected to converge (as $n \to \infty$) their beliefs to the extant perceived value st. When there is a shift in costs $(\Delta t \neq 0)$, initially people are likely to have larger valence of transaction utilities (positive or negative depending on δ and the direction of the cost shift). Over time and with repeated transactions, they may update their reference point so that the valence of transaction utilities starts to reduce towards zero, modulated of course by any perceptions of trust. Of note is that this shift is relative to $s\Delta t$, highlighting the role of the salience s of changes. Soliman et al. (2014) showed the importance of learning in how individuals dynamically evolve to comply with tax laws, and observe that this process of learning over time is often the reason that empirical observations of a dynamically evolving population behavior do not converge with the static predictions of many neoclassical models of tax compliance. The framework provided here specifically addresses this issue as it predicts a

changing rate of compliance behavior over time.

9.2.4 Confirmation bias and asymmetric hedonic adaptation

The rate of hedonic adaptation can affect how behavior evolves dynamically over time in response to policy changes. Apart from heterogeneity across individuals, there may also be systemic variations in L_t within individuals.

$$t_{\theta}^{(n)} = st + (1-(mv+(1-v))L_t)^n(t_0 - st) + L_f(2\alpha^{(n)} - 1)n \tag{9.4}$$

I propose that the rate of hedonic adaptation L_t is asymmetric, and depends on whether movement of the reference point towards the perceived value of t supports or inhibits current behavior. For instance, if the transaction utility for a consumption decision is reduced by the direction of adaptation, the rate of adaptation is likely to be slower than if the direction of adaptation increased transaction utility. Thus, confirmation bias will manifest as a *consumption bias*, an emergent property of asymmetric reference point adaptation. Let m reflect a bias that reduces the rate of adaptation $0 \leq m \leq 1$ when adaptation inhibits current behavior. This bias is introduced in equation 9.4.

Here, v is an indicator function that serves to identify whether or not the bias applies in a particular situation. The indicator $v = 1$ if ($st < t_{\theta}$, $\delta > 0$), or if ($st > t_{\theta}$, $\delta < 0$), and 0 otherwise. These situations reflect the self-interest and utilitarian modes when a hedonic adaptation of the reference point would reduce transaction utility, resulting in a lower effective rate mL_t of adaptation.

9.3 Incorporating the cognitive framework into econometric analysis

Chetty (2009) describes how sufficient statistics form a compromise between structural models and reduced-form approaches in econometric analysis. I use the suggested approach to evaluate whether sufficient statistics for the proposed cognitive framework can be generated as functions of empirically observable elasticities. I outline 2 key examples based on consumption (sales) tax, and income tax. In subsequent analysis, I treat δ as a discrete parameter $\delta \in [1, 0, -1]$, so that it only carries information about the mode of processing (self-interest, deontological, and utilitarian), although more sophisticated analysis can allow this to vary on a continuous spectrum. The key idea here is to develop sufficient statistics for these structural models that allow us to create an estimator for key parameters of interest, and compare this to existing sufficient statistics developed within this domain.

9.4 Application 1: Consumption tax

Chetty et al. (2009b); Chetty (2015), describe a representative-agent model to incorporate behavioral aspects into a consumption decision, specifically, the impact of whether sales tax is included or excluded from the consumption decision. The key sufficient statistic is an estimator for θ, which defines the proportion of agents that are assumed to include sales tax into their consumption decision. I approach the same problem using the cognitive framework described above, making the following assumptions:

$x =$ The number of units of the numeraire

$U(x) =$ The utility function for $x = ax^{1-b}/(1-b)$

$t =$ Actual sales tax

$s =$ Salience of sales tax

t_s = Perceived sales tax = st

t_θ = Internal reference point for sales tax

p = Actual pre-tax unit price of the numeraire

$f(x) = px$ = Total purchase price, the value on which tax is applied.

Substituting the above values in equation 9.1, we obtain equation 9.5. Since $V(x)$ is dependent on whether $x > 0$ or $x < 0$, we add an indicator function η. Here $\eta = 1$ if $st \geq t_\theta$ and $\eta = 0$ if $st < t_\theta$.

$$T(x) = \left(\frac{ax^{1-b}}{1-b}\right)^{\gamma} - (px)^{\gamma} + \delta\,(px)^{\gamma}\,(1-2\eta)\,((1-2\eta)(t_\theta - st))^{\gamma} \tag{9.5}$$

Since we assume that a decision maker maximizes $T(x)$, we evaluate x at $\partial T(x)/\partial x = 0$. This gives us equation 9.6.

$$x = \frac{(1-b)^{(1-\gamma)/b\gamma}}{a^{-1/b}}\; p^{-1/b}\left(1 - \delta(1-2\eta)((1-2\eta)(t_\theta - st))^{\gamma}\right)^{-1/b\gamma} \tag{9.6}$$

Taking logs, defining $\beta = -1/b$ and $A = \beta\; log(a(1-b)^{(\gamma-1)/\gamma})$ we obtain equation 9.7 which is in the form of a log-linear model with price elasticity $\varepsilon_p = -\beta$.

$$log(x) = A + \beta\; log(p) + \beta\; log\left(\left(1 - \delta(1-2\eta)((1-2\eta)(t_\theta - st))^{\gamma}\right)^{1/\gamma}\right) \tag{9.7}$$

We seek to understand how x varies with t. The key quantity to estimate empirically is t_θ. If we can estimate t_θ within a *steady state* environment, that is, when reference point adaptation is minimal, then a sensitivity analysis of welfare is possible by using informative priors about s (from surveys), and δ. We compare empirical strategies under this model with similar strategies used in previous literature.

9.4.1 Empirical manipulation 1:

Now consider an empirical scenario where the sales tax is completely included in the displayed price (Chetty et al. (2009b); Chetty (2015)). In the presented model, this corresponds to:

$p = p(1+t)$

$s = 0$

$t_\theta = 0$

Here, the transaction utility does not come into the picture, and the acquisition utility is assumed to include the tax that is incorporated into the displayed sales price. In this scenario, the consumption quantity x_1 is given by equation 9.8.

$$log(x_1) = A + \beta\ log(p) + \beta\ log(1+t) \tag{9.8}$$

Equation (9.8) - (9.7) gives,

$$log(x_1) - log(x) = \beta\ log(1+t) - \beta\ log\left(\left(1-\delta(1-2\eta)((1-2\eta)(t_\theta - st))^\gamma\right)^{1/\gamma}\right) \tag{9.9}$$

Defining $\rho_1 = -(log(x_1) - log(x))/log(1+t)$,

$$\rho_1\ log(1+t) = \beta\ log\left(\left(1-\delta(1-2\eta)((1-2\eta)(t_\theta - st))^\gamma\right)^{1/\gamma}\right) - \beta\ log(1+t) \tag{9.10}$$

$$(1+t)^{\gamma(1-\rho_1/\varepsilon_p)} = 1-\delta(1-2\eta)((1-2\eta)(t_\theta - st))^\gamma \tag{9.11}$$

$$(t_\theta - st) = (1-2\eta)\left(\frac{1-(1+t)^{\gamma(1-\rho_1/\varepsilon_p)}}{\delta}\right)^{1/\gamma} \tag{9.12}$$

Applying the fact that t is small, we can use a Taylors series approximation so that $(1+t)^k = 1+kt$. Note that here δ is either +1 or -1, since a reference point is undefined when $\delta = 0$, that is, under deontological processing. When $\delta = 0$, we would expect $\rho_1 = \varepsilon_p$.

$$(t_\theta/t) = s + (1-2\eta)\left(\frac{-\gamma t^{1-\gamma}(1-\rho_1/\varepsilon_p)}{\delta}\right)^{1/\gamma} \tag{9.13}$$

Now note that in the special case where $\gamma = 1$ and $s = 1$, we obtain equations 9.14 and 9.15. Note that ρ_1 and ε_p are empirically observable quantities under such a manipulation Chetty et al. (2009b).

If $\delta = 1$, that is, for self-interest based processing,

$$(t_\theta/t) = (1-\eta)\left(\frac{\rho_1}{\varepsilon_p}\right) + \eta\left(2 - \frac{\rho_1}{\varepsilon_p}\right) \tag{9.14}$$

If $\delta = -1$, that is, for utilitarian processing,

$$(t_\theta/t) = \eta\left(\frac{\rho_1}{\varepsilon_p}\right) + (1-\eta)\left(2 - \frac{\rho_1}{\varepsilon_p}\right) \tag{9.15}$$

Note that under these constraints $\gamma = 1$, $s = 1$, and $\eta = 0$, the estimator for t_θ/t in the self-interest mode is ρ_1/ε_p. This is identical to the estimator for $1-\theta$ obtained by Chetty et al. (2009b), where θ is the proportion of the population that takes taxes into account in a rational sense. In fact, under these constraints, the mode of processing δ, drives the relationship between ρ_1 and ε_p. A deontological decision maker ($\delta = 0$) will always have $\rho_1 = \varepsilon_p$. A utilitarian decision maker ($\delta = -1$) will have $\rho_1 \leq \varepsilon_p$, and a self-interested decision maker ($\delta = 1$) will have $\rho_1 \geq \varepsilon_p$.

Comparing the inference made by Chetty et al. (2009b), if under such an empirical manipulation, $\rho_1 \approx \varepsilon_p$, this implies a value of $\theta = 0$ in their model, which suggests that everyone ignores taxes. In our model this corresponds to either a deontological mode of processing, where the transaction utility is not computed, or to a self-interest mode of processing ($\delta = 1$) with $t_\theta = t$, or to a utilitarian model of processing ($\delta = -1$) with $t_\theta = t$. In all three cases, the qualitative interpretation remains the same, that people have a transaction utility of zero, and thus implicitly behave as if they are ignoring taxes.

If on the other extreme, if we find that $\rho_1 \approx 0$, this implies $\theta = 1$ in their model, which in turn implies that everyone takes taxes into account in the same way as they do prices, and accordingly adjust consumption. In the model presented here, this situation gives rise to interesting constraints. If t are costs (i.e. taxes, $t \geq 0$), this situation implies that people are not in either a deontological or self-interest processing mode. This situation can only manifest if people are in a utilitarian processing mode, and either $t_\theta = 2t$ or $t_\theta = 0$. Behaviorally, this implies that the taxes are quite different from the reference point, that people ascribe either a large positive or negative transaction utility to taxes, and is thus compatible with the view that people take taxes into account in a way similar to prices.

Thus the sufficient statistic identified by Chetty et al. (2009b) is subsumed by the model presented

here, under the special case $s = 1$, $\gamma = 1$. However, we highlight that the sufficient statistics described here potentially enable us to gain additional information about the cognitive approach that drives behavior.

9.4.2 Empirical manipulation 2:

Next, consider a scenario where it is ensured through prior education or surveys that people have knowledge of the true tax rate t, such that $s = 1$. In this scenario, the consumption quantity x_2 is given by equation 9.16.

$$log(x_2) = A + \beta\ log(p) + \beta\ log\Big(\big(1 - \delta(1-2\eta)((1-2\eta)(t_\theta - t))^\gamma\big)^{1/\gamma}\Big) \tag{9.16}$$

Equation (9.16) - (9.7) and defining $\rho_2 = -(log(x_2) - log(x))/log(1+t)$ gives equation 9.17.

$$(1+t)^{-\rho_2\gamma/\varepsilon_p} = \frac{1 - \delta(1-2\eta)((1-2\eta)(t_\theta - st))^\gamma}{1 - \delta(1-2\eta)((1-2\eta)(t_\theta - t))^\gamma} \tag{9.17}$$

We note that from equation 9.12

$$(t_\theta - st) = (1-2\eta)\left(\frac{1-(1+t)^{\gamma(1-\rho_1/\varepsilon_p)}}{\delta}\right)^{1/\gamma} \tag{9.18}$$

Using this, and the fact that t is small allows using the approximation $(1+t)^k = 1+kt$, we get equation 9.19.

$$(1+\rho_2\gamma t/\varepsilon_p)(1+\gamma t(1-2\eta)(1-\rho_1/\varepsilon_p)) = 1 - \delta(1-2\eta)((1-2\eta)(t_\theta - t))^\gamma \tag{9.19}$$

For self-interested 1d decision makers, $\delta = 1$,

$$(t_\theta/t) = \eta\left(\frac{\rho_1+\rho_2}{\varepsilon_p} - \frac{\rho_2 t}{\varepsilon_p}\left(1 - \frac{\rho_1}{\varepsilon_p}\right)\right) + (1-\eta)\left(2 - \frac{\rho_1-\rho_2}{\varepsilon_p} + \frac{\rho_2 t}{\varepsilon_p}\left(1 - \frac{\rho_1}{\varepsilon_p}\right)\right) \tag{9.20}$$

For utilitarian decision makers, $\delta = -1$,

$$(t_\theta/t) = (1-\eta)\left(\frac{\rho_1-\rho_2}{\varepsilon_p} - \frac{\rho_2 t}{\varepsilon_p}\left(1-\frac{\rho_1}{\varepsilon_p}\right)\right) + \eta\left(2 - \frac{\rho_1+\rho_2}{\varepsilon_p} + \frac{\rho_2 t}{\varepsilon_p}\left(1-\frac{\rho_1}{\varepsilon_p}\right)\right) \tag{9.21}$$

Equations 9.20 and 9.21 provide a more refined estimator for t_θ/t that still assumes $\gamma = 1$, but allows s to vary. In fact, with both empirical manipulations, where ρ_1, ρ_2, and ε_p are observable quantities, it is possible to get an estimator for s based on combining these equations with equation 9.12.

$$s = (t_\theta/t) + \delta(1-2\eta)\left(1-\frac{\rho_1}{\varepsilon_p}\right) \tag{9.22}$$

Note that if $s = 1$, then $\rho_2 = 0$, in which case, equations 9.20 and 9.21 reduce to equations 9.14 and 9.15. This highlights the face that even if $\gamma = 1$ is a reasonable assumption, if $s \neq 1$, the estimate for t_θ/t obtained as a sufficient statistic based on simply the first empirical manipulation will be biased, dependent on the true value of s.

9.4.3 Implications for Pigovian taxes

Consider Pigovian taxes, for instance, on consumption of goods that bring personal utility but can harm the environment. Higher taxes on such goods may reduce consumption for a while, but eventually people get conditioned to these tax levels, if $L_t > 0$ (see equation 9.4). Alternatively, consider providing subsidies for consumption of an environmentally friendly substitute. Assuming perfect information ($s = 1$), this corresponds to a negative tax t. If initially the reference point $t_\theta = 0$, this implies $st < t_\theta$ for the subsidies. Once again, if $L_t > 0$, people may eventually get conditioned to these subsidies. For self interested decision makers, if $L_t(taxes) > L_t(subsidies)$, the subsidies are likely to be more effective in shifting demand away from the harmful products, and vice versa. The model directly predicts that this should be the case because of the confirmation bias inherent in hedonic adaptation. Interestingly, for utilitarian decision makers, higher taxes may

increase transaction utility for potentially harmful goods, rationalizing the decision to purchase. Confirmation bias would work in the reverse direction for utilitarian decision makers, $L_t(taxes) < L_t(subsidies)$, but this would not constrain consumption of harmful goods effectively.

Thus, the model predicts that subsidies on substitutes might have a more prolonged impact on limiting consumption of harmful goods compared to increased taxes on the harmful goods. There is evidence (Wächter et al. (2009)) to show that the effect of reward and punishment have distinct neural substrates, and that reward and punishment engage separate motivational systems. These differences are an emergent property of the model presented here.

A resulting implication is that subsidies may provide a more sustained route to reducing consumption of environmentally harmful goods. However subsidies worsen fiscal deficit as compared to environmental taxes that improve the fiscal position. The decision to tax or subsidize needs to balance the potential reduction in harmful consumption with the different in cost budgets required to implement these, and whether reinvestment of the environmental taxes can contribute towards other environmental goals. Rather than prescribe one mode over the other, the takeaway from the predictions of the model should be to take into account the asymmetric nature of impact for comparable policy actions.

Villas-Boas et al. (2016) measured the change in consumption patterns of bottled-water following an increase and subsequent drop in consumption tax, levied from an environmental perspective. They find low elasticity of bottled water to environmental taxes, thus leading to increased revenue but meeting only limited success in reducing harmful consumption. This aligns with the qualitative predictions of our model. The study reports that there is a drop after the increase in tax, and an increase in consumption after the tax increase was rolled back, however this increase is lower than the initial fall, and consumption did not go back to initial levels (controlling for other factors). Self-interested decision makers are likely to show a larger drop after the initial tax change, and an even larger increase after roll back. That this was not the case indicates a significant proportion of utilitarian decision makers.

Fletcher et al. (2015) share similar findings where imposing large taxes on sodas in order to improve health results in increasing revenue, but does not seem to positively influence health. This again provides evidence for confirmation bias in hedonic adaptation.

9.5 Application 2: Income tax

Elasticity of taxable income has often been considered as a sufficient statistic for the welfare analysis of changes in marginal tax rates. However recent analysis shows that this may not be the case (Saez et al. (2012); Doerrenberg et al. (2014)). I show how viewing tax compliance through the lens of the structural consumption model presented here throws light on this issue. Sensitivity of reported taxable income to marginal tax rates may comprise many factors, including labor choices and tax evasion choices, amongst others. A comprehensive welfare analysis should incorporate all of these aspects. As proof of concept, I focus here specifically on the decision to evade taxes. Let x be the proportion of income evaded. Welfare analysis would need a sufficient statistic to estimate $\partial x/\partial t$. Here, I frame the consumption problem with the numeraire, x, being the proportion of income, I, that the taxpayer wants to *evade*. For the purpose of this paper, we assume that the *perceived* probability of getting caught (say via an audit or other means), is proportional to the % of income hidden x, and to the total income level I. This assumption reflects the fear that evading a greater share of income increases the probability of being caught, and that people with higher income levels are more likely to be audited. Thus the perceived probability of being caught is defined as kxI, where k is an individual level scaling factor. Let t be the actual applicable average tax rate, and $r = ct$ be the penalty rate, $c \geq 1$ applied on the under-reported income, if caught. The utility $V(U(x))$ of reporting $x\%$ income is given by equations 9.23 and 9.24. The utility is evaluated as the utility of paying tax only on the reported income, $[-(1-x)tI]$, if not caught, with a perceived probability of $1-kxI$, and the utility of having to pay tax on the reported income plus

penalty tax on the unreported income, $[-(1-x)tI - cxtI]$, if caught, with a perceived probability of kxI. The total utility $T(x)$ is given by equations 9.25 to 9.26. There is no purchase cost involved in the transaction. Since $V(z)$ is dependent on whether $z > 0$ or $z < 0$, I add an indicator function η. Here $\eta = 1$ if $st \geq t_\theta$ and $\eta = 0$ if $st < t_\theta$.

$x =$ Proportion of income that is evaded

$I =$ Income level

$t =$ Actual average income tax rate

$s =$ Salience of income tax rate

$t_s =$ Perceived income tax rate$= st$

$t_\theta =$ Internal reference point for income tax

$f(x) = I(1-x) =$ The value on which the tax rate is applied for transaction utility calculation

$$V(U(x)) = -(1-kxI)((1-x)tI)^\gamma - (kxI)((1-x)tI + cxtI)^\gamma \tag{9.23}$$

$$V(U(x)) = \left(kxI((1-x)^\gamma - (1-x+cx)^\gamma) - (1-x)^\gamma\right)t^\gamma I^\gamma \tag{9.24}$$

$$T(x) = V(U(x)) + \delta\, V(t_\theta f(x) - t_s f(x)) \tag{9.25}$$

$$T(x) = \left(kxI((1-x)^\gamma - (1-x+cx)^\gamma) - (1-x)^\gamma\right)t^\gamma I^\gamma + \delta((1-x)I)^\gamma(1-2\eta)((1-2\eta)\ (t_\theta - st))^\gamma \tag{9.26}$$

For the special case $\gamma = 1$.

$$T(x) = (x - ckx^2 I - 1)tI + \delta(1-x)(t_\theta - st)I \tag{9.27}$$

Taking partial derivatives with respect to x and equating to 0, we get equations 9.28 and 9.29.

$$x = \frac{1 - \delta\left(\frac{t_\theta}{t} - s\right)}{2ckI} \tag{9.28}$$

$$(t_\theta/t) = s + \delta(1 - 2ckxI) \tag{9.29}$$

Now consider a change in the tax rate from t_1 to t_2, that results in a change in reported taxable income from z_1 to z_2. Note that the reported taxable income $z = (1-x)I$. Let ε be the elasticity of reportable income, measured as the percentage change in reportable income for a 1% change in tax rate. This quantity is observable based on z_1 and z_2 Saez et al. (2012). By definition of ε, we get equation 9.30.

$$\varepsilon = \frac{(x_1 - x_2)t_1}{(1-x_1)(t_2 - t_1)} \tag{9.30}$$

Under the assumption that $\gamma = 1$, $s = 1$ and $\delta = 1$, that is, for self-interested decision makers, we first assume that t_θ is *not* adaptive, and remains constant as the tax rate changes from t_1 to t_2. We then get equation 9.31 as an estimator for t_θ. Here, c is known and k can be elicited by survey or proxy to a reasonable degree of approximation. This makes elasticity ε a sufficient statistic for t_θ.

$$t_\theta = \frac{2\varepsilon t_1 t_2(1 - ckI)}{t_1 + \varepsilon t_2} \tag{9.31}$$

However, if we assume (as the model proposes), that t_θ is not constant but adapts, then the elasticity, ε, is no longer a sufficient statistic. Assuming the simplest of adaptive processes, using equation 9.2 with n=1 and L_f=0, we get equations 9.32 and 9.33.

$$t_{\theta 2} = t_{\theta 1}(1 - L) + st_2 L \tag{9.32}$$

$$t_{\theta 1} = \frac{2\varepsilon t_2(t_2 - t_1)(ckI - 1) - t_1 t_2 L}{(1-L)t_1 - t_2 - \varepsilon(t_2 - t_1)\dfrac{t_2}{t_1}} \tag{9.33}$$

Critically, t_{θ_1} and hence the resulting change in tax evasion levels as t_1 changes to t_2 cannot be determined using the reportable tax elasticity ε, because of the unknown rate of adaptation L. I thus propose that the insufficiency of reported taxable income elasticity arises from the adaptive reference points. This can be solved by modeling the change in reference points based on informative priors, and getting a probabilistic estimate of t_θ as qualitatively outlined in earlier.

9.5.1 A cognitive explanation for the Slippery Slope Framework

The slippery slope framework is a popular framework for analysis of tax compliance (Kastlunger et al. (2013); Kirchler et al. (2008); Kogler et al. (2013); Prinz et al. (2014)). I propose that the generic consumption model applied to income tax evasion behavior maps nicely into this existing framework, and in fact, provides some additional insights. The slippery slope framework posits that people may comply with tax regulations either in a voluntary or an enforced manner. Voluntary compliance is higher if people have trust in the government, and enforced compliance is higher if people perceive that the government has a high level of legitimate or coercive power. Legitimate (fairly implemented) power drives both enforced and voluntary compliance. Lozza et al. (2013) found that enforced compliance through coercive rather than legitimate power can lead to increased tax evasion, especially among people with a left-wing political ideology. On the other hand, there is a stronger path to tax evasion via voluntary compliance in the case of people with a right wing ideology. Wahl et al. (2010) found that increasing trust in authorities increased voluntary but decreased enforced compliance, whereas increasing power of authorities decreased voluntary but increased enforced compliance. Muehlbacher et al. (2011) found that voluntary compliance increased with age and education levels, whereas enforced compliance was negatively correlated with education.

Within the currently presented framework, trust directly affects the adaptive reference point via the α and L_f parameters in equation 9.4, and thus impacts transaction utility. Power may indirectly affect the rate of hedonic adaptation, but primarily affects the acquisition utility, since it affects perceptions of the probability and consequences of getting audited, caught, or fined (parameters k and c in equation 9.23). Thus deontological decision ($\delta = 0$) makers can be affected by changes in perceptions of power, but are unlikely to be swayed by changes in trust or fairness of taxes, since they systemically ignore transaction costs. Utilitarian decision makers on the other hand are concerned with the fairness of the tax and redistribution policies. For such decision makers, having a high level of trust implies a higher reference point (t_θ). Such tax payers are likely to have

a high sensitivity to political climate, that is, a high value of L_f (see equation 9.4). As a result, their reference points may show larger or more extreme movements depending on the direction of trust α. In general, utilitarian tax evasion decreases with increasing values of the reference point. Thus, increasing trust would lead to increasing voluntary compliance in utilitarian tax payers, and vice versa Sidani et al. (2014). For self-interested decision makers, increased trust in the government, leading to a higher reference point, does not increase tax compliance. If anything, a higher reference point allows them to derive the same psychological transaction utility with an even lower rate of compliance. This suggests that tax evasion in self-interested decision makers is primarily through the enforced compliance route. Since self-interest decision makers may not be overtly concerned with the fairness of tax or redistribution, they are likely to have a low sensitivity to government trust L_f, in equation 9.4. If such decision makers have a low reference point, both confirmation bias and increasing trust would lead to faster adaptation towards the actual rate of tax, thus increasing the propensity to evade taxes.

Finally, a key question is what attributes drive people towards utilitarian, deontological, or self-interested tax decisions, since this seems to determine the critical slippery slope pathway. Frecknall-Hughes et al. (2016) measured propensity for deontological and utilitarian decision making for tax and non-tax specialists in tax and non-tax scenarios. They found that the relative propensity for deontological decision making in tax scenarios is higher for tax specialists than for non-tax specialists. This suggests that the influence of tax specialists on final tax related decision making behavior of individuals should be taken into account.

9.5.2 Accuracy of perceived income tax rates

The model predicts that if the perceived tax rate (measured by s) is higher than the actual (for instance when there is confusion between marginal and average tax rates), this does not significantly affect people with a higher reference point. However, it reduces the level of evasion for

self-interested tax payers with a lower reference point, while increasing evasion in utilitarian tax payers with a lower reference point. Since utilitarian tax payers are generally expected to have a higher reference point, this suggests that overestimating the true tax rate may in fact reduce evasion at an overall portfolio level. On the other hand, if the perceived tax rate is significantly lower ($s << 1$)than actuals, self-interest tax payers with low reference points may significantly increase evasion, as they underestimate the potential costs of getting caught, since they believe the amount in question is much lower than it actually is. Lower than actual perceived tax rates affect tax evasion much less in utilitarian tax payers. Thus transparency of taxes, while obviously a virtuous act in itself, may have negative economic consequences. Since this impact is estimated to be higher on self-interest rather than utilitarian taxpayers, efforts to improve tax transparency should be coupled with publicity about enforcement mechanisms. Information about redistributive justice and fairness of taxes will only be more effective to utilitarian taxpayers. Further, this messaging based on trust is very likely to be mediated by political affiliations. A Gallup poll (Gallup (2015)) showed that perceptions of fairness of income tax fell amongst low and middle income Republicans after Obama took office, but perceptions of fairness of tax amongst low income Democrats were unchanged, and amongst middle income Democrats improved. Actual tax changes during this particular period were not deemed significant. While perceptions of trust and fairness seem to play an important psychological role, increased conflation with political ideologies may blunt this as a policy tool. It is likely that perceptions of trust are mediated by a general psychological contract (Robinson and Morrison (1995)) rather than tax specific issues.

9.5.3 Tax evasion elasticity versus tax rate

The model also predicts that increasing tax rates also have opposing effects, with the propensity to evade increasing for utilitarian but dropping for self-interest decision makers, in response to a tax increase. A significant increase in tax rate will result in a slower reference point adaptation in utilitarian tax payers on account of confirmation bias. On the other hand, a significant drop in tax

rates will result in a slower adaptation of reference point by self-interested payers. This goes to explain the mixed results in literature (Freire-Serén and Panadés (2013)) on how tax changes affect tax evasion behavior. Bernasconi et al. (2014) found that in a lab experiment, individuals adapted faster to tax cuts than to tax increases. The model suggests that this behavior is typical of utilitarian rather than self-interested decision makers. Since most of the participants in this experiment were economics students, this seems a likely corroboration. The model also shows that sensitivity of tax evasion behavior is highest at low reference points, and that tax evasion elasticity flattens out as the tax rate increases beyond the reference point levels. Thus, tax cuts are expected to lead to higher tax evasion elasticity than similar tax increases. However, the key aspect to note is that response to tax cuts may result in an increase or decrease depending on the mode of processing. Assuming that the population proportions stay constant however, any net directional change of tax compliance levels on reduction of taxes is less likely to be fully recovered on a tax increase. There is one exception to the general trend of increasing / decreasing tax evasion versus the reference point. The model predicts a kink (in the opposite direction) where the reference point crosses the actual tax rate which results in a relatively large shift in the propensity to evade taxes. The hedonic adaptation mechanism saturates as the adaptive reference point approaches the actual tax rate from any direction. Hence, the model predicts that sharp changes in tax evasion levels are most likely caused when there are significant changes made to tax brackets. This is in line with previous empirical evidence (Saez et al. (2012); Chetty et al. (2009a)).

9.6 An initial proof of concept - Consumption decisions

In this section, we provide a proof of concept rather than a detailed empirical application, which will be considered in future work. We implement the log-linear demand specification using the structural model advocated by Chetty et al. (2009b), and also implement a modified version of this demand specification by incorporating the adaptive reference point based cognitive-framework

proposed in this chapter. We implement both models within a Bayesian inferential framework. We use data from Chetty et al. (2009b) that includes panel data on the per capita consumption of beer by state in the US for a period of 34 years, along with the corresponding price, tax, and other regulatory changes. The cognitive-framework essentially adds in temporal structure to reference points, so that the continuous adaptation in reference points for each state is inferred by the model.

Figure 9.1 shows the marginal improvement in error obtained by applying this framework. Whilst in the right direction, this level of improvement may not justify a significantly more complex approach. However, we propose that this approach can be significantly strengthened by making a full model identifiable (e.g. the current implementation assumes within state homogeneity of preferences, but there are various data sources, including surveys and publicly available preference data that can form strong priors and identify proportions of divergent behavior within each state). To make the cognitive approach more powerful, future work should incorporate such informed priors that identifies systematic heterogeneity across states, in terms of how people respond to trust in the government, their consumption bias, and similar aspects.

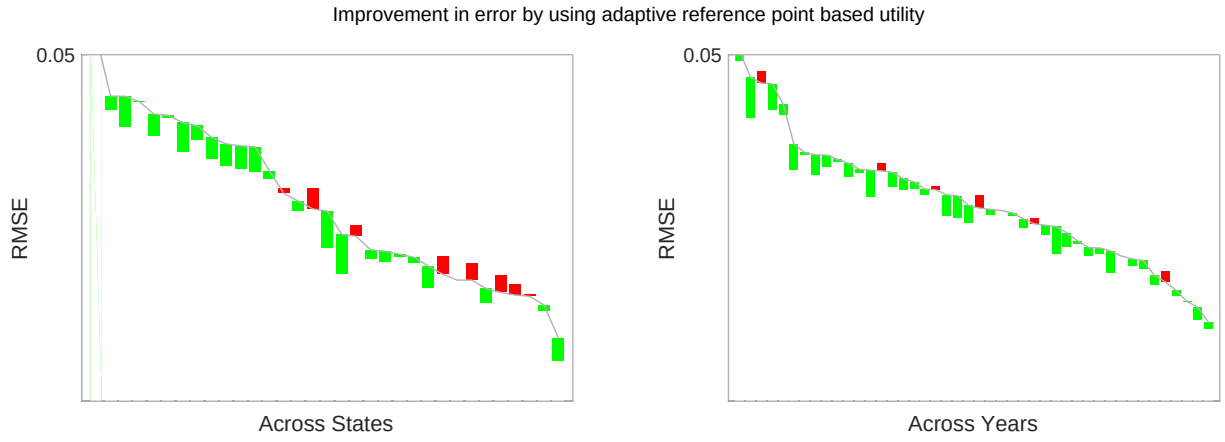

Figure 9.1: Descriptive RMSE for baseline regression is the gray line, in order of improving error across states and across years. The green and red bars show the improvement and deterioration respectively for each state (left subplot) and each year (right subplot), by using the adaptive reference point based utility model.

9.6.1 Inferences

Figure 9.2 shows the inferred structural assumptions about the movement of latent reference points for tax for two different sources, the reaction to tax changes and the perception of government. The background is color coded to reflect Democratic and Republican Presidents during the blue an pink years. The values are a mean across states, and there are variations across states. The bars represent the mean change in the tax rate during the year.

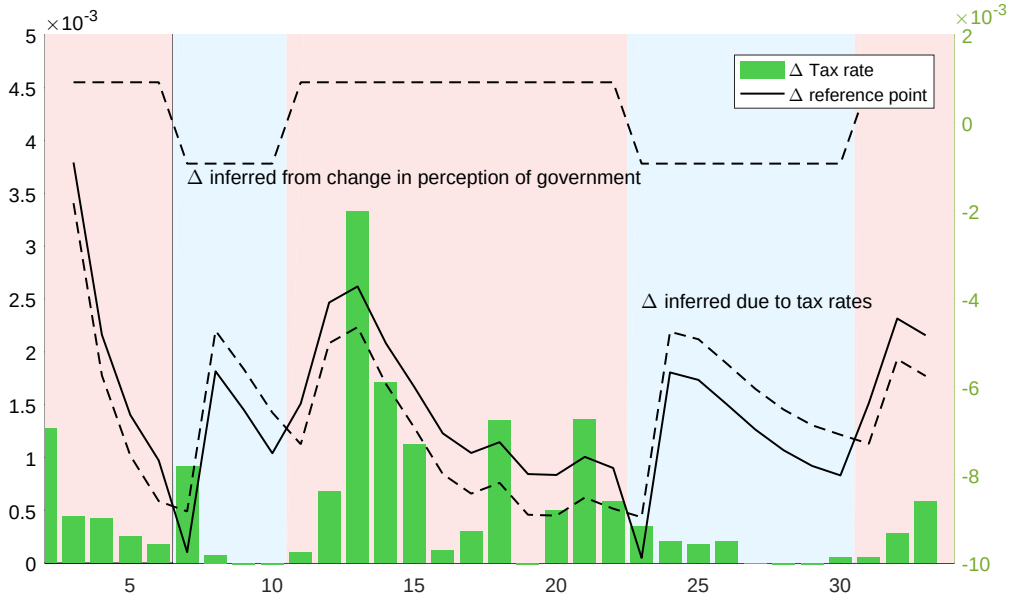

Figure 9.2: Inferred structural changes to reference point for consumption tax based on political regime and change in tax rate.

9.7 Conclusions

The methodological contributions of this chapter are to specify an extensive structural model of adaptive population behavior based on insights from cognitive psychology that highlights the limitations of a sufficient statistics based approach typically used to analyze such behavior in economics. Specifically, sufficient statistics seem restricted to special cases of the possible parameterization of the cognitive structural models. Under alternate assumptions of some of these parame-

ters, the sufficient statistics may be highly biased, inefficient, or both. In some cases, it is possible to infer multiple parameters from a structural model based on multiple empirical manipulations. This can however be an extremely tedious exercise. Another promising approach is the ability to combine sufficient statistics with informative priors about the parameters of the structural model within a Bayesian inference framework. This allows measuring the bias and noise in an estimator in a structured manner. More importantly it provides an avenue for counter-factual as well as predictive behavioral analysis.

Part IV

Inducing adaptive behavior - Behavioral nudging

Chapter 10

Behavioral Nudging in a resource allocation task

10.1 Introduction

Consider the problem of constrained resource allocation, a decision paradigm that is pervasive across many different domains. People distribute money across a range of retirement savings and investments options, government agencies allocate funding to different scientific research areas, the military allocates human resources across strategic geographies, water supplies are distributed across different end uses, and cyber security resources are allocated across different layers of avoidance and recovery mechanisms, to name only a few.

Often, investing a resource into a particular choice option may involve a riskless transaction cost or benefit, and a set of risky payoffs that may depend on stochastic outcomes. For example, cyber security allocations between avoidance and recovery have riskless implementation costs for both, and risky payoffs that depend on the frequency and intensity of cyber attacks. Many real word decisions however combine risky and riskless components into compound choice structures where the boundary between riskless and risky components is less clear. For instance, insurance plans may vary the combination of premium (riskless cost component) and deductible (risky cost component) in tandem. Most savings and investment options have a riskless transaction cost and a risky investment return payoff. In some cases, the risky payoffs are stated after netting off the riskless transaction costs. In other cases, the riskless components may be charged upfront and the risky components realized later. Sometimes, different risky choice options within a choice set often have proportional riskless components. For example, on an average, the riskless transaction costs on high-risk stock purchases are empirically higher than for medium-risk mutual funds, which are higher than for low-risk debt bonds. Regulating agencies may also aim to incentivize behavior by making costs or incentives proportional to risk levels, for example, in the case of capital charges that banks face depending on the risk level of their assets.

The motivating question for this paper is how risk-framing, in terms of the framing of risky and riskless components, can influence allocation decisions. To investigate this, we consider the sim-

plest representation of a prospect, in terms of a two-outcome gamble, where each outcome is associated with a different risky payoff, and each prospect may have a riskless acquisition cost. The decision maker is expected to distribute finite resources across a set of such two-outcome prospects which differ in terms of their risk-reward-cost profile, to construct a portfolio. There is a large amount of literature on how people might select between one of many such gambles, but not as much focus on how people distribute resources across a set of such gambles. The normative version of this allocation problem is extensively studied, in terms of algorithms for optimization of the risk-reward characteristics of the constructed portfolio.

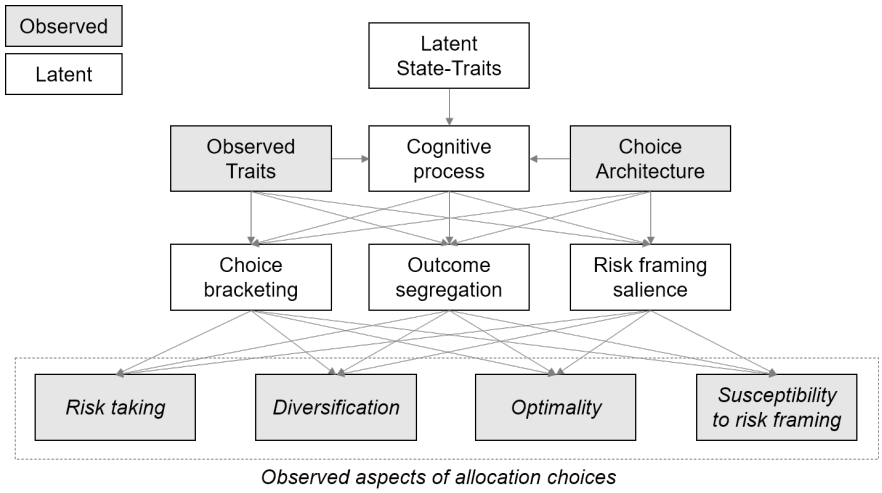

Figure 10.1: Key aspects of a resource allocation decision

We focus on the descriptive account of how people make these decisions. This is a far more nuanced decision than simply selecting one of many such gambles. The resource allocation decisions can be measured in terms of the level of risk undertaken, the level of diversification, how optimal or efficient the allocation is, and how easily such decisions can be manipulated by changing the risk-framing, between risky and riskless components.

To investigate possible structural sources of heterogeneity across all of these aspects, we adopt a cognitive modeling approach. We use this to infer how allocation decisions may be influenced (a) by observed aspects such as individual traits that can be measured, (b) by different aspects of choice architecture that can be experimentally manipulated, and (c) by latent aspects such as

whether people evaluate choice options together or individually (choice bracketing), whether people segregate the risky and riskless components (outcome segregation), the level of salience on different components of the risk-frame, and the underlying cognitive process people use to evaluate options (such as prospect theory). Figure 10.1 shows a schematic of the key observed and latent aspects that define resource allocation decisions.

In subsequent sections, we define the resource allocation task, show that based on theoretical considerations, the appropriate manipulation of risky and riskless components of choices can be the basis for a behavioral nudge, define key behavioral measures within this paradigm, provide a theoretical basis for both observed traits and latent cognitive factors that can be a source of heterogeneity in this task. We then report a novel experiment that manipulates risk framing to create behavioral nudges, and analyze behavior by incorporating observed and latent factors into a comprehensive computational model of cognition.

10.2 Defining the basic risk-resource allocation problem

Consider allocation of a finite amount of resources to a set of N two-outcome prospects. The proportion of total resources allocated to each prospect (i) is the allocation weight (w_i), with weights across all prospects summing to one. Since these are two outcome gambles, the outcomes are denoted as success and failure, with the payoff under a success outcome always being greater than that under a failure outcome. Each individual prospect is defined independently (no correlation between outcomes of different prospects), and each prospect may also have an associated riskless transaction cost, that reduces the net payoff. *The risky payoffs and riskless costs are defined in terms of percentage returns on the amount of resources invested in that prospect.* We define the following quantities:

p_{iS} = Probability of success for prospect i

$p_{iF} = 1 - p_{iS}$ = Probability of failure for prospect i

v_{iS} = Percentage payoff if the outcome is success

v_{iF} = Percentage payoff if the outcome is failure $(v_{iF} < v_{iS})$

C = Riskless acquisition cost (%) for the prospect $(C <= 0)$

Each prospect can thus be characterized in terms of its expected value v_i and standard deviation d_i.

$v_i = p_S v_S + p_F v_F + C$

$d_i = \sqrt{p_S(v_S + C)^2 + p_F(v_F + C)^2 - v_i^2}$

Since prospects are independent, a portfolio created by assigning weight w_i to the i^{th} prospect will have expected value V and standard deviation D. Further, the portfolio will have a *Herfindahl index* H, which measures the degree of diversification (Rhoades (1993)). When allocation is such that all weights are equal, H takes on the minimum value of 1/N (maximum diversification), and when all resources are allocated to a single prospect, H=1 (highest value, representing maximum concentration).

$V = \Sigma_{i=1}^{N} \; w_i \; v_i$

$D = \sqrt{\Sigma_{i=1}^{N} \; w_i \; d_i^2}$

$H = \Sigma_{i=1}^{N} \; w_i^2$

As an illustration, we provide an example in table 10.1, with $N = 4$ possible options. The table illustrates how the same normative choice set can be framed differently. Choice set 2 is equivalent to choice set 1, and is obtained by adding a riskless acquisition cost component (C) to each prospect, and increasing both v_S and v_F by the same amount. The corresponding prospects in the two choice sets have the same expected value and standard deviation. The bottom part of the table provides an illustration of 5 (out of an infinite) possible portfolios constructed using different allocation weights (w is a vector of allocation weights to prospects i_1 to i_4), and the resulting portfolio level characteristics. We will continue to use this example to illustrate concepts over the rest of the paper.

Table 10.1: Illustrative example: The two choice sets are equivalent, differing only in the framing of risky and riskless components. The last part of the table lists 5 example portfolios.

Choice Set Framing 1							
Prospect	p_S	v_S	p_F	v_F	C	v_i	d_i
i_1	0.5	5.2	0.5	4.8	0	5.0	0.2
i_2	0.6	7.5	0.4	3.8	0	6.0	1.8
i_3	0.7	9.5	0.3	1.1	0	7.0	3.8
i_4	0.8	10.9	0.2	-3.5	0	8.0	5.8
Choice Set Framing 2							
Prospect	p_S	v_S	p_F	v_F	C	v_i	d_i
i_1	0.5	6.2	0.5	5.8	-1.0	5.0	0.2
i_2	0.6	9.5	0.4	5.8	-2.0	6.0	1.8
i_3	0.7	12.5	0.3	4.1	-3.0	7.0	3.8
i_4	0.8	14.9	0.2	0.5	-4.0	8.0	5.8
Portfolio	Weights				V	D	H
P_1	w = [.5 .5 0 0]				5.5	0.9	0.50
P_2	w = [.25 .25 .25 .25]				6.5	1.8	0.25
P_3	w = [0 .5 .5 0]				6.5	2.1	0.50
P_4	w = [.5 0 0 .5]				6.5	2.9	0.50
P_5	w = [0 0 .5 .5]				7.5	3.5	0.50

10.3 Choice bracketing

The first latent factor in how people might make such decisions is choice bracketing. The concept of choice bracketing has been proposed in paradigms that involve selecting one out of many choice options (Hsee et al. (1999)), where people may evaluate these options separately (narrow choice bracketing) or jointly by comparison (wide choice bracketing). We propose an extension of this to a resource allocation paradigm. Here, narrow choice bracketing involves evaluation of each individual prospect separately, and deciding on the allocation based on a comparison of these individual evaluations. Wide choice bracketing on the other hand involves evaluating aggregate characteristics of a constructed portfolio of prospects, with the aim of optimizing one or more of

such portfolio level measures. Using the examples in table 10.1, narrow choice bracketing would involve independently evaluating i_1, i_2, i_3, and i_4, focusing on features in the top 2 parts of the table, and comparing these evaluations to arrive at an allocation. Wide choice bracketing would involve evaluating choices at a portfolio level, for example thinking about V and D, or other portfolio characteristics, and deciding between portfolios based on these comparisons.

Most *normative* optimization algorithms would involve wide choice bracketing, focusing on some portfolio level characteristics. However, this is cognitively very demanding. We might expect that people tend to evaluate these decisions using narrow choice bracketing, or some form of heuristics, unless the selected choice architecture makes wide choice bracketing relatively simple. Narrow and wide choice bracketing may apply regardless of the specific valuation approach, for example, people can apply cumulative prospect theory in both a narrow or wide choice bracket (as illustrated later in the paper). The choice bracketing decision can significantly influence allocation, since narrow and wide bracketing essentially reflect different cognitive processes used to evaluate the problem.

10.4 Outcome segregation and risk framing

The next latent factor we consider is outcome bracketing, which speaks directly to the primary issue in the paper - how people treat the risky versus riskless components. Early discussions on choice framing, including the original formulation of prospect theory (Kahneman and Tversky (1979)), discuss an editing phase, where people might combine, cancel, segregate, or otherwise edit aspects of different options before evaluating them. One of the proposed components of such an editing phase was the *segregation* of outcomes into riskless and risky components. Some studies (Marquis and Holmer (1996)) conduct model-based inference to suggest that people may indeed be performing some form of segregation of these components. On the other hand, Fischhoff (1983) suggests that segregation may not occur if *"people were overwhelmed by the surface structure*

of a problem". We propose that the "surface structure" of the problem may be instrumental in determining both, whether the risky and riskless components of outcomes are segregated, and whether they are evaluated with different levels of salience. To define this formally, we draw upon principles from the theory of mental accounting (Thaler (1985)), which proposes that the overall utility (U) of a prospect incorporates both transaction (TU) and acquisition utility (AU). Here, the two utilities are evaluated in a segregated manner, and the acquisition utility may be weighted differently (difference in salience), by a factor k.

Let $f()$ define a utility function or valuation process (for example, prospect theory). Then normatively, evaluating an individual prospect would yield a normative utility (NU).

$$NU_i = f(p_{iS}v_{iS} + p_{iF}v_{iF} + C_i) = f(v_i) \tag{10.1}$$

On the other hand, segregated outcome evaluation based on the theory of mental accounting would yield an experienced utility (EU). In terms of the prospects considered in this paper, the riskless cost C defines the acquisition utility.

$$EU_i = TU + k\,AU = f(p_{iS}v_{iS} + p_{iF}v_{iF}) + k\,f(C_i) = f(v_i - C_i) + k\,f(C_i) \tag{10.2}$$

Table 10.2: Normative and experienced utilities for the two choice sets presented in Table 10.1

Choice Set Framing 1		
Option	NU	EU
i_1	$f(5)$	$f(5)$
i_2	$f(6)$	$f(6)$
i_3	$f(7)$	$f(7)$
i_4	$f(8)$	$f(8)$
Choice Set Framing 2		
Option	NU	EU
i_1	$f(5)$	$f(6)+k\,f(-1)$
i_2	$f(6)$	$f(8)+k\,f(-2)$
i_3	$f(7)$	$f(10)+k\,f(-3)$
i_4	$f(8)$	$f(12)+k\,f(-4)$

Outcome segregation between riskless and risky components can influence allocation decisions if $EU \neq NU$. This can occur (a) because of the difference in salience between risky and riskless components (i.e. $k \neq 1$), and (b) because of non-linear utility functions $f()$, even if $k = 1$. The normative (combined) and experienced (segregated) utilities for the two example choices sets are illustrated in table 10.2. Since the acquisition cost only enters in choice set 2, if outcomes for this set are evaluated in a segregated manner, it could lead to allocation decisions that are different from those made for choice set 1, although the two choice sets are equivalent.

10.5 Behavioral nudges based on risk framing

Thaler and Sunstein (2008) define a nudge, as *"any aspect of the choice architecture that alters people's behavior in a predictable way without forbidding any options or significantly changing their economic incentives"*. Consider the two choices sets in table 10.1. For people who do decide to segregate outcomes for choice set 2, a change in the choice architecture (presentation of riskless and risky components) may result in altering allocation decisions, without any economic change, since the choice sets are economically equivalent, but $EU \neq NU$. When comparing multiple prospects, the quantity of interest is $\Delta U_i = EU_i - NU_i$. The higher this difference, the more desirable prospect i will seem to be, under outcome segregation.

$$\Delta U_i = f(v_i - C_i) - f(v_i) + k\, f(C_i) \tag{10.3}$$

We assume that $f(-x) = -\lambda f(x)$, as in prospect theory, where λ is a loss aversion parameter. Then we can write,

$$\Delta U_i = f(v_i - C_i) - f(v_i) - k\lambda f(-C_i) \tag{10.4}$$

Equation 10.4 shows that ΔU_i is a linear function of individual preferences $k\lambda$, with the slope of the

function being $-f(-C_i)$. Note that C_i is non-positive, hence as $k\lambda$ increases, ΔU_i will fall. Now consider any two prospects i_1 and i_2, with $C_1 \neq C2$. The values of ΔU_1 and ΔU_2 when plotted as a function of $k\lambda$, will always intersect. This point of intersection, $k\lambda = \tau_{12}$, is when $\Delta U_1 = \Delta U_2$. This gives,

$$\tau_{12} = \frac{\Big(f(v_2) - f(v_1)\Big) - \Big(f(v_2 - C_2) - f(v_1 - C_1)\Big)}{-\Big(f(-C_2) - f(-C_1)\Big)} \tag{10.5}$$

If individual preferences $k\lambda < \tau_{12}$, the change in preference ΔU_i under outcome segregation will be higher for the prospect with higher value of C_i, and vice versa for individual preferences $k\lambda > \tau_{12}$. We extend this to the resource allocation problem with multiple prospects, and define two specific choice architectures.

In both cases, we define prospects which are not dominated, such that prospects with higher standard deviation d_i do not have lower expected returns v_i. In the first, we place a constraint that the acquisition cost C_i and for prospects increases monotonically with an increase in risk-return (v_i, d_i). We denote this as $'C_i \uparrow (v_i, d_i)'$. In this case, people with lower values of individual preference $k\lambda$ will have a relative increase in preference for riskier prospects under outcome segregation. People with higher values of $k\lambda$ will have a relative increase in preference for safer prospects under outcome segregation. In the second choice architecutre, acquisition cost C_i decreases monotonically with (v_i, d_i), denoted as $'C_i \downarrow (v_i, d_i)'$. Here, the effect is opposite, as summarized in table 10.3. Essentially, depending on the level of salience to segregated acquisition costs, determined by $k\lambda$, the monotonic change in acquisition costs with a change in risk-reward profile of prospects may thus act as a behavioral nudge.

We return to the example in table 10.1. Choice set framing 2 is an example of $'C_i \uparrow (v_i, d_i)'$, as acquisition costs increase with increasing risk-reward profiles. To illustrate the points made, we plot ΔU_i for each of the four prospects as a function of different values of $k\lambda$ in figure 10.2, under

Table 10.3: Nudge effects as a function of individual preferences and choice architecture

		Choice Architecture	
		$C_i \uparrow (v_i, d_i)$	$C \downarrow (v_i, d_i)$
Individual	Low $k\lambda$	Riskier	Safer
	High $k\lambda$	Safer	Riskier

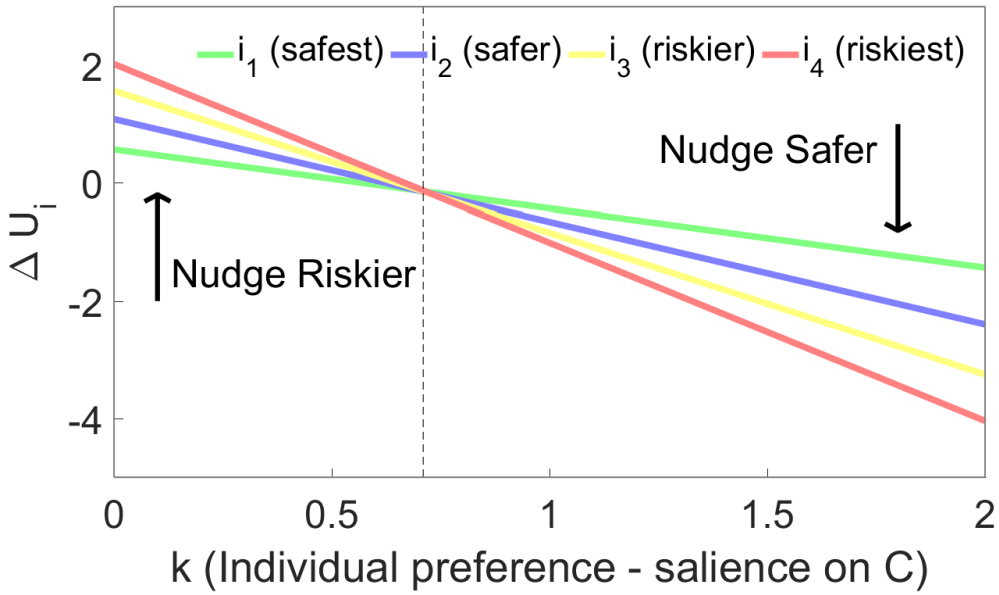

Figure 10.2: Difference between segregated experience utility (EU) and normative utility (NU) for the 4 options in choice set 2: $C_i \uparrow (v_i, d_i)$. This illustration assumes a concave utility function $f(x) = x^{0.8}$. The plots show that depending on the value of $k\lambda$, a matter of individual preference, the shift in utilities may progressively favor either the riskier or safer prospects.

the assumption of a concave utility function $f(x) = x^{0.8}$. It can be seen that while the individual preference levels for $k\lambda$ remain close to 0, that is, very low levels of salience of the acquisition cost, the riskier prospects begin to look more and more desirable under the framing for choice set 2. If individual preferences are high levels of $k\lambda$, that is, very high salience of the segregated acquisition costs, the safer prospects begin to look progressively more desirable under choice set 2.

We might expect that people demonstrate a wide range of behavior in terms of salience to acquisition costs, and the resulting value of $k\lambda$. We again look to principles from the theory of mental accounting. This theory posits the principle of hedonic segregation. This states that people will

be prone to segregate outcomes if such segregation leads to higher experienced utility. Thus, it might be expected that people are more likely to segregate outcomes and evaluate prospects using EU rather than NU, if $\Delta U_i > 0$. If $f()$ is concave, then, $0 \geq f(v_i - C_i) - f(v_i) \geq f(-C_i)$. Hence $\Delta U_i > 0$ only when $k\lambda << 1$. Thus, if people generally behave as predicted by the principle of hedonic segregation, they are more likely to consider utilities on a segregated outcome basis when their individual preferences are such that $k\lambda << 1$. Thus the principle of hedonic segregation strongly favors behavior where people are expected to be nudged towards riskier prospects under $'C_i \uparrow (v_i, d_i)'$ and towards safer prospects under $'C_i \downarrow (v_i, d_i)'$. Mathematically, the principle of hedonic segregation suggested that outcome segregation will be pursued for evaluation only if,

$$k\lambda < \frac{f(v_i - C_i) - f(v_i)}{f(-Ci)} < 1 \tag{10.6}$$

These observations are the premise of creating a behavioral nudge based on risk framing. In general, under the principle of hedonic segregation, a behavioral nudge pushing people towards safer options may be created if the level of translation of risky and riskless components in opposite directions (in order to maintain equivalence) is progressively higher for safer options, and vice versa for a nudge pushing people towards riskier options. Of course, it is possible that people do not underweight the transaction costs (k is high), and hence under hedonic segregation, may be evaluating the options normatively. On the other hand, it is also possible that some people do not behavior in line with the principle of hedonic segregation, and may display opposite behavior. The experimental design and computational modeling approach later in the paper attempt to identify and isolate these effects.

10.6 Experimental Methods

In this experiment (Mistry and Trueblood (2017)), we test people's preferences for allocating a fixed set of resources between multiple risky prospects, with a focus on testing the effects of behavioral nudges created using risk-framing effects.

10.6.1 Task

The cover story for the task was that participants had to play the role of the head of a company that had the opportunity to invest a fixed amount of money (hypothetical $100,000) into one or more of 4 possible projects. Participants were advised that all projects had the same expected time to completion and their objective was to maximize the return on the invested amount. They were required to invest all the money, but could distribute this in any proportion between the 4 projects, including allocating no resources to one or more projects. Each project had two possible outcomes - success or failure. They were provided with the probability of success (p_S) and failure (p_F = 1-p_S) for each project, as well as the percentage returns on their investment depending on whether a project succeeded or failed. A successful project always had a positive return (v_S), whereas a failed project resulted in either a lower positive or a negative return (v_F). The 4 projects always varied in terms of the variability (standard deviation) of return outcomes.

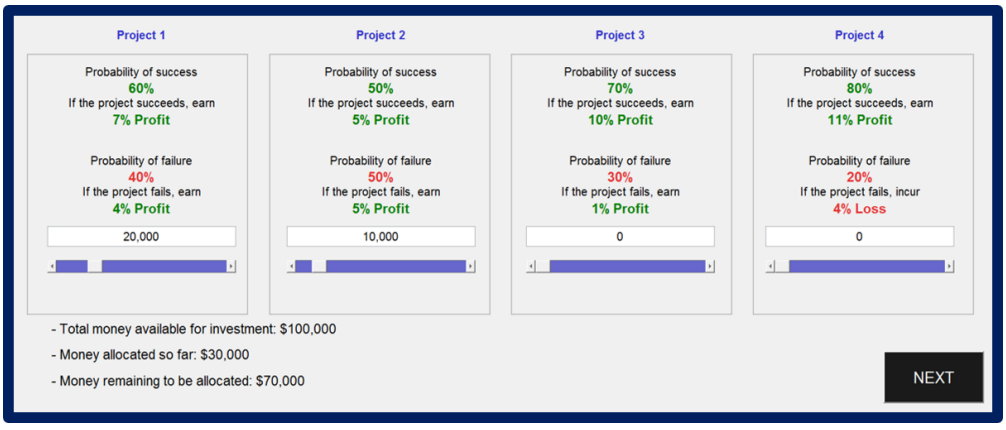

Figure 10.3: Illustration of experimental interface for portfolio choice allocation decisions

Figure 10.4: Illustration of experimental interface for providing feedback to participants after each trial

Participants were given an example and a practice trial to familiarize themselves with the interface (see Figure 10.3). After each trial, participants were provided feedback on the outcome. The outcome was based on the described probability of success and dynamically (randomly) picked by the computer program. The process of realization of the outcome for each project was graphically displayed to the participants. For each project, they were shown a box containing 100 balls, of which $100p_S$ were green and $100p_F$ were red. The computer program randomly traversed the box space and eventually picked one of the balls. A green ball implied success, and a red ball implied failure of the project (see Figure 10.4). This was done independently for each project. The returns on the investment for each project and for the weighted portfolio were updated based on these outcomes before moving on to the next trial.

10.6.2 Participants

50 undergraduate students from Vanderbilt University participated in the experiment, either for credit (n = 25) or for financial incentives (n = 25). Of the 50 participants, 38 were female, 11 were male, and 1 reported other. Their ages ranged from 18 to 49 years, with a mean of 20.6 (SD = 5.4).

10.6.3 Between-subjects conditions

Participants were split into 2 groups of 25 students each. The between subjects design entailed different rewards, with the rest of the design factors being identical between the two groups. Group 1 participated for course credit, and group 2 for financial compensation. This between-subjects condition, which we denote as *CI*, tests whether financial incentives affect portfolio choice allocation in laboratory tasks. There is mixed prior evidence for this in tasks involving risky choices Beattie and Loomes (1997). Participants in group 2 received a fixed payout of $5 plus an incentive ranging from $0 to $10 that was linked to their performance on the task. At the end of the experiment, one of the trials was randomly selected. The incentive component was calculated as $5 plus or minus $0.10 times the %returns achieved on that trial, but limited to the range $0-$10. So for example, achieving a loss of 20% resulted in an incentive of $5 - 0.1(20) = $3 (and a total payout of $8), and achieving a gain of 20% resulted in an incentive of $5 + 0.1(20) = $7 (and a total payout of $12). This allowed for the incentive to be framed as reductions for losses and increments for gains. The total payout for each participant in group 2, including the fixed and incentive components, ranged between $5-$15. The hypothetical in-task payoffs were incentive compatible, and the method of calculation was explained to the participants before the task.

10.6.4 Within-subjects factorial design

Each participant completed 36 portfolio choice decisions. To maximize risk preference information obtained from these decisions, there were 12 unique decisions based on a 2 X 2 X 3 within-subject factorial design. This entailed a 2 (second order stochastic dominance - present vs absent) X 2 (domain - gains vs mixed) X 3 (skew - none, positive, negative) manipulation. Each of these 12 decisions was repeated in 3 blocks, with the choices randomized within each block. Although the underlying decision remained equivalent across blocks, the three blocks varied in terms of a risk-framing effect, by translating outcomes and adding a corresponding acquisition cost:

Second order stochastic dominance (SOSD; 2 levels):

The *SOSD* manipulation included 2 levels. In the first, all prospects in a trial had equal expected value, but the 4 prospects had progressively higher standard deviation. As a result, each prospect had SOSD over the subsequent riskier prospects. In the second level, the prospects were not mean preserving, and riskier prospects (higher standard deviation) also had higher expected values. Thus there was no SOSD. Any behavioral account that is based on a weakly increasing concave utility function, or mean-variance optimization, predicts a strong preference for prospects that have SOSD. Decisions involving SOSD choices allow a straightforward measure of inefficiencies in optimization. Decisions that do not involve SOSD choices allow a nuanced measure of risk tolerance.

Domain (DM; 2 levels):

The *DM* manipulation included 2 levels. In the first, all-gain domain, all outcomes including project failures resulted in positive returns. In the second, mixed domain, the average returns across prospects on *failure* were negative. Even for the latter, the safest option had both outcomes positive, so that investing in this option was always preferable to not making any investment. Domain manipulation allows testing for the effects of asymmetric gain-loss utilities within the

portfolio choice framework.

Skew (SK, 3 levels):

The *SK* manipulations included 3 levels. In the first, all prospects had zero skew, that is, success and failure were equally likely. In the second, all except the safest prospect had negative skew, that is, failure outcomes were more likely. In the third, all except the safest prospect had positive skew, that is, success outcomes were more likely. Symmonds et al. (2011) showed that risk and skewness are differently encoded in the brain. People have been shown to be relatively averse to negatively skewed gambles (Mellers et al. (1992); Deck and Schlesinger (2010)). Manipulation of skew allows us to test whether these effects extend to the portfolio choice paradigm, as well as test whether people employ safety-first approaches to minimize the probability of their portfolio value dropping below a certain risk governed threshold.

Risk-framing (PC, 3 levels):

The *PC* manipulation included 3 levels, implemented within a blocked design. In the first level, there are no extraneous purchase costs. In the second and third levels, the outcomes from the first level were translated and re-framed into higher gross outcomes accompanied by an appropriate purchase cost. This re-framing led to prospects that were expected-value-equivalent to the prospects presented in the first level. In the second level, the amount of re-framing was increasingly higher with increasing variability (risk) of prospects outcomes, representing choice architecture '$C_i \uparrow (v_i, d_i)$'. In the third, the re-framing decreased with increasing variability (risk), representing choice architecture '$C_i \downarrow (v_i, d_i)$'. The decisions between blocks were equivalent in the sense that the same percentage allocation across prospects resulted in the construction of identical portfolios, net of purchase costs. The second and third blocks thus attempted to create behavioral nudges pushing people towards riskier and safer prospects respectively, under the principle of hedonic segregation. The trials were presented in a blocked design with three blocks corresponding to the 3 purchase costs conditions, with the order of the 12 problems in each being randomized.

10.6.5 Behavioral measures

The resource allocation paradigm offers many degrees of freedom for people to express individual differences. This observed heterogeneity includes the level of risk people are ready to take on (e.g. in table 10.1, the risk of the portfolios P_1 to P_5 continuously keeps increasing), the level of diversification (e.g. P_2 is more diversified compared to the remaining illustrative portfolios), and how optimal people's choices are (e.g. in table 10.1, $P2$ and $P3$ have the same expected value but $P2$ has lower standard deviation, and may be consider more optimal). Most importantly in this paper, we also want to evaluate individual differences in how susceptible people might be to changing their preferences under the risk framing effects, that is, how robust are these potential behavioral nudges proposed in the previous section? These aspects are represented in the bottom row of the schematic in figure 10.1. In this section we specify how each of these observed aspects is measured:

Segregated measures (S, R):

The simplest way of measuring a resource allocation decision is to look at the *allocation weights* (w_i) for each (i^{th}) prospect, where $\Sigma_{i=1}^{N}\ w_i\ = 1$, and N is the total number of choices available. 691999Lopes and Oden () proposed that there are individual differences in whether people approach risky decision making from a perspective of security (protecting low outcomes) or potential (maximizing high outcomes). A simplistic measure of people's security and aspiration levels are measured by the weight allocated to the two extreme prospects - *safest* (S) and *riskiest* (R) respectively.

Aggregated measures (V, D, T):

Often, the emergent characteristics of the aggregated portfolio are of greater interest than the individual choices. Most normative theories of portfolio choice are based on optimizing some function of the portfolio characteristics. Since a portfolio can be represented as a probability distribution

over outcomes, the most common characteristics are derived from the moments of the resulting portfolio. We calculate the *expected value (V)*, and the *standard deviation (D)* of the aggregate portfolio. We also measure the actual *realized returns* (T) which are the portfolio weighted realized outcomes on each trial based on resolution of the selected prospects. This is a function of the resolution of stochastic outcomes that the participants actually observe.

Portfolio diversification (H)

In addition, the *Herfindahl index* $(H) = \Sigma_{i=1}^{N} w_i^2$, where N is the total number of prospects in the choice set, measures the degree of diversification Rhoades (1993). When all weights are equal, H takes the minimum value of 1/N (maximum diversification) and when all resources are allocated to a single prospect, H=1. For N=4, values close to 0.25 indicate naive diversification, values close to 0.5 indicate some form of conditional diversification (equal allocation to 2 of 4 prospects), and values close to 1 indicate concentration in a single prospect.

Optimality of portfolio ($\varepsilon_w, Q_w, \varepsilon_w - \varepsilon_N, Q_w/Q_N$)

We consider two broad normative solutions, which we propose should be considered as the lower and upper benchmarks, when evaluating how well people allocate resources. As a lower benchmark, we consider naive diversification, or the 1/N heuristic (Benartzi and Thaler (2001); DeMiguel et al. (2007)). This heuristic basically proposes equal allocation to all prospects under consideration. A conditional diversification heuristic is similar, but restricts equal allocation to a subset of the prospects under consideration, after eliminating some prospects based on a selected criteria. Naive or conditional diversification heuristics are cognitively simple to implement, have been shown to perform comparably to more sophisticated optimization rules in some real-world environments (DeMiguel et al. (2007)). People have also been shown to have some bias towards diversification (Bardolet et al. (2011)). It stands to reason that any cognitive effort in optimizing a resource allocation decision should try to improve performance over a naive diversification heuristic.

As an upper benchmark on performance, we use one of the most popular normative theories of resource allocation, modern portfolio theory (MPT), characterized by mean-variance optimization (Markowitz (1952)). This states that people should select weights that optimize the balance between the expected value V and standard deviation D of the resulting portfolio. The optimization is a function of a *risk tolerance factor* Q, which is a matter of individual preference. There have been several subsequent optimization algorithms that build upon this framework. However, we base our measure of optimality on the MPT for two reasons. First, optimization under this framework can be viewed as based upon simple to understand concepts of expected value and standard deviation, and thus define cognitively plausible behavior, compared to more complex optimization algorithms. Secondly, there is a large body of research that proposes behavioral decision analysis, manipulation of choice architecture, providing incentives and subsidies, or training and debiasing human decisions under the assumption that people may be carrying out, have the potential to carry out, or should be carrying out, some kind of mean-variance optimization. This approach has been taken in several fields including improving the efficiency in allocating resources for water conservation, combating terrorism, managing fossil fuel use, and capital investments, among others (Roques et al. (2008); Gaydon et al. (2012); Phillips (2009); Byers et al. (2015)).

Given a set of risky prospects, MPT proposes an *efficient frontier* of possible weight allocations that result in mean-variance optimization. Given an implicit objective to maximize expected value V and minimize standard deviation D, the frontier represents portfolio choices such that no other combination of weights can result in an increase in V without an increase in D, or a decrease in D without a decrease in V. The efficient frontier can comprise a wide variety of possible weight allocation, and can be viewed as parametric on a risk tolerance factor Q, where $Q >= 0$. A value of $Q = 0$ indicates preference only for the safest prospect and extremely high values of Q tending to ∞ indicate preference only for the riskiest prospect. From a perspective of human behavior, risk tolerance Q can be viewed as an individual preference, however, it is *not* a reliable psychometric trait, because it varies significantly depending on the specific set of choices under consideration. As per MPT, the set of weights (x) on the efficient frontier for a particular value of Q can be found

by minimizing the expression:

$$x_{optimal}|Q = argmin_x\left(x^T \Sigma\, x \;-\; Q\, E^T x|Q\right) \tag{10.7}$$

Here $E = [v_1, v_2, ..., v_N]$ is a vector of expected returns on the individual prospects and Σ is the covariance matrix for the returns on the prospects. For the purpose of this paper, since prospects are uncorrelated, Σ has non-diagonal elements zero, and the diagonal elements are $[d_1^2, d_2^2, ..., d_N^2]$.

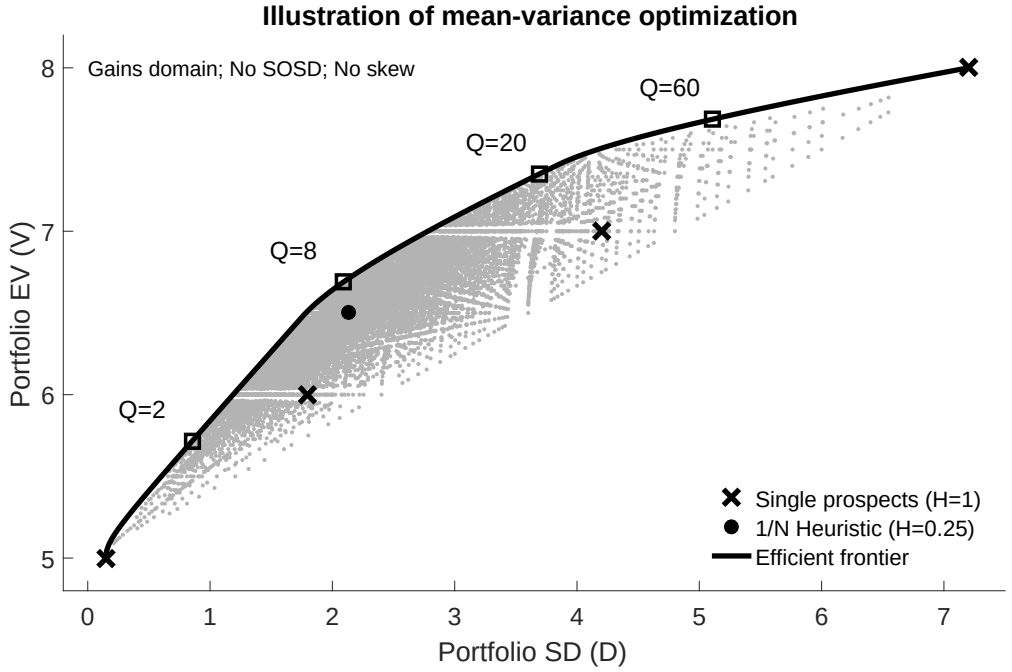

Figure 10.5: Illustration of the efficient frontier for one of the portfolio choice problems. The gray dots represent different portfolio allocation weights and the resulting portfolio characteristics. The black line marks the efficient frontier under a mean-variance optimization framework. The squares illustrate the optimal portfolio for some sample levels of risk tolerance (Q).

Figure 10.5 illustrates the range of portfolio characteristics (V and D) for different portfolio allocation weights (gray dots) for a representative resource allocation problem involving four risk prospects. The crosses mark selection of single-prospect portfolios, and the circle shows the 1/N or naive diversification based portfolio. The thick line shows the efficient frontier for this choice problem. As risk tolerance Q increases, the optimal portfolio shifts to increase both V and D along the efficient frontier.

For any observed portfolio allocation (w_i) we can calculate the minimum euclidean distance of the observed allocation weights to the set of portfolio weights that comprise the efficient frontier. This gives the minimum distance to optimality, ε_w. The risk tolerance value corresponding to the closest point on the efficient frontier is denoted Q_w, and can be inferred to be a model-free estimate of risk tolerance level for that choice. Similar estimates can also be obtained for the naive diversification heuristic, which we notate as ε_N and Q_N.

We propose that cognitive effort towards improving optimality is reasonable if $\varepsilon_w < \varepsilon_N$, since the constructed portfolio is closer to optimality than a cognitively simple naive diversification heuristic. Secondly, a comparison of of Q_w and Q_N may provide a more reliable psychometric indication of risk tolerance across different forms of choice sets, where the absolute value of Q_w may vary significantly.

Susceptibility to risk-framing nudges (N)

Finally, we measure differences between different choice set framing conditions. The framing conditions are setup so that correctly accounting for the costs and translation of outcomes should result in no difference between behavior across the three conditions. However, segregating and placing differential levels of salience on the costs framed separately would result in a preference for prospects with a higher degree of framing (low salience) or lower degree of framing (high salience). In one condition, riskier prospects are subject to higher framing (we denote this condition as F_1), and in the other, safer prospects are subject to higher framing, denoted as F_2. Discounting the costs would result in higher selection of riskier prospects in the first and safer in the second framing condition. We calculate *susceptibility to nudges* as, $N = mean[(S_{F_2} - S_{F_1}), (R_{F_1} - R_{F_2})]$. A value of N close to 0 indicates that people are *not* susceptible to cost framing nudges. A high positive value indicates that people under-weight separately framed costs, and thus are nudged towards options with higher framing (larger translation of outcomes). A high negative value indicates that people over-weight separately framed costs, and thus are nudged towards options with lower framing (smaller translation of outcomes).

10.6.6 Measured individual traits

After the resource allocation task was complete, participants were required to complete surveys or tasks in to observe the following set of traits:

Elicited risk propensity (HL)

Multiple price lists present a series of choices between two risky gambles to elicit ranges of risk aversion. We use the price list measures from Holt et al. (2002) to measure elicited risk propensity, with higher scores indicating higher risk aversion. This is a direct task based measure and has been used to reveal risk preferences which might be applicable to a variety of domains. We test whether this measure based on simpler risk-based choices is robust enough to predict the level of risk undertaken in more complex resource allocation decisions, and whether this trait impacts how people react to risk-framing nudges.

Self-reported financial risk propensity (DF)

The financial risk-taking subscale within the DOSPERT scale (Blais and Weber (2006)), is a self-reported measure of risk propensity, based on survey type questions. This self-reported measure has been used in many financial applications. The nature of the resource allocation task in this paper is financial, and we test whether this self-reported measure influences allocation decisions.

Risk congruency (RCI)

Self-reported and elicited risk propensity measures may not always coincide, and may point in opposite directions. We propose a new metric in terms of risk congruency, the degree of difference between the elicited and self-reported measures above. We measure the differences based on normalized scores, and highlight that risk congruency is bi-directional, so that a score of zero would indicate people that are perfectly risk congruent, high scores in one direction indicate people who believe that they are far more risk seeking than their elicited measures suggest, and high scores

in the opposite direction indicate people who believe that they are far more risk averse than their elicited measures suggest. We propose that this metacognitive judgment of their own risk preferences may impact how people make resource allocation decisions. Within risk congruency, we also classify risk-congruent (RC=1) people as risk-seeking versus risk-averse. On the other hand, risk incongruent (RC=0) people are classified as those that think they are more risk-seeking than they actually are and those that think they are more risk averse than they actually are.

Internal versus external control (SL)

Finally, the locus of control scale Rotter (1966), measures internal versus external control and has been shown to influence risky decision making in single choice tasks, with increasing external locus of control associated with increase risk seeking behavior in simpler experimental tasks.

10.7 Experimental Results

Table 10.4 shows the summary statistics and correlations for the key behavioral measures defined. Table 10.5 shows the breakup of mean values by experimental condition and choice architecture. Reporting is based on Bayesian statistics, with a log Bayes Factor (LBF) used to determine significance. Table 10.6 shows the LBF for whether the key measures depend in individual risk traits measured, with LBF *ge* 1 indicating that the risk traits influence the behavioral measure.

The left panel in figure 10.6 shows the distribution of H across participants and trials. The color shading also shows the distribution of the number of unique prospects selected on any trial. A large mass of the distribution lies between the range of 0.25 and 0.5, with further peaks at 0.5 and 1.0 indicating choices where people selected 2 of the 4 prospects equally, or invested all their resources in a single prospect. The right panel in figure 10.6 shows the joint distribution of allocation to S and R. It shows a large variety in patterns of allocation with peaks at 25%, 50% or 100% allocation to either R or S, but also a wide spread of intermediate behavioral patterns.

Table 10.4: Summary statistics: Correlations with LBF ≥ 1 highlighted in bold

			Correlation with							
Measures	Mean	SD	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
(1) Herfindahl Index (H)	0.48	0.23	**0.22**	**0.20**	0.02	**0.12**	-0.04	-0.05	**0.10**	**-0.11**
(2) % Safest (S)	0.34	0.28		**-0.56**	**-0.30**	**-0.44**	**-0.10**	0.05	**0.10**	-0.06
(3) %Riskiest (R)	0.27	0.26			**0.25**	**0.61**	0.06	-0.08	**-0.12**	0.01
(4) Portfolio EV (V)	3.2	2.8				**-0.38**	**0.40**	-0.01	-0.02	0.01
(5) Portfolio SD (D)	4.1	4.0					**-0.20**	-0.04	**-0.08**	0.00
(6) Realized returns (T)	3.5	6.8						-0.00	-0.02	0.01
(7) Locus of Control (SL)	13.1	3.8							**-0.16**	0.07
(8) Risk aversion (HL)	4.9	1.9								**-0.23**
(9) Risk seeking (DF)	3.0	1.1								

Table 10.5: Mean values by condition: Values highlighted in bold indicate that the difference within levels for that design factor have LBF ≥ 1

	SOSD		Domain		Skew			Cost framing		
	(Yes)	(No)	(Gain)	(Mixed)	(0)	(Neg)	(Pos)	(None)	(Riskier)	(Safer)
H	0.48	0.47	0.48	0.47	0.46	0.48	0.49	**0.44**	**0.49**	**0.49**
S	**0.37**	**0.31**	**0.29**	**0.39**	**0.34**	**0.37**	**0.31**	0.32	0.33	0.36
R	**0.25**	**0.29**	**0.30**	**0.24**	0.26	0.28	0.27	0.27	0.28	0.26
V	**2.50**	**3.98**	**5.84**	**0.64**	3.23	3.25	3.24	3.27	3.26	3.19
D	**3.83**	**4.42**	**2.06**	**6.19**	**4.38**	**3.60**	**4.39**	4.23	4.32	3.83
N	0.01	0.04	0.00	0.05	0.02	0.03	0.02	-	-	-
ε_w	**0.76**	**0.34**	0.55	0.55	**0.57**	**0.51**	**0.57**	0.53	0.57	0.55
Q_w	-	63.7	**19.9**	**107.5**	64.7	52.4	73.9	**67.0**	**76.1**	**48.0**

Table 10.6: LBF for influence of risk traits (HL, DF, SL, RCI) on key behavioral measures. In addition, the column $RC = 1$ measures the influence of congruent risk seekers versus congruent risk aversiveness, and $RC = 0$ measures the influence of levels of incongruency - whether people overestimate their risk seeking or risk aversive behavior

	HL	DF	SL	RCI	RC=1	RI=1
H	0.9	**6.2**	-1.5	**1.1**	**10.6**	-2.0
S	**4.9**	-0.1	**3.4**	**1.3**	**3.6**	-1.5
R	**11.0**	-2.2	**6.8**	**1.0**	**1.5**	**2.2**
V	**4.4**	-1.1	**5.0**	0.7	**1.3**	-1.2
D	**5.5**	-2.0	**3.1**	**2.6**	**1.2**	**1.3**
N	**4.3**	-1.6	-1.3	**1.0**	-1.3	**7.7**
ε_w	-0.7	-2.9	-2.3	-2.4	-1.6	-1.2
Q_w	-1.5	-2.0	**2.7**	**2.7**	-1.4	-1.0

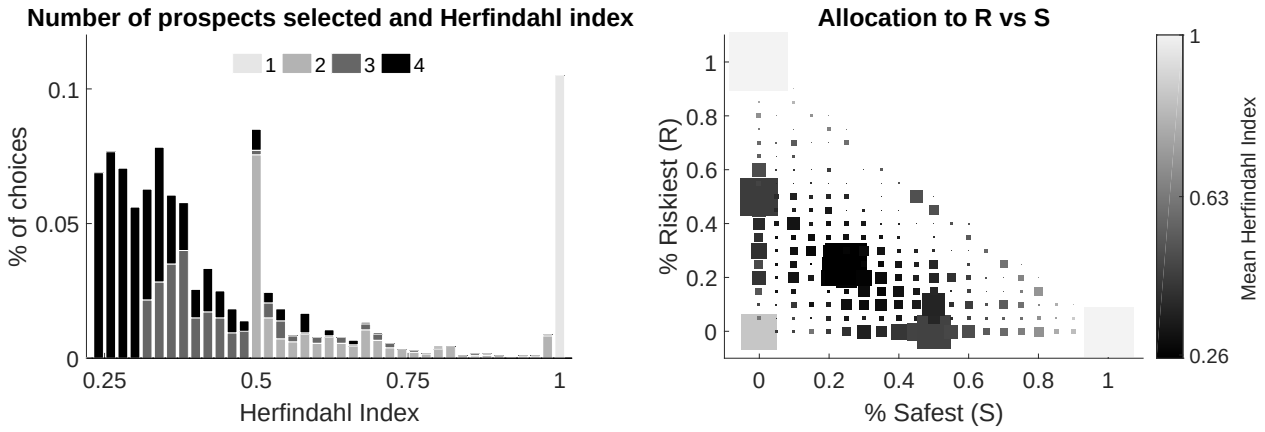

Figure 10.6: *Left Panel:* Distribution of Herfindahl index across participants and trials. The color shading shows the number of unique prospects (1 to 4) selected on each trial. *Right Panel:* Proportion of resources allocated to the riskiest prospect (R) vs to the safest prospect (S) on each trial. Size of the squares reflects the proportion of trials and the color of the squares reflects the mean Herfindahl index for that combination of R vs S.

10.7.1 Impact of choice manipulations:

A Bayesian repeated measures ANOVA JASP-Team (2016) is used to test the effect of choice manipulations on these measures. Figure 10.7 shows the mean allocations to each of the 4 prospects split by type of choice manipulation. There is no evidence that the incentive condition had any effect on S, R, H, V, or D. There is no evidence that the level of diversification as measured by H is affected by the domain, skew, or SOSD manipulations. There is evidence for a main effect of domain (LBF 28.1), SOSD (LBF 12.8), and skew (LBF 3.2) on S. Allocation to S is higher in the mixed domain (mean 39%) than in the gains domain (mean 29%), higher in the SOSD (mean 37%) compared to non-SOSD (mean 31%) condition, and higher in negative skew (mean 37%) than positive skew (mean 31%) conditions. There is evidence for a main effect of domain (LBF 9.6) and SOSD (LBF 4.3) on R. Allocation to R is higher in the gains domain (mean 30%) than in the mixed domain (mean 24%), and higher in the non-SOSD (mean 29%) compared to the SOSD (mean 25%) condition. V and D are expected vary with domain and SOSD by design. There is no evidence for a main effect of skew on V, but there is evidence (LBF 6.9) for a main effect of skew on D. Participants exhibit the lowest D (mean 3.6) in the negative skew condition and highest D

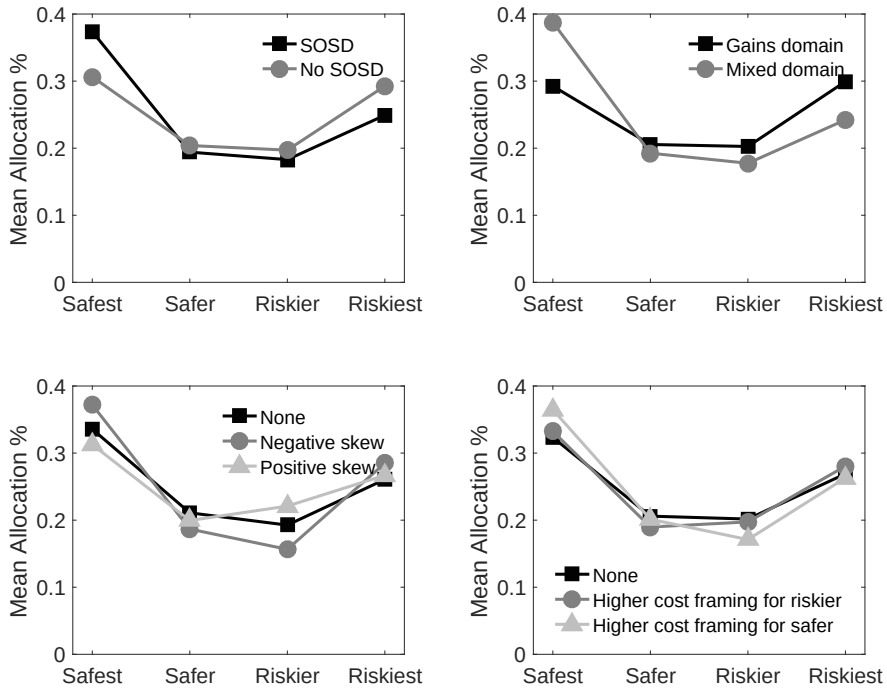

Figure 10.7: Mean allocation to the 4 prospects ranging from safest to riskiest, split by choice manipulation.

(mean 4.39) in the positive skew condition, indicating a marked preference for lower variability in the negative skew condition.

10.7.2 Impact of purchase cost framing effects:

First, we test for differences between the framing conditions and no framing condition is more complicated, since there is the possibility of order effects confounding results. To test this, we first run a Bayesian ANCOVA analysis testing for an order effect of item presentation order within each of the three blocks. We find evidence of no order effects on H, S or R within any of the three blocks. A repeated measures Bayesian ANOVA reveals strong evidence (LBF 10.3) for an effect of cost framing on H. Framing (of either type) reduces diversification, with mean H increasing from 0.44 in the no framing condition to 0.49 in both the cost framing conditions (see figure 10.8). A possible hypothesis is that whilst people may not be susceptible to cost framing effects, the

introduction of an additional cognitive element induces people to reduce their diversification. This is supported by the observation that in the no framing condition, people selected all 4 prospects on 55% and either 1 or 2 prospects on 22% of the trials. Compared to this, in the framing conditions (combined), people selected all 4 prospects on 45% and either 1 or 2 prospects on 33% of the trials. Overall, the measure of susceptibility to nudges (N) shows values close to 0, indicating that people on an average may be susceptibility to cost framing effects. However, a deeper analysis shows that there is in fact a significant amount of heterogeneity in peoples' responses and susceptibility to the cost framing nudges, as seen in figure 10.9.

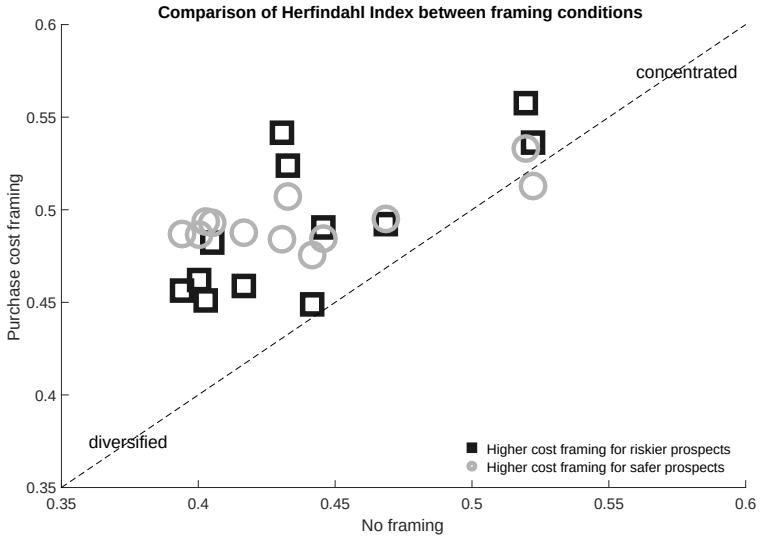

Figure 10.8: Reduced diversification (higher H) in cost framing conditions compared to the no-framing condition

Table 10.7: Mean values of susceptibility to nudges by risk congruency traits

Traits	% of sample	mean N
Risk congruent	60%	4.3%
Risk incongruent	40%	-0.2%
- Self-report risk averse	26%	**5.5**%
- Self-report risk seeking	14%	**-11.0**%

To identify whether individual traits can explain this heterogeneity, we look at whether people were risk congruent (i.e. whether peoples' self-reported measures of risk propensity (DF) and elicited measures of risk propensity (HL) were congruent). We find that risk incongruence is a key driver

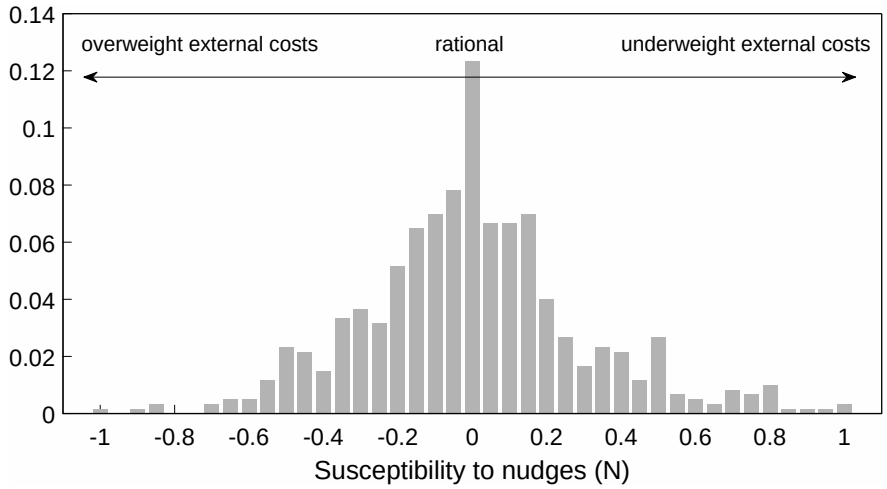

Figure 10.9: Distribution of N, showing significant heterogeneity in whether people are rational, overweight or underweight external costs, and hence their likely susceptibility to cost framing nudges

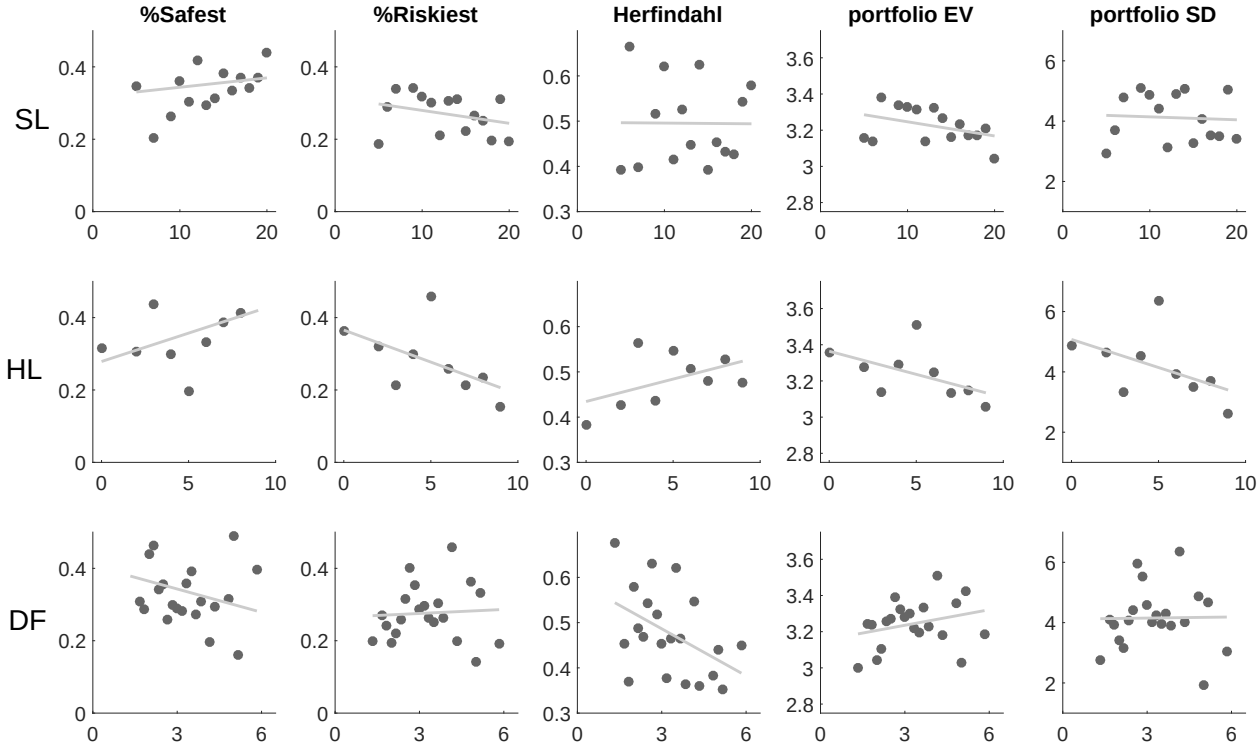

Figure 10.10: Mean values of S, R, H, V, and D for unique values of locus of control (SL), risk aversion (HL), and financial risk seeking (DF). The lines indicate best fit regression.

of how susceptible people were to nudges, as summarized in table 10.7. The key implication here is that risk incongruency (conflict between self-reported and elicited risk preferences) gives rise to significant differences, so that risk incongruent participants that believed they were risk averse

discounted external framed costs whereas risk incongruent participants that believed they were risk seeking over-weighted external transaction costs, leading to opposite reactions between groups for the two types of nudges. This difference was not observed in risk-congruent participants.

10.8 Discussion

10.8.1 Do previous outcomes influence behavior?

We test whether the status of returns (whether gains or losses) on a previous trial has any impact on behavior in the subsequent trial, by conducting Bayesian independent sample t-tests. There is evidence that previous trial outcome has no impact on *H* (LBF -2.6), *R* (LBF -2.6), *S* (LBF -1.8), or *D* (LBF -1.6). There is mild evidence that previous trial outcome has an impact on *V* (LBF 1.2), with slightly higher values (mean 3.6) after a loss trial than after a gain trial (mean 3.1).

10.8.2 Do traits influence behavior?

Figure 10.10 shows the mean values of dependent measures for unique values of the three trait-based measures (SL, HL, and DF). To test if the risk based traits influence behavior in the portfolio choice task, we use a Bayesian ANCOVA analysis treating the between and within subject choice manipulation factors as random effects and testing for the effects of covariates locus of control (SL), risk aversion (HL), and financial risk seeking (DF). Separately, we also measure the influence of risk congruency (RCI). Table 10.6 shows the LBF for the influence of risk traits on behavior. We find evidence that elicited risk averseness (HL) and locus of control (SL) have an impact on S, R, V, and D so that higher elicited risk averseness in simple gamble choices does carry over to an extent and result in relatively risk averse behavior in allocation tasks. Increasing external locus of control also leads to similar effects as increasing risk averseness. Directionally, this is in

contrast to findings based on risky gambles (Rotter (1966)) which showed that increasing external locus of control was associated with waging more money on riskier bets. On the other hand, the level of diversification seems to be influenced only by the self-elicited risk propensity (DF), with higher self-reported risk seeking behavior linked to higher diversification. This behavior is again contrary to the popular notion that diversification leads to reduced risk. Interestingly, risk congruent individuals show lower levels of risk taking in allocation tasks compared to risk incongruent individuals. The higher risk taking in the latter group is driven by individuals that believe they are risk averse based on self-reporting, but whose behavior even in simple gamble selection elicitation tasks demonstrates risk seeking behavior.

Next, we look at the susceptibility to nudges, N. Elicited risk propensity (HL) significantly influences susceptibility to nudges, with increasing risk averseness linked to lower susceptibility to nudges. The susceptibility to nudges also differs significantly within the risk incongruent group - individuals who believe they are risk seeking but show risk averse behavior behave very differently from everyone else. They seem to significantly overweight acquisition costs and behave in a way that seems opposite to hedonic segregation. They hence show behavior in the direction opposite to the intended nudges.

10.8.3 How close to mean-variance optimization do people get?

Conducting a Bayesian ANOVA analysis and comparing against a null model that included participants as random effects, we find evidence for the effect of SOSD (LBF 36) and skew (LBF 3.8) on ε_w. We also find evidence for the effect of domain (LBF 36.7) and cost framing (LBF 1.1) on Q_w. On an average, the actual allocations that people make are less optimal from a mean-variance optimization standpoint than what a simple 1/N heuristic would result in. Surprisingly, the distance to optimality is significantly lower for the more complex non-SOSD choice sets. The mean inferred risk tolerance Q_w is 20 in the gains domain and 107 in the mixed domain, indicating that risk tol-

erance is highly contextual, rather than a stable trait. The mean inferred value of risk tolerance Q_w is 67 in the no-framing condition, 76 in the higher-riskier-framing condition, and 48 in the higher-safer-framing condition is 48, reflecting sensitivity of risk tolerance to framing effects. There is evidence that measures of risk traits (SL, HL, DF) do not have any effect on the closest distance to optimality, but SL does impact risk tolerance Q_w. Note that if people indeed discount purchase costs, but maintain an otherwise constant level of risk tolerance Q, this would however *appear* to produce the same behavior as a shift in the risk tolerance. Specifically, discounting costs would appear to increase risk tolerance in the riskier framing condition and decrease tolerance in the safer framing condition. This is exactly what is observed. As we shall see in subsequent sections, by assuming that risk tolerance for a particular choice set is fixed, regardless of cost framing, we can use the observed choices to impute the amount of cost discounting carried out by individuals.

10.9 Behavioral models that explain resource allocation

We outline a set of behavioral models that combine assumptions about choice bracketing, outcome segregation, and people's utility and valuation functions. Since this is a latent cognitive process, we implement these as computational models of cognition, with a mixture model setup, to identify how likely each of these models is to have generated the observed behavior for each individual. In this section, we first outline the different behavioral models considered, by categorizing them into narrow and wide choice bracketing.

10.9.1 Narrow choice bracketing

Here, each prospect is evaluated independently based on one or more objective functions. These valuations for each prospect are then combined into a set of weights, as below:

$$w_i \;=\; \Sigma_j \upsilon_j \omega_{ij} \tag{10.8}$$

$$\omega_{ij} \;=\; \frac{1}{1 \;+\; \Sigma_k \, e^{-\eta(V_{ij} - V_{kj})}} \tag{10.9}$$

$$i,k \in [1:N] \;\; k \neq i; \quad j \in [1:J]; \quad \eta \geq 0$$

Here we consider N prospects, with each prospect indicated by the index i, and J concurrently held objective functions, each represented by the index j. The overall allocation weight for prospect i is w_i and is given by equation 10.8. Here, υ_j is a free parameter that represents the weight placed on the objective function j. Equation 10.9 shows how $w_{i,j}$, the allocation weight given to prospect i under the function j is computed. This depends on $V_{i,j}$, the valuation for prospect i under objective function j, and η, a free parameter that determines the bias for diversification versus concentration. For $\eta = 0$ this reduces to the 1/N heuristic, where each prospect is allocated an equal weight. For very high values of η this reduces to a greedy choice, where all the resources are placed in a single prospect. In practice, the number of concurrently utilized objective functions J would typically be 1, and rarely be more than 2. There are of course many possible objective functions that people might be using; for the purpose of this paper we consider the following:

1. Expected value (EV); J=1:

This is the baseline model, and known to be poor at explaining choices under risk, but is included as a lower reference point. This assumes that people diversify between prospect based on their diversification bias and the relative EV of prospects, in accordance with equation 10.9.

2. Cumulative prospect theory value (CPT); J=1 Tversky and Kahneman (1992)

This assumes that people diversify between prospects based on their diversification bias and the relative CPT derived value (V_{CPT}) of individual prospects, in accordance with equation 10.9. The CPT valuation is based on equations 10.10 - 10.14. In equation 10.10, π are the decision weights that are obtained by calculating the difference in cumulative capacities (w^+ or w^-) between a sequentially ordered set of prospects (see (Tversky and Kahneman (1992)) for details).

$$V_{CPT} = \Sigma\pi^+ v^+ + \Sigma\pi^- v^- \tag{10.10}$$

$$v^+(x) = x^\alpha; \; x \geq 0 \tag{10.11}$$

$$v^-(x) = -\lambda(-x)^\alpha; \; x < 0 \tag{10.12}$$

$$w^+(p) = \frac{p^\gamma}{(p^\gamma + (1-p)^\gamma)^{1/\gamma}} \tag{10.13}$$

$$w^-(p) = \frac{p^\delta}{(p^\delta + (1-p)^\delta)^{1/\delta}} \tag{10.14}$$

3. Security-Potential / Aspiration theory (SPA); J=2 (Lopes and Oden (1999))

This theory proposes individual differences in the balance of whether people are security (safety) minded or aspiration minded. Further, people may have high or low safety and aspirational thresholds. People with higher thresholds and high skew towards safety or aspiration would accordingly concentrate their allocations. People with lower thresholds indulge in greater diversification. The theory proposes two criteria; *SP* based on a balance of security-potential, and *A*, based on aspiration. We use each of these two criteria ($J = 2$), to calculate possible weight allocation in accordance with equation 10.9. these weights are then combined based on people's inferred balance between *SP* and *A* using equation 10.8. The criteria for aspiration, *A* is given by equation 10.15, where α_{SPA} is the aspiration level and *A* represents the probability that a prospect will yield an outcome that is

at least as high as this aspiration level. The criteria for SP is obtained from equations 10.16-10.17, where D_i is the decumulative probability of outcome v_i, that is, the probability of obtaining an outcome at least as high as v_i. Here $h(D)$ is a decumulative weighting function, and w_G determines whether a decision maker is security minded (biased towards w_G = 1) or potential minded (biased towards w_G=0). For further details, see Lopes and Oden (1999).

$$A = p(\ v \geq \alpha_{SPA}\ ) \tag{10.15}$$

$$SP = \Sigma_{i=1}^{N}\ h(D_i)\ (v_i - v_{i-1}) \tag{10.16}$$

$$h(D) = w_G\ D^{1+q_s} + (1 - w_G)\ (1 - (1 - D)^{1+q_p}) \tag{10.17}$$

10.9.2 Wide choice bracketing

In aggregated portfolio choices, prospects are not evaluated independently. Rather, the combined portfolio based on a set of allocation weights is evaluated based on some objective functions, with the aim that the constructed portfolio optimizes a set of objective functions. For the purpose of this paper, we consider that people might be using one of 2 possible classes of aggregated models:

1. *Mean-variance optimization*: Here, the optimal set of weights are not unique, but a subset of all possible combinations and form an *efficient frontier* (Markowitz (1952)). Any combination of weights along this efficient frontier is considered optimal, and the actual selection along this frontier depends on a free risk tolerance parameter Q. The efficient frontier is the set of combination of weights such that the portfolio EV cannot be increased without increasing the variance, and the portfolio variance cannot be reduced without reducing the EV. This assumes that higher EV and lower variance are desirable features of the portfolio. The efficient frontier consists of solutions to equation 10.18, where E^T is the vector of expected returns on prospect and Σ is the covariance

matrix for prospect returns.

$$w_i \mid arg\ min_{w_i}\ [w_i^T \Sigma w_i\ -\ Q\ E^T w_i \mid Q] \qquad Q \geq 0 \tag{10.18}$$

2. *Maximizing aggregated CPT*: Here, the selected weights w_i maximize the CPT objective function V. This function is applied to the portfolio as a whole rather than individual prospects. Given a particular set of parameters that define an objective function, there will typically be a unique combination of allocation weights that maximizes the objective function. This can be represented as:

$$w_i \mid arg\ max_{w_i}\ [V_{w_i}] \tag{10.19}$$

The key difference between segregated and aggregated CPT is that in the former, each prospect is independently evaluated under CPT to yield n different CPT values which are then combined in some form. In the latter, the weighted portfolio is treated as a single multiple outcome prospect, where the value of each *aggregated outcome* is the underlying outcome for the prospect multiplied by the allocation weight to the prospect. Thus, a portfolio with four prospects, each with 2 possible outcomes, will be treated as a single 8-outcome prospect, with the total probabilities of the 8 outcomes adding up to 1. This 8-outcome prospect is then subject to a CPT valuation. Since CPT evaluation is rank dependent, the allocation weights can change the sorted order of outcomes and it is possible that a small change in weights results in a relatively large change in aggregated CPT valuation. The link between Δw and ΔV_{CPT} is thus not constant, linear or monotonic under an aggregated mode of evaluation.

10.9.3 Cognitive parameterization of the behavioral models

The 3 narrow bracketing and 2 wide bracketing behavioral models are implemented as a mixture model within a hierarchical Bayesian inference framework. The models were implemented using MCMC sampling in JAGS (Plummer et al. (2003)). We make specific assumptions based on the

plausibility of cognitive processes, and link behavioral patterns to personality traits. To do this, we use a cognitive latent variable modeling approach (Vandekerckhove (2014)), where unobserved latent variables such as diversification bias and preference for choice bracketing are inferred based on a combination of observed behavior and measured traits. Specifically, we impose the following constraints:

1. *Diversification bias (*η*)* as applied in equation 10.9 under a given model (m) for a particular individual (i), is treated as a random-effect individual trait that remains constant across choice sets. Individual differences in diversification bias are modeled as being dependent on (z-transformed) measured traits SL, HL, and DF, as shown in equation 10.20.

$$\eta_{im} \sim \beta_{1m}^{(\eta)} z_{SL} + \beta_{2m}^{(\eta)} z_{HL} + \beta_{3m}^{(\eta)} z_{DF} + \beta_{0m}^{(\eta)} + \varepsilon^{(\eta)} \tag{10.20}$$

2. We propose that *purchase cost framing* is accounted for by separately evaluating the objective function based on the translated outcomes, and then adding a transaction utility (a negative amount that depends on the transaction cost), similar to the mental accounting proposed by Thaler (1985). In the mental accounting framework, transaction utility is is given by $\kappa V(-c, -c^*)$, where κ is the relative weight on transaction utility, and V is the objective (utility) function, c is the actual cost, and c^* is the reference point for the cost. Within our framework, c^* is 0. We modify the transaction utility function to $\big(V(-c)\big)^{\kappa}$, where $(0 \leq \kappa \leq 1)$. Here $\kappa = 0$ implies that any cost, if present, is treated as a single equivalent unit of cost. Non-linear utilities (CPT, SPA) may thus result in naturally emergent cost framing effects even if costs are not discounted. In addition, discounting of purchase costs ($\kappa \neq 1$) is treated as an individual trait (random effect) that remains constant across choice sets for an individual. Note that the same approach to cost framing may also result in different behavior under segregated or aggregated objective function evaluation.

3. *Risk tolerance* (Q) within the mean-variance optimization framework is not allowed to vary freely across all choices sets. For experiment 2, the viable range of risk tolerance parameters are the same for all choice sets, and risk tolerance is modeled as a random-effect individual trait that

is constant for an individual across all choice sets. For experiment 1, the choice sets differ significantly depending on the domain (gains vs mixed), and risk tolerance is assumed to be contextual depending on the domain. Further, based on previous studies that emphasize changing risk preference for negatively and positively skewed prospects, we allow risk tolerance to be contextually changed depending on the skew of the choice set. Importantly, risk tolerance is assumed to remain constant under different purchase cost framing choices of the same underlying problem. This allows us to infer the effects of purchase cost discounting under the optimization framework, and attribute apparent changes in risk preference to cost discounting or accounting rather than changes in underlying risk preferences.

4. *Preference for aggregated versus segregated* evaluation is treated as an mixed-effect, influence both by personality traits (SL, HL, DF) and nature of choice sets. We implement the behavioral models as a mixture model within a hierarchical Bayesian modeling framework. This allows us to measure the likelihood of a particular behavioral model being used by each individual and choice type. Model inference is performed separately for experiments 1 and 2 (since the choice sets are very different), but combined for the different between-subject conditions in experiment 1. We model the specific preference (ρ) for each of the 5 considered models (m) for each indivudal (i), as ρ_{im}, as shown in equations 10.21, 10.22 (for experiment 1) and 10.23 (for experiment 2). This is based on individual level random effects ψ_{im} and choice set based effects.

$$\psi_{im} \sim \beta_{1m}^{(\psi)} z_{SL} + \beta_{2m}^{(\psi)} z_{HL} + \beta_{3m}^{(\psi)} z_{DF} + \beta_{0m}^{(\psi)} + \varepsilon^{(\psi)} \tag{10.21}$$

$$\rho_{im} \propto \Phi\left(\psi_{im} + \beta_{domain}^{(\rho)} + \beta_{sosd}^{(\rho)} + \beta_{skew}^{(\rho)} + \beta_{framing}^{(\rho)}\right) \tag{10.22}$$

$$\rho_{im} \propto \Phi\left(\psi_{im} + \beta_{ev}^{(\rho)} + \beta_{sosd}^{(\rho)} + \beta_{skew}^{(\rho)} + \beta_{correlation}^{(\rho)}\right) \tag{10.23}$$

5. For parameterized CPT (both segregated and aggregated), the parameters $[\alpha, \gamma, \delta, \lambda]$ are treated as individual random effects within a hierarchical population distribution. Similarly for segregated

SPA, the parameters $[\alpha_{SPA}, w_G, q_S, q_P, \upsilon_{SP}]$ are treated as individual random effects within a hierarchical population distribution. Since J=2 for SPA, υ_{SP} is the weight parameter used in equation 10.8 to balance the allocations derived separately from the SP and A criteria.

10.10 Application of behavioral models to experimental data

Using Bayesian inference, we infer the parameters of the mixture model used to separately fit data from both experiments. The 5 models considered are segregated EV (EV-S), segregated CPT (CPT-S), segregated SPA (SPA-S), aggregated mean variance optimization (MVO-A), and aggregated CPT (CPT-A). The mixture model allows us to make inferences about the nature of choice bracketing choices and preference for segregated versus aggregated models. We also fit each of the five behavioral models individually to each experimental dataset to infer the full posterior distribution of parameters. Aside from the inference on choice bracketing, the parameter inferences in the below section are based on separately inferred models.

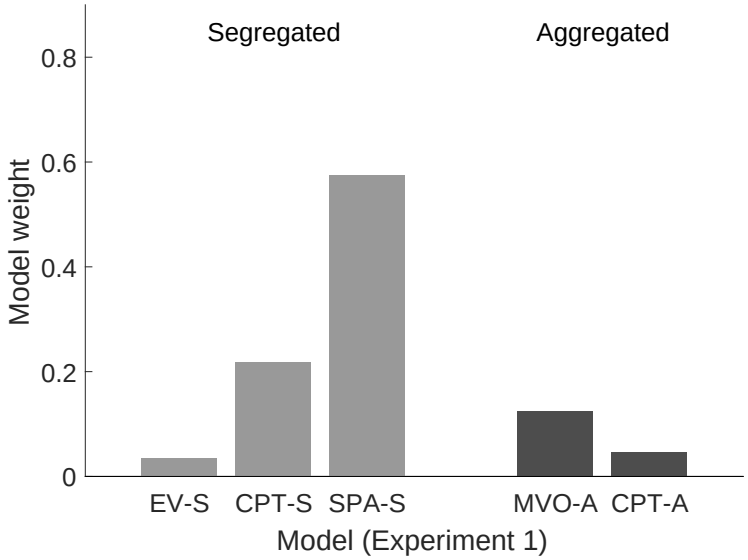

Figure 10.11: Model preference in the experiment over all participants and choice decisions. Segregated models, (specifically, SPA-S) are the dominant preference

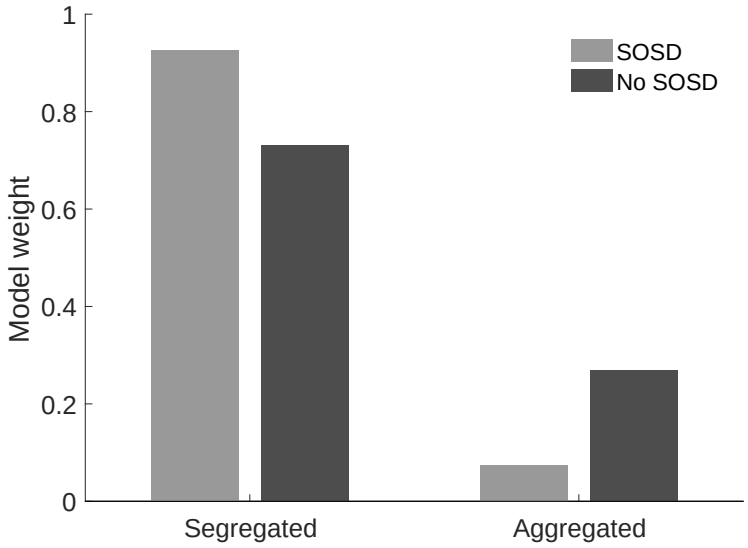

Figure 10.12: Model preference by choice set. Non SOSD choice sets, and choice sets with skewed prospects increase relative preference for aggregated choice bracketing.

10.10.1 Choice bracketing and model preferences

We obtain the model weights, or the probabilities for each of the 5 models being used for every individual choice decision. Figure 10.11 show the weights (summing to 1) over all individuals and trials. The mean weights were 57%, 22%, and 13% for SPA-S, CPT-S, and MVO-A in experiment 1, and 74%, 19%, and 4% for SPA-S, CPT-S, and CPT-A in experiment 2. Segregated models are the dominant preference, with SPA-S being the single dominant choice model in both experiments. We combine the preference for segregated and aggregated choice models, and find that in experiment 1, aspects of choice architecture had a significant influence on choice bracketing. Figure 10.12 shows that in experiment 1, the preference for aggregated models, although still lower, relatively increases in non-SOSD choice sets, and in choice sets that include skewed prospects. The choice architecture factors in experiment 2 do not seem to have a significant influence on choice bracketing.

Figure 10.13 shows individual differences in preference for aggregate models in both the experiments, as measured by the difference between participants that had low or high SL, HL, and DF scores. High and low scores are based on z-transforms above and below zero. Locus of con-

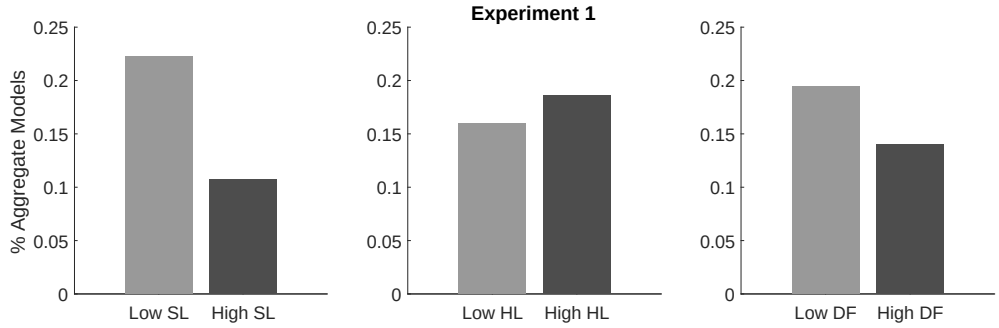

Figure 10.13: Individual differences in choice bracketing by choice set. Non SOSD choice sets increase relative preference for aggregated choice bracketing.

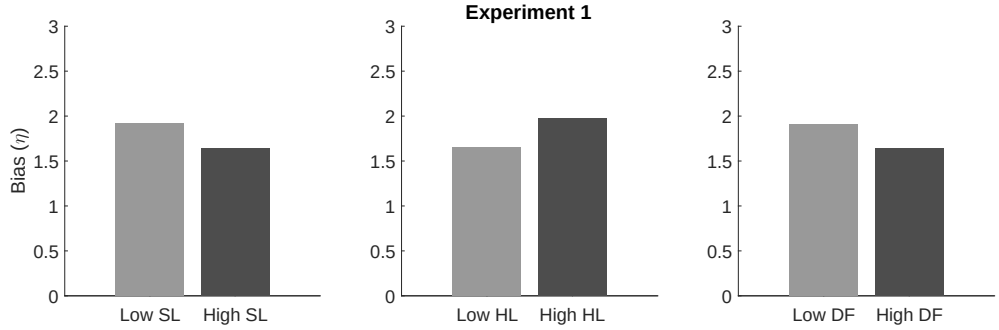

Figure 10.14: Individual differences in diversification / concentration bias (η)based on the SPA-S model. Higher values of η indicate higher concentration.

trol (SL) seems to have a strong effect, so that people with higher internal locus of control (low SL scores) have a higher propensity to for aggregated choice bracketing decisions in both experiments. The differences based on risk averseness (high HL and low DF) versus risk seeking (low HL and high DF) are not stable or strong enough to provide any conclusion.

10.10.2 Diversification Bias

Diversification bias is assumed to hold constant for an individual across all choice sets. Our cognitive latent variable modeling approach allows us to directly model this as dependent on the measured . We evaluate this bias, η, where higher values of η indicate a higher bias towards concentration for the SPA-S model, given that this is the most dominant model across experiments. The mean value of η_{SPA-S} (see table 10.10)is similar in experiment 1 (mean 1.45) and 2 (mean 1.37).

As seen in figure 10.14, the results are directionally consistent between the experiments for the influence of locus of control and HL, but the effects are stronger in experiment 2. Higher internal locus of control (lower SL) seems to drive higher concentration. Higher risk aversion (HL) also seems to drive higher concentration.

10.10.3 MVO-A Risk tolerance

We allowed the inferred risk tolerance to depend upon the skew and domain of the choice sets. Table 10.8 shows the median risk tolerance Q across participants depending on the choice set types (gains or mixed domain, and none, negative or positive skew choices). The risk tolerance is much higher in the mixed domain, suggesting that risk tolerance is a context specific measure. In the gains domain, risk tolerance drops for both positive and negative skew. In the mixed domain, risk tolerance drops for negative skew and increases for positive skew choice sets. This is consistent with previous findings that people have a preference and are willing to take higher risks for positive compared to negative skew prospects. The risk tolerance depends heavily on the specifics of the choice sets, and is not comparable across experiments. We find strong evidence of the influence of individual traits on the inferred risk tolerance, with increasing external locus of control, and increasing risk aversion as measured by personality scales linked to higher risk tolerance.

Table 10.8: Median inferred MVO-A parameters

		Gains			Mixed		
	Overall	None	Neg	Pos	None	Neg	Pos
Risk tolerance (Q)	27	14	11	7	66	41	77

Table 10.9: Median inferred CPT parameters

	α	γ	δ	λ	κ	η
CPT-S	0.63	0.78	0.98	1.1	0.31	0.41
CPT-A	0.50	0.70	0.70	2.10		

Table 10.10: Median inferred SPA parameters

	υ_{SP}	w_G	q_P	q_S	α_{SPA}	κ	η
SPA-S	0.26	0.58	0.14	3.05	-0.43	0.96	1.45

10.10.4 CPT parameters

Table 10.9 shows the median inferred parameters under the CPT-S and CPT-A models. The surprising finding is the lack of loss aversion under CPT-S, which is the second most preferred model. The median λ values for both experiments are 1.1, which indicates almost no loss aversion. The CPT-S model also demonstrates a higher bias towards diversification (lower η) compared to the SPA-S model.

10.10.5 SPA-S parameters

Table 10.10 shows the median inferred parameters under the SPA-S model, the most preferred model in these experiments. Under both experiments, υ_{SP} is low, about 0.26, indicating that people put greater weight on the aspiration criterion A, rather than the security-potential criterion SP. The criterion A is defined as $p(v \geq \alpha_{SP})$. The median values of α_{SP} in the two experiments is very close to zero (-0.43 and 0.01). Taken together, this suggests that participants were trying to maximize their probability of not incurring a loss.

10.10.6 Purchase cost discounting

Purchase cost discounting, is measured by κ, where the effective cost evaluated by participants is $[v(cost)]^{\kappa}$. Higher values of κ indicate lower susceptibility to cost framing on account of cost discounting, and values range from 0 to 1. In experiment 1, the median value of κ under the dominant SPA-S model was 0.96 (see table 10.10), indicating that people do not discount costs.

This does not preclude framing effects - framing effects can still be present because of segregated mental accounting of transaction costs which gives rise to framing effects when the valuation model is non-linear (as in SPA-S and CPT-S). Note that under the CPT-S model, median κ is much lower, at 0.31 (see table 10.9), indicating greater susceptibility to purchase cost discounting effects.

10.11 Conclusion

We report that in the experiment, susceptibility to transaction utility framing effects is highly heterogeneous, and is moderated by risk congruency, the degree of dissonance between self-expressed and elicited risk propensity. Depending on risk congruency, participants either under-weight or over-weight outcomes that are re-framed, resulting in a wide range of behavioral responses to these nudges. This has implications for the impact of transaction costs and incentive framing in situations where individuals face a conflict in terms of their perceived and experienced risk preferences, either because of internal or externally imposed risk preferences. In both framing cases, people reduced their level of diversification.

A novel finding was that measures of risk traits typically linked to risk aversion versus risk seeking behavior in single gamble choices were in fact linked to the level of diversification in our experiments, with risk seeking traits correlating to higher diversification, a counter intuitive trend. We find that people certainly did not follow the naive diversification or 1/N heuristic. How close participants came to mean-variance optimization of portfolio choices depended heavily on choice architecture, with better performance in mixed domains, and in non second order stochastic dominant choice sets (details in below sections). Performance in terms of optimization was not influenced by traditional risk aversion or risk seeking traits. Personality traits linking locus of control (internal vs external) to portfolio choice decisions showed increasing external locus of control leading to safer choices.

We analyze these behavioral patterns through the lens of a set of behavioral accounts of portfolio choice. We define a set of models based on choice bracketing, that is, whether prospects are evaluated individually or as a combined portfolio, and incorporate traditional measures such as traditional prospect theory, mean-variance optimization, and security-potential-aspiration (SPA) theory into these models. Our results show that people strongly preferred segregated evaluation of prospects, with preference for aggregated modes of evaluation being very low across both experiments. Under certain choice set conditions (such as non SOSD, or skewed choice sets), the preference for aggregated choice bracketing increased, but only marginally. We show that the level of diversification bias as measured formally in our model is linked to personality traits such as locus of control and risk aversion. We find evidence that preference for segregated versus aggregated choice evaluation is modulated by personality traits such as locus of control.

References

Arkes, H. R., Hirshleifer, D., Jiang, D., and Lim, S. (2008). Reference point adaptation: Tests in the domain of security trading. *Organizational Behavior and Human Decision Processes*, 105(1):67–81.

Arkes, H. R., Hirshleifer, D., Jiang, D., and Lim, S. S. (2010). A cross-cultural study of reference point adaptation: Evidence from china, korea, and the us. *Organizational Behavior and Human Decision Processes*, 112(2):99–111.

Bardolet, D., Fox, C. R., and Lovallo, D. (2011). Corporate capital allocation: A behavioral perspective. *Strategic Management Journal*, 32(13).

Beattie, J. and Loomes, G. (1997). The impact of incentives upon risky choice experiments. *Journal of Risk and Uncertainty*, 14(2).

Beesley, T., Nguyen, K. P., Pearson, D., and Le Pelley, M. E. (2015). Uncertainty and predictiveness determine attention to cues during human associative learning. *The Quarterly Journal of Experimental Psychology*, pages 1–25.

Benartzi, S. and Thaler, R. H. (2001). Naive diversification strategies in defined contribution saving plans. *American economic review*.

Bergert, F. B. and 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.

Bernasconi, M., Corazzini, L., and Seri, R. (2014). Reference dependent preferences, hedonic adaptation and tax evasion: Does the tax burden matter? *Journal of Economic Psychology*, 40:103–118.

Betsch, T., Brinkmann, B. J., Fiedler, K., and Breining, K. (1999). When prior knowledge overrules new evidence: Adaptive use of decision strategies and the role of behavioral routines. *Swiss Journal of Psychology/Schweizerische Zeitschrift für Psychologie/Revue Suisse de Psychologie*, 58(3):151.

Betsch, T., Haberstroh, S., Glöckner, A., Haar, T., and Fiedler, K. (2001). The effects of routine strength on adaptation and information search in recurrent decision making. *Organizational behavior and human decision processes*, 84(1):23–53.

Blais, A.-R. and Weber, E. U. (2006). A domain-specific risk-taking (dospert) scale for adult populations.

Bröder, A., Glöckner, A., Betsch, T., Link, D., and Ettlin, F. (2013). Do people learn option or strategy routines in multi-attribute decisions? the answer depends on subtle factors. *Acta psychologica*, 143(2):200–209.

Bröder, A. and Schiffer, S. (2003). Bayesian strategy assessment in multi-attribute decision making. *Journal of Behavioral Decision Making*, 16(3):193–213.

Bröder, A. and Schiffer, S. (2006a). Adaptive flexibility and maladaptive routines in selecting fast and frugal decision strategies. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32(4):904.

Bröder, A. and Schiffer, S. (2006b). Adaptive flexibility and maladaptive routines in selecting fast and frugal decision strategies. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32(4):904.

B'Tselem. https://www.btselem.org/statistics.

Byers, S. S., Groth, J. C., and Sakao, T. (2015). Using portfolio theory to improve resource efficiency of invested capital. *Journal of Cleaner Production*, 98:156–165.

Chetty, R. (2009). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annu. Rev. Econ.*, 1(1):451–488.

Chetty, R. (2015). Behavioral economics and public policy: A pragmatic perspective. *The American Economic Review*, 105(5):1–33.

Chetty, R., Friedman, J. N., Olsen, T., and Pistaferri, L. (2009a). Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: Evidence from danish tax records. Technical report, National Bureau of Economic Research.

Chetty, R., Looney, A., and Kroft, K. (2009b). Salience and taxation: Theory and evidence. *The American economic review*, 99(4):1145–1177.

Conway, P. and Gawronski, B. (2013). Deontological and utilitarian inclinations in moral decision making: a process dissociation approach. *Journal of personality and social psychology*, 104(2):216.

Deck, C. and Schlesinger, H. (2010). Exploring higher order risk effects. *The Review of Economic Studies*, 77(4).

DeMiguel, V., Garlappi, L., and Uppal, R. (2007). Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? *The review of Financial studies*, 22(5):1915–1953.

Doerrenberg, P., Peichl, A., and Siegloch, S. (2014). Sufficient statistic or not? the elasticity of taxable income in the presence of deduction possibilities.

Erev, I. (1998). Signal detection by human observers: a cutoff reinforcement learning model of categorization decisions under uncertainty. *Psychological review*, 105.

Erev, I. and Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological review*, 112(4):912.

Fischhoff, B. (1983). Predicting frames. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 9(1):103.

Fletcher, J. M., Frisvold, D. E., and Tefft, N. (2015). Non-linear effects of soda taxes on consumption and weight outcomes. *Health economics*, 24(5):566–582.

Frecknall-Hughes, J., Moizer, P., Doyle, E., and Summers, B. (2016). An examination of ethical influences on the work of tax practitioners. *Journal of Business Ethics*, pages 1–17.

Frederick, S. and Loewenstein, G. (1999). 16 hedonic adaptation. *Well-being: Foundations of hedonic psychology*, page 302.

Freire-Serén, M. J. and Panadés, J. (2013). Do higher tax rates encourage/discourage tax compliance? *Modern Economy*, 2013.

Friston, K. J., Lawson, R., and Frith, C. D. (2013). On hyperpriors and hypopriors: comment on pellicano and burr. *Trends Cogn. Sci*, 17.

Frydman, C. and Nave, G. (2016). Extrapolative beliefs in perceptual and economic decisions: evidence of a common mechanism. *Management Science*, 63(7):2340–2352.

Gallup (2015). http://www.gallup.com/poll/182423/perceptions-tax-fairness-diverging-income.aspx.

Gaydon, D., Meinke, H., Rodriguez, D., and McGrath, D. (2012). Comparing water options for irrigation farmers using modern portfolio theory. *Agricultural water management*, 115:1–9.

Gigerenzer, G. and Gaissmaier, W. (2011). Heuristic decision making. *Annual review of psychology*, 62:451–482.

Gigerenzer, G. and Todd, P. M. (1999). Fast and frugal heuristics: The adaptive toolbox. In *Simple heuristics that make us smart*, pages 3–34. Oxford University Press.

Glöckner, A. and Betsch, T. (2008). Multiple-reason decision making based on automatic processing. *Journal of experimental psychology: Learning, memory, and cognition*, 34(5):1055.

Gluth, S., Rieskamp, J., and Büchel, C. (2013). Neural evidence for adaptive strategy selection in value-based decision-making. *Cerebral Cortex*, page bht049.

Gneezy, U. (2005). Updating the reference level: Experimental evidence. In *Experimental business research*, pages 263–284. Springer.

Greene, J. D. (2007). Why are vmpfc patients more utilitarian? a dual-process theory of moral judgment explains. *Trends in cognitive sciences*, 11(8):322–323.

Greene, J. D. (2009). The cognitive neuroscience of moral judgment. *The cognitive neurosciences*, 4:1–48.

Haushofer, J., Biletzki, A., and Kanwisher, N. (2010). Both sides retaliate in the israeli–palestinian conflict. *Proceedings of the National Academy of Sciences*, 107(42):17927–17932.

Helson, H. (1948). Adaptation-level as a basis for a quantitative theory of frames of reference. *Psychological review*, 55(6):297.

Helson, H. (1964). Adaptation-level theory: an experimental and systematic approach to behavior.

Holt, C. A., Laury, S. K., et al. (2002). Risk aversion and incentive effects. *American economic review*, 92(5).

Hsee, C. K., Loewenstein, G. F., Blount, S., and Bazerman, M. H. (1999). Preference reversals between joint and separate evaluations of options: A review and theoretical analysis. *Psychological bulletin*, 125(5):576.

Jaeger, D. A. and Paserman, M. D. (2006). Israel, the palestinian factions, and the cycle of violence. *American Economic Review*, 96(2):45–49.

JASP-Team (2016). Jasp (version 0.8.0.0). *[Computer software]*.

Jeliazkov, I., Poirier, D. J., et al. (2008). Dynamic and structural features of intifada violence: A markov process approach. *Bayesian Analysis*, 3(1):63–77.

Johnson, M. D., Ilies, R., and Boles, T. L. (2012). Alternative reference points and outcome evaluation: The influence of affect. *Journal of Applied Psychology*, 97(1):33.

Jones, M. and Sugden, R. (2001). Positive confirmation bias in the acquisition of information. *Theory and Decision*, 50(1):59–99.

Kabanikhin, S. I. (2008). Definitions and examples of inverse and ill-posed problems. *Journal of Inverse and Ill-Posed Problems*, 16(4):317–357.

Kahneman, D. (1992). Reference points, anchors, norms, and mixed feelings. *Organizational behavior and human decision processes*, 51(2):296–312.

Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the econometric society*, pages 263–291.

Kahneman, D. and Tversky, A. (1984). Choices, values, and frames. *American psychologist*, 39(4):341.

Kastlunger, B., Lozza, E., Kirchler, E., and Schabmann, A. (2013). Powerful authorities and trusting citizens: The slippery slope framework and tax compliance in italy. *Journal of Economic Psychology*, 34:36–45.

Kirchler, E., Hoelzl, E., and Wahl, I. (2008). Enforced versus voluntary tax compliance: The slippery slope framework. *Journal of Economic Psychology*, 29(2):210–225.

Klayman, J. and Ha, Y.-W. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review*, 94:211.

Kleinhans, N. M., Johnson, L. C., Richards, T., Mahurin, R., Greenson, J., Dawson, G., and Aylward, E. (2009). Reduced neural habituation in the amygdala and social impairments in autism spectrum disorders. *American Journal of Psychiatry*, 166.

Kogler, C., Batrancea, L., Nichita, A., Pantya, J., Belianin, A., and Kirchler, E. (2013). Trust and power as determinants of tax compliance: Testing the assumptions of the slippery slope framework in austria, hungary, romania and russia. *Journal of Economic Psychology*, 34:169–180.

Lawson, R. P., Rees, G., and Friston, K. J. (2014). An aberrant precision account of autism. *Frontiers in human neuroscience*, 8.

Le Pelley, M., Beesley, T., and Griffiths, O. (2011). Overt attention and predictiveness in human contingency learning. *Journal of Experimental Psychology: Animal Behavior Processes*, 37:220.

Lee, M. D. (in press). Bayesian methods in cognitive modeling. In *The Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience*. Fourth edition.

Lee, M. D., Newell, B. R., and Vandekerckhove, J. (2014). Modeling the adaptation of search termination in human decision making. *Decision*, 1(4):223.

Lieder, F. and Griffiths, T. L. (2015). When to use which heuristic: A rational solution to the strategy selection problem. In *Proceedings of the 37th annual conference of the cognitive science society*, pages 1362–1367.

Lin, J. (1991). Divergence measures based on the shannon entropy. *IEEE Transactions on Information theory*, 37(1):145–151.

Lopes, L. L. and Oden, G. C. (1999). The role of aspiration level in risky choice: A comparison of cumulative prospect theory and sp/a theory. *Journal of mathematical psychology*, 43(2):286–313.

Lozza, E., Kastlunger, B., Tagliabue, S., and Kirchler, E. (2013). The relationship between political ideology and attitudes toward tax compliance: The case of italian taxpayers. *Journal of Social and Political Psychology*, 1(1):51–73.

Manning, C., Kilner, J., Neil, L., Karaminis, T., and Pellicano, E. (2016). Children on the autism spectrum update their behaviour in response to a volatile environment. *Developmental Science*.

Markowitz, H. (1952). Portfolio selection. *The journal of finance*, 7(1).

Marquis, M. S. and Holmer, M. R. (1996). Alternative models of choice under uncertainty and demand for health insurance. *The Review of Economics and Statistics*, pages 421–427.

Matthews, W. J. and Stewart, N. (2009). Psychophysics and the judgment of price: Judging complex objects on a non-physical dimension elicits sequential effects like those in perceptual tasks. *Judgment and Decision Making*, 4(1):64.

Mellers, B. A., Ordoñez, L. D., and Birnbaum, M. H. (1992). A change-of-process theory for contextual effects and preference reversals in risky decision making. *Organizational Behavior and Human Decision Processes*, 52(3).

Mistry, P. K., Skewes, J., and Lee, M. D. (2018). An adaptive signal detection model applied to perceptual learning. *In proceedings of the 40th Annual Conference of the Cognitive Science Society. Austin, TX: Cognitive Science Society*.

Mistry, P. K. and Trueblood, J. S. (2015). Reconstructing the bayesian adaptive toolbox: Challenges of a dynamic environment and partial information acquisition. *In proceedings of the 37th Annual Conference of the Cognitive Science Society. Austin, TX: Cognitive Science Society*.

Mistry, P. K. and Trueblood, J. S. (2017). An investigation of factors that influence resource allocation decisions. *In proceedings of the 39th Annual Conference of the Cognitive Science Society. Austin, TX: Cognitive Science Society*.

Mistry, P. K. and Trueblood, J. S. (2018). An investigation of factors that influence resource allocation decisions. *In proceedings of the 40th Annual Conference of the Cognitive Science Society. Austin, TX: Cognitive Science Society*.

Muehlbacher, S., Kirchler, E., and Schwarzenberger, H. (2011). Voluntary versus enforced tax compliance: Empirical evidence for the slippery slope framework. *European Journal of Law and Economics*, 32(1):89–97.

Newell, B. R. and Lee, M. D. (2009). Learning to adapt evidence thresholds in decision making. In *Proceedings of the 31st Annual Conference of the Cognitive Science Society. Austin, TX: Cognitive Science Society*.

Newell, B. R. and Shanks, D. R. (2003). Take the best or look at the rest? factors influencing" one-reason" decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29(1):53.

Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of general psychology*, 2(2):175.

Norman, K. A., Detre, G. J., and Polyn, S. M. (2008). Computational models of episodic memory. In Sun, R., editor, *The Cambridge handbook of computational psychology*, pages 189–224. Cambridge University Press, New York.

OConnor, K. (2012). Auditory processing in autism spectrum disorder: a review. *Neuroscience & Biobehavioral reviews*, 36.

Palmer, C. J., Paton, B., Kirkovski, M., Enticott, P. G., and Hohwy, J. (2015). Context sensitivity in action decreases along the autism spectrum: a predictive processing perspective. *Proceedings of the Royal Society of London B: Biological Sciences*, 282.

Palminteri, S., Lefebvre, G., Kilford, E. J., and Blakemore, S.-J. (2016). Confirmation bias in human reinforcement learning: evidence from counterfactual feedback processing. *bioRxiv*, page 090654.

Paxton, J. M., Bruni, T., and Greene, J. D. (2014). Are counter-intuitivedeontological judgments really counter-intuitive? an empirical reply to. *Social cognitive and affective neuroscience*, 9(9):1368–1371.

Pearson, J. M., Heilbronner, S. R., Barack, D. L., Hayden, B. Y., and Platt, M. L. (2011). Posterior cingulate cortex: adapting behavior to a changing world. *Trends in cognitive sciences*, 15(4):143–151.

Pellicano, E. and Burr, D. (2012). When the world becomes too real: a bayesian explanation of autistic perception. *Trends in cognitive sciences*, 16.

Phillips, P. J. (2009). Applying modern portfolio theory to the analysis of terrorism. computing the set of attack method combinations from which the rational terrorist group will choose in order to maximise injuries and fatalities. *Defence and Peace Economics*, 20(3):193–213.

Plummer, M. et al. (2003). Jags: A program for analysis of bayesian graphical models using gibbs sampling. In *Proceedings of the 3rd international workshop on distributed statistical computing*, volume 124, page 125. Vienna.

Prinz, A., Muehlbacher, S., and Kirchler, E. (2014). The slippery slope framework on tax compliance: An attempt to formalization. *Journal of Economic Psychology*, 40:20–34.

Racey, D., Young, M. E., Garlick, D., Pham, J. N.-M., and Blaisdell, A. P. (2011). Pigeon and human performance in a multi-armed bandit task in response to changes in variable interval schedules. *Learning & behavior*, 39(3):245–258.

Rhoades, S. A. (1993). The herfindahl-hirschman index. *Fed. Res. Bull.*, 79:188.

Rieskamp, J. (2008). The importance of learning when making inferences. *Judgment and Decision Making*, 3(3):261–277.

Rieskamp, J. and Otto, P. E. (2006). Ssl: a theory of how people learn to select strategies. *Journal of Experimental Psychology: General*, 135(2):207.

Robinson, S. L. and Morrison, E. W. (1995). Psychological contracts and ocb: The effect of unfulfilled obligations on civic virtue behavior. *Journal of organizational behavior*, 16(3):289–298.

Rolison, J. J., Evans, J. S. B., Walsh, C. R., and Dennis, I. (2011). The role of working memory capacity in multiple-cue probability learning. *The Quarterly Journal of Experimental Psychology*, 64(8):1494–1514.

Roques, F. A., Newbery, D. M., and Nuttall, W. J. (2008). Fuel mix diversification incentives in liberalized electricity markets: A mean–variance portfolio theory approach. *Energy Economics*, 30(4):1831–1849.

Rotter, J. B. (1966). Generalized expectancies for internal versus external control of reinforcement. *Psychological monographs: General and applied*, 80(1).

Saez, E., Slemrod, J., and Giertz, S. H. (2012). The elasticity of taxable income with respect to marginal tax rates: A critical review. *Journal of economic literature*, 50(1):3–50.

Sidani, Y. M., Ghanem, A. J., and Rawwas, M. Y. (2014). When idealists evade taxes: the influence of personal moral philosophy on attitudes to tax evasion–a lebanese study. *Business Ethics: A European Review*, 23(2):183–196.

Sitzia, S. and Zizzo, D. J. (2012). Price lower and then higher or price higher and then lower? *Journal of Economic Psychology*, 33(6):1084–1099.

Skewes, J. C. and Gebauer, L. (2016). Brief report: suboptimal auditory localization in autism spectrum disorder: support for the bayesian account of sensory symptoms. *Journal of autism and developmental disorders*, 46.

Soliman, A., Jones, P., and Cullis, J. (2014). Learning in experiments: Dynamic interaction of policy variables designed to deter tax evasion. *Journal of Economic psychology*, 40:175–186.

Steyvers, M., Lee, M. D., and Wagenmakers, E.-J. (2009). A bayesian analysis of human decision-making on bandit problems. *Journal of Mathematical Psychology*, 53(3):168–179.

Stüttgen, M. C., Yildiz, A., and Güntürkün, O. (2011). Adaptive criterion setting in perceptual decision making. *Journal of the experimental analysis of behavior*, 96.

Symmonds, M., Wright, N. D., Bach, D. R., and Dolan, R. J. (2011). Deconstructing risk: separable encoding of variance and skewness in the brain. *Neuroimage*, 58(4).

Teder-Sälejärvi, W. A., Pierce, K. L., Courchesne, E., and Hillyard, S. A. (2005). Auditory spatial localization and attention deficits in autistic adults. *Cognitive Brain Research*, 23.

Thaler, R. (1985). Mental accounting and consumer choice. *Marketing science*, 4(3):199–214.

Thaler, R. H. (1999). Mental accounting matters. *Journal of Behavioral decision making*, 12(3):183.

Thaler, R. H. (2008). Mental accounting and consumer choice. *Marketing Science*, 27(1):15–25.

Thaler, R. H. and Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press.

Treisman, M. and Williams, T. C. (1984). A theory of criterion setting with an application to sequential dependencies. *Psychological review*, 91.

Turi, M., Burr, D. C., Igliozzi, R., Aagten-Murphy, D., Muratori, F., and Pellicano, E. (2015). Children with autism spectrum disorder show reduced adaptation to number. *Proceedings of the National Academy of Sciences*, 112.

Turner, B. M., Van Zandt, T., and Brown, S. (2011). A dynamic stimulus-driven model of signal detection. *Psychological review*, 118.

Tversky, A. and Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, 5(4):297–323.

Van de Cruys, S., Evers, K., Van der Hallen, R., Van Eylen, L., Boets, B., de Wit, L., and Wagemans, J. (2014). Precise minds in uncertain worlds: predictive coding in autism. *Psychological review*, 121.

Vandekerckhove, J. (2014). A cognitive latent variable model for the simultaneous analysis of behavioral and personality data. *Journal of Mathematical Psychology*, 60:58–71.

Villas-Boas, S. B., Berck, P., Stevens, A., and Moe-Lange, J. (2016). Measuring consumer responses to a bottled water tax policy. *American ournal of Ag Economics*.

Wächter, T., Lungu, O. V., Liu, T., Willingham, D. T., and Ashe, J. (2009). Differential effect of reward and punishment on procedural learning. *Journal of Neuroscience*, 29(2):436–443.

Wahl, I., Kastlunger, B., and Kirchler, E. (2010). Trust in authorities and power to enforce tax compliance: An empirical analysis of the slippery slope framework. *Law & Policy*, 32(4):383–406.

Walsh, M. M. and Gluck, K. A. (2016). Verbalization of decision strategies in multiple-cue probabilistic inference. *Journal of Behavioral Decision Making*, 29(1):78–91.

Wichmann, F. A. and Hill, N. J. (2001). The psychometric function: I. Fitting, sampling and goodness-of-fit. *Perception & Psychophysic*, 63:1293–1313.

Wilson, R. C. and Niv, Y. (2011). Inferring relevance in a changing world. *Frontiers in human neuroscience*, 5.

Yechiam, E. and Busemeyer, J. R. (2008). Evaluating generalizability and parameter consistency in learning models. *Games and Economic Behavior*, 63(1):370–394.