

UC Irvine

UC Irvine Electronic Theses and Dissertations

Title

Response time and variability in the psychological inference of decision making under uncertainty and memory.

Permalink

<https://escholarship.org/uc/item/0ft9v090>

Author

Banavar, Nidhi Venkatanarayan

Publication Date

2023

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA,
IRVINE

Response time and variability in the psychological inference of decision making under
uncertainty and memory.

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Cognitive Sciences

by

Nidhi Venkatanarayan Banavar

Dissertation Committee:
Assistant Professor Aaron M. Bornstein, Chair
Professor Michael D. Lee
Professor Joachim Vandekerckhove

2023

DEDICATION

To my parents.

TABLE OF CONTENTS

	Page
LIST OF FIGURES	v
LIST OF TABLES	ix
ACKNOWLEDGMENTS	xii
VITA	xiv
ABSTRACT OF THE DISSERTATION	xvii
1 Introduction	1
2 Variability in Complex Constructs: Inferring Risk Preference and Temporal Discounting	5
2.1 Introduction	6
2.2 Delineating Risk	13
2.3 Delineating Temporal Discounting	17
2.4 Variability in Experiments: Measurements	20
2.5 Trial-Trial Temporal Dependence	23
2.6 Conclusion	29
3 Independent not Irrelevant: Trial order causes systematic misestimation of economic choice traits in humans	31
3.1 Introduction	32
3.2 Methods	36
3.2.1 Experiments and Data	36
3.2.2 Models: Choice Behavior	39
3.2.3 Models: Response Times	43
3.3 Results	45
3.3.1 Sequential Effects in Choice Behavior	45
3.3.2 Choice and Response Time	48
3.4 Discussion	54
4 A response time model of memory discrimination.	62
4.1 Introduction	62
4.2 Methods	65

4.2.1	Experiments and Data	65
4.2.2	Response Time Model: The Linear Ballistic Accumulator	65
4.3	Results	68
4.3.1	Model Free.	68
4.3.2	Model Based.	70
4.4	Discussion	81
5	Putting it together: Motivating sequential effects in older adults and preliminary analyses in economic choice and memory discrimination.	84
5.1	Introduction	84
5.2	Methods	89
5.2.1	Task Descriptions and Analyses	90
5.3	Results	91
5.3.1	Ambiguity	91
5.3.2	Mnemonic Similarity Task	94
5.3.3	Decision making under uncertainty and MST	97
5.4	Discussion	98
6	Conclusion	100
	Bibliography	103

LIST OF FIGURES

	Page
<p>3.1 Example of how trial order may influence decision making (choice and RT modeling parameters) in ITC. The subject participates in the experiment on November 22nd. (A) On trial $j-1$, they have to choose between \$5 today and \$10 in 30 days. They simulate their future 30 days from now, December 22nd, and imagine what they could buy. This is a difficult choice for the subject and deliberation is captured here by a slow rate of evidence accumulation. The subject ultimately decides to take the \$10 and wait. (B) On the next trial j, the choice is between \$5 today and \$13 in 35 days. Standard models assume that the previous trial should have no influence on the current problem. The spillover hypothesis predicts that a subject could perceive and evaluate the current decision relative to the preceding choice, effectively reusing the outcome of their previous simulation instead of starting afresh: here demonstrated by a steeper drift rate and then a repeated choice.</p>	34
<p>3.2 Example trials for all three tasks. (A) E1: Example RISK trial in the Loss domain where a subject has an unlimited amount of time from stimulus onset to choose between either the gamble on the Left or the Right. The Expected Value maximizing (and correct) choice is the gamble on the Left (figure from (Guan, 2019)). (B) E2: Example ITC trial where the subject has up to 6s to chose between \$9 today or \$30 in 42 days (figure from (L. Hunter et al., 2018)). (C) E3: Example AMB ambiguous trial where the subject has up to 3 seconds to make a choice between a certain reward of \$3 and a chance to win \$11 by playing the lottery. The Expected Value maximizing (but not necessarily “correct”) choice is the lottery reward.</p>	38
<p>3.3 Participants in RISK and AMB tasks make distinct choices depending on trial sequence. (A) In the RISK task, they are more likely to make the EV maximizing choice between two gambles when the previous trial was more difficult (i.e. choices less distinct). (B) In AMB, they are more likely to pick the lottery option as opposed to the fixed when the previous choice was easier (i.e. less difference between options). Conversely, there was no significant difference between EV maximizing choices when considering relative increases or decreases in trial entropy (A) or trial ambiguity (B). Overlaying the violin plots are the median and IQR.</p>	47

3.4	Choice and RT in ITC change as a function of trial properties. In the intertemporal choice task, subjects took a median of 1.64s when they chose the Smaller Sooner (SS) option, and 1.71s when they chose the Larger Later (LL) option. Grey histograms are RTs for all subjects, all trials. Negative RTs correspond to making the SS choice, while positive correspond to LL. (A) When the current trial increased in value difference from the previous trial, subjects took a median of 1.64s to make a SS choice and 1.68s to make a LL choice (overlaid orange histograms). (B) However, when the current trial decreased in value difference from the previous trial, subjects took a median of 1.35s to make a SS choice and 1.77s to make a LL choice (overlaid brown histograms).	48
3.5	Participants in the AMB task show sequential dependencies on the drift rate EV term: “main effects” of value difference but not ambiguity difference. Sorted posterior 95% Credible Intervals of sequential effects on the drift rate Expected Value term β_1 , Equation 8, when successive trials (A) increase in value difference, (B) decrease in value difference, (C) increase in ambiguity difference, and (D) decrease ambiguity difference as summarized in table S4.	55
3.6	AMB task: Magnitude and interpretation changes in (A, B) ambiguity and (C, D) risk tolerance. <i>Left:</i> Ratio of (A) ambiguity tolerance and (C) risk tolerance estimates: $\log(\text{simulated choice set}/\text{observed data})$. We plot median ratios and IQRs from 1000 simulated choice sets for 98 subjects. <i>Right:</i> The percentage of simulation-fit parameters that change interpretations in (B) ambiguity and (D) risk attitudes when compared to parameter fits in the original data. Subjects re-sorted by effect size in each plot.	56
4.1	Schematic of the Linear Ballistic Accumulator. As the standard MST is 3AFC, we allow for three accumulators for each response type. We further allow the drift rate (rate of evidence accumulation) and the upper limit of the starting point (tendency to make a type of response) to vary for each subject and accumulator. Boundary (amount of evidence needed to make a response) and non-decision time (non-decision-relevant processes, not pictured) vary only at the subject level.	68
4.2	Choice proportions change as a function of response time. We find, in general, that faster choices tend to be <i>Repeats</i>	71
4.3	Model posteriors for Experiments 1-3, Experiment 5 We find overall that LBA posteriors follow qualitatively the same patterns across experiments. The left most three grey violins are the start point upper bound (color coded by response type: Repeat, Lure, Foil). Next are the drift rates, similarly color coded. Finally, we have the boundary and non-decision time.	72

4.4	LDI correlates with drift rate <i>and</i> start point upper bounds in E:1-4. We collapse across all 4 experiments ($n = 233$) and correlate mean drift rate and start point upper bound with LDI. We find statistically significant correlations between the LDI and drift rates for all accumulators. We also find significant correlations between the LDI and Repeat Foil accumulator start point upper bounds. Correlations shown in plots are statistically significant ($p < 0.01$) after adjusting for multiple comparisons.	74
4.5	LDI correlates with drift rate <i>and</i> start point upper bounds in lifespan sample. In a dataset comprised of older and younger adults, we find similar qualitative relationships between the LDI and accumulator drift rates/start point upper bounds. Correlations shown in plots are statistically significant ($p < 0.01$) after adjusting for multiple comparisons.	75
4.6	E1-E4: LDI-Drift Rate correlations for Repeat and Lure accumulators are stronger than LDI-Start Point correlations. E5: No significant correlation differences. We plot histograms showing bootstrapped correlation differences between LDI and drift rate, and LDI and start point (left: Stark 1-4, right: E5: Wahlheim et al). Repeat and Lure accumulator 95% CIs in E1 – 4 do not contain 0, suggesting we can reject the null of no correlation difference.	76
4.7	E5: LDI tends to correlate negatively with posterior hippocampal RSFC and positively with other regions. We show a symmetric correlation matrix plot where each square represents the correlation between the row-column hippocampal resting state functional connectivity and the LDI. Cells colored darker yellow show stronger positive correlations, and darker purple stronger negative correlations.	79
4.8	E5: Lure accumulator drift rates mostly correlate positively with hippocampal RSFC while Repeat accumulator drift rates correlate negatively. We show a symmetric correlation matrix plot where each square represents the correlation between the row-column hippocampal resting state functional connectivity and the respective accumulator drift rate. The Repeat and Lure accumulator – RSFC correlations also seem to be stronger than the Foil accumulator – RSFC correlations.	80
4.9	E5: Start Point – hippocampal RSFC correlations appear to be weaker than Drift Rate – hippocampal RSFC correlations. We show a symmetric correlation matrix plot where each square represents the correlation between the row-column hippocampal resting state functional connectivity and the respective accumulator drift rate. We also see greater variability in the qualitative patterns in the start point – RSFC correlations (i.e. evidence of both negative and positive correlations).	81
5.1	Example trial of new AMB task for Older Adults. The visual changes we make include delineating the left and right options more clearly (instead of having the lottery placed in the center as in Chapter 3) and making the font sizes generally larger.	90

5.2	AMB Task: Magnitude and interpretation changes in (A,B) ambiguity and (C,D) risk tolerance. Ratio of (A) ambiguity tolerance and (C) risk tolerance estimates: $\log(\text{simulated choice set}/\text{observed data})$. We plot median ratios and IQRs from 1000 simulated choice sets. <i>Right</i> : The percentage of simulation-fit parameters that change interpretations in (B) ambiguity and (D) risk attitudes when compared to parameter fits in the original data. Subjects re-sorted by effect size in each plot (i.e. differently sorted in A/C vs. B/D).	94
5.3	Choice proportions change as a function of response time. In this dataset, we find that fastest responses tend to be Repeat and the slowest tend to be Lures – indeed participants make no Lure responses in the fastest 2 RT bins.	95
5.4	LBA model posteriors. We find that the Lure accumulator start point upper bound and drift rate tend to be lower than the equivalent parameters for the Repeat and Foil accumulators, consistent with expectations.	96
5.5	LDI only correlates with Lure and Foil accumulators. In the data we have collected thus far, we find different qualitative relationships between LDI and the LBA parameters of interest. Correlations shown in plots are statistically significant ($p < 0.01$) after adjusting for multiple comparisons.	97

LIST OF TABLES

	Page
<p>3.1 Trial properties considered for sequential effects and the choice parameters on which we test for these effects. <i>Upper</i> Differences in trial properties considered as potential drivers of sequential effects. $H = -\sum p \log(p)$ is the Shannon Entropy of a gamble. <i>Lower</i> Parameters we simultaneously test for sequential effects by allowing them to vary trial-trial as a function of relative differences in properties as described in Table 1 a. We test the first three parameters in all three tasks. The final three are task-dependent (RISK: α; ITC: k; AMB: α, β).</p>	42
<p>3.2 ITC + AMB: Drift rate decompositions capture meaningful variance in both ITC and AMB. Each cell shows aggregate posterior means (95% Credible Intervals) for Drift Rate decompositions and other DDM parameters. Bolded parameters are ones we also test for sequential effects. . . .</p>	49
<p>3.3 ITC + AMB: Most participants are sensitive to trial sequences. The majority of subjects show sensitivity to trial sequences in both ITC and AMB tasks. Each cell shows the proportion of subjects that demonstrated sequential effects ($BF > 3$) on any one of the drift rate decomposition parameters or bias. Each row represents specific successive trial properties (e.g. <i>value</i> \uparrow subsets successive trials that increased in value difference as noted in Equation 3.4).The top four rows can be thought of as “main effects” of specific trial properties and the bottom four “interactions.”</p>	50
<p>3.4 E3: Evolution of sequential effects over task. We find that the proportion of subjects that show sensitivity to sequential effects changes over the course of the task, with the majority of recovered sensitivities crossing a statistically meaningful threshold presenting when considering all trials in the experiment. Overall, subjects tend to be sensitive increases in ambiguity . . .</p>	52
<p>3.5 ITC: Widespread sequential effects across drift rate and bias terms. Participants in ITC show sequential effects across all trial sequence types and all four RT parameters (see Equation 3.5). All results reported in a cell are posterior sequential effect group means (<i>mean (standard deviation)</i>) and have a $BF > 3$ if they are non-zero. If instead there is a H_0 then we find evidence ($BF > 3$) in favor of the null. Finally, if a cell contains an asterisk (*) then the data does not contain enough evidence to favor either the null or alternative hypothesis.</p>	53

3.6	AMB: Sequential effects are restricted to drift rate decomposition parameters. Sequential effects are more selective compared to ITC. Specifically, they are related to online evidence accumulation and are only on the terms that include trial property information (see Equation 3.6). Cell interpretations are as in Table 3.5.	54
4.1	MST experiments modeled. The <i>Source</i> column indicates the researchers who originally collected the data. Experiment 5 has the additional suffix “rsfc” to indicate that the researchers also collected resting state functional connectivity measures. E5 also included a lifespan sample.	66
4.2	Accuracy across experiments by accumulator. We show median(IQR) accuracy for each response type. Unsurprisingly across all experiments, subjects are more accurate when identifying repeat and foil stimuli vs. lure stimuli. However, lure stimuli accuracy is also the most variable.	69
4.3	Median Response Time across experiments by accumulator. We show median(IQR) RT for each response type. We find that, on average, participants tend to take longer when they make a Lure response especially when compared to when they make a Repeat response	69
4.4	Choice Proportions across experiments by accumulator. We show median(IQR) choice proportions for each response type. We see that across experiments, participants to most often classify stimuli as Repeat (recall that the true proportion of repeated stimuli presented during test is 0.33).	69
4.5	Lure Discrimination Indices for each experiment. We show median(IQR) LDI, which is the current standard metric for summarizing choice behavior in this experiment. We find no statistical differences in median LDI across experiments due to how much the LDI varies within each experiment.	70
4.6	Younger adult LDI-LBA correlations follow the same qualitative patterns as in E1-E4. None of the correlations are statistically significant, however, after adjusting for multiple comparisons.	77
4.7	Older adult LDI-LBA correlations follow the same qualitative patterns as in E1-E4. None of the correlations are statistically significant, however, after adjusting for multiple comparisons.	78
4.8	Significant Hippocampus resting state functional connectivity and LDI/LBA parameters: <i>Only drift rates correlate with same RSFC regions as LDI.</i> Each row delineates which parameter correlated significantly ($p < 0.05$) with RSFC between the two listed hippocampus subregions. LH and RH are shortform for Left Hemisphere and Right Hemisphere respectively.	82
5.1	AMB: Like in Chapter 2, drift rate decompositions capture meaningful variance in AMB. Each cell shows aggregate posterior medians (Interquartile Range) for Drift Rate decompositions and other DDM parameters. Bolded parameters are ones we also test for sequential effects.	92

5.2 **Older adults are sensitive to most trial combinations.** We find that older adults are most sensitive to relative differences in reward *values* across trials. 93

ACKNOWLEDGMENTS

Thank you to everyone who has been a part of this journey whether expressly mentioned here or not. This thesis would not have been possible without your support.

I first thank my advisor Aaron Bornstein to whom I am extremely grateful. Thank you for your patience, knowledge, and for the opportunity to think freely. It has further been my fortune to work with someone who strives to be a good person as well as a good scientist.

Thank you to my secondary advisor Michael Lee for your support, intellectually and otherwise, over the course of my graduate career. Thank you also for the way you carefully structured lab meetings. I recognize it as a significant source of professional growth.

Finally, to the third member of my committee, Joachim Vandekerckhove. Thank you for always lending an ear, for the sound advice, and the levity. The path to Kamata in spring.

I would not have been able to begin a graduate program in Cognitive Sciences were it not for the three people that gave me a start in the field. Thank you to Anna Konova, Candace Raio, and Paul Glimcher for my formative two years at the Glimcher Lab. Thank you for teaching me what good science and good mentorship is. Thank you, most of all, for giving me a chance and helping me grow as a researcher and as an individual with so much care.

Over the course of my graduate career, I have been a member of two labs: the Computational Cognitive Neuroscience Lab and the Cognitive Modeling Lab. Thank you to all the lab members and affiliates, including those not in academia, for fruitful discussions ranging from scientific to mundane.

This dissertation is made possible through generous support from the Social Science Merit Fellowship, the Associate Dean's Fellowship, and the Graduate Dean's Recruitment Fellowship. Thank you to Aaron Bornstein for additional funding from NIA R21AG072673 and the Brain and Behavior Research Foundation NARSAD Young Investigator Award.

I would like to thank many individuals for their particular impact on and support of me. Outside of academia, I thank Laura Noceta, Tanvi Shetty, and Shyane Weaver for their years of friendship. Thank you to Pradeep Bhatta, who has always prompted me to critically consider the "value" of my research in the kindest way.

Thank you to Aakriti Kumar for your camaraderie and mischief especially as we forged our way through the statistics program. Thank you to Holly Westfall for your kindness and particularly for our conversations during the peak of the pandemic. Thank you to Jungsun Yoo for your earnestness and your company in shenanigans. Thank you to Priyam Das for your support and for the brain-storming sessions on how to navigate some of the more complicated circumstances in academia. Thank you to Ellen Shi for rich discussions and loud laughs. Thank you to Lauren Montgomery for your comprehension, joie de vivre, and generosity. Thank you to Jingyi Wu for far-ranging conversations, inspiration, and shared vulnerability. Thank you to Adriana Felisa Chávez De la Peña for your mirth, thoughtfulness,

and empathy. Thank you to Jaime Osvaldo Islas Farias for helping my inner nerd to resurface, for the delicious food, and for your unwavering friendship. Finally, thank you to Nora Harhen, whose solidarity, piercing insight, and gregarious laughter have been invaluable.

Thank you to my paternal grandparents Leela Venkataramaiah and B. V. Venkataramaiah for their support even after being freed from the confines of the physical world. I would like to give a particular acknowledgement to Leela Ajji for everything you did and still do. Thank you to my maternal grandparents B. R. Ramamurthy and Jaya Ramamurthy for continuing to motivate and support me.

Thank you to my sister, Sharanya Banavar, for teaching and learning with me some of life's most valuable lessons.

Finally, my deepest gratitude to my parents B. V. Venkatanarayan and Kala Venkatanarayan. Anything worthwhile in this thesis and in "me" is because of and for you.

VITA

Nidhi Venkatanarayan Banavar

EDUCATION

Doctor of Philosophy in Cognitive Sciences University of California, Irvine	2023
Master of Science in Statistics University of California, Irvine	2022
Bachelor of Arts in Mathematics Thomas Edison State University	2016

RESEARCH EXPERIENCE

Graduate Research Assistant Cognitive Computational Neuroscience Lab, University of California, Irvine	2019–2023 <i>Irvine, California</i>
Assistant Research Scientist Glimcher Lab, New York University	2017–2018 <i>New York, NY</i>
Research Assistant Glimcher Lab, New York University	2016–2017 <i>New York, NY</i>

TEACHING EXPERIENCE

Teaching Assistant University of California, Irvine	2018–2022 <i>Irvine, California</i>
---	---

REFEREED JOURNAL PUBLICATIONS AND PREPRINTS

Multi-plasticities: Distinguishing context-specific habits from complex perseverations. 2023

Banavar, N. V., Bornstein, A.M.

<https://doi.org/10.31234/osf.io/t7vsc>

Independent, not irrelevant: Trial order causes systematic misestimation of economic choice traits. 2023

Banavar, N. V., Bornstein, A.M.

<https://doi.org/10.31234/osf.io/a8gz3>

Variability in Complex Constructs: Inferring risk preference and temporal discounting. 2023

Banavar, N. V., Bornstein, A.M.

<https://doi.org/10.31234/osf.io/zdq5v>

A neuroeconomic signature of opioid craving: how fluctuations in craving bias drug-related and nondrug-related value 2021

Biernacki, K., Lopez-Guzman, S., Messinger, J., *Banavar, N. V.*, Rotrosen, J., Glimcher, P.W., Konova, A.B.

Neuropsychopharmacology

Measurement of 21cm brightness fluctuations at $z = 0.8$ in cross-correlation 2013

Masui, K.W., Switzer, E.R., *Banavar, N.*, et. al.

The Astrophysical Journal Letters 763(1)

REFEREED CONFERENCE PUBLICATIONS

Response time modeling provides stable and mechanistically interpretable measures of individual differences in behavioral pattern separation. 2022

Banavar, N. V., Bornstein, A.M.

Proceedings of the 20th International Conference on Cognitive Modeling.

Decision difficulty modulates the re-use of computations across trials in non-sequential decision tasks. 2022

Banavar, N. V., Bornstein, A.M.

Proceedings of the 5th Multidisciplinary Conference on Reinforcement Learning and Decision Making.

Sequential effects in non-sequential tasks.

2021

Banavar, N. V., Lee, M.D., & Bornstein, A.M.

Proceedings of the 19th International Conference on Cognitive Modeling.

ABSTRACT OF THE DISSERTATION

Response time and variability in the psychological inference of decision making under uncertainty and memory.

By

Nidhi Venkatanarayan Banavar

Doctor of Philosophy in Cognitive Sciences

University of California, Irvine, 2023

Assistant Professor Aaron M. Bornstein, Chair

In this dissertation, we consider ways to incorporate “meaningful” variability in our inferences about human decision making under uncertainty and memory. In particular, we use response time modeling to incorporate theoretically and empirically motivated sources of variability to develop cognitive models of decision making. In the first two chapters, we motivate and demonstrate the importance of explicitly considering sequential effects in decision making under uncertainty. We show that, even in experiments designed to maximize the perceived independence in successive choices an individual makes, individuals are sensitive to previous trials. Importantly, some individuals are substantively sensitive to the point that our qualitative interpretations about their behavior changes (e.g. an individual may be re-classified from risk seeking to risk averse). In the next chapter, we develop a model of choice and response time to allow for more psychologically nuanced interpretations of individual abilities in memory discrimination. We apply this model across several datasets and use it to further test the validity of the current gold standard measure, derived from choice behavior alone. In the final chapter, we combine our previous analyses and explore how sequential effects and memory discrimination may change in aging populations. Together, we demonstrate how response time models can be easily adapted to incorporate psychologically important considerations of and novel inferences about human decision making.

Chapter 1

Introduction

Every empirical researcher is tasked with the careful consideration of her object of study: how to appropriately define it, to operationalize it, to measure it, to model it, to generalize it. Such questions often fall under the domain of measurement and its science, metrology.

Any measurement result reports information that is meaningful only in the context of a metrological model, such a model being required to include specification for all the entities that explicitly or implicitly appear in the expression of the measurement result – Luca Mari

Such a strict definition of what it means to build a model that appropriately measures something, however, may have limited use when considering processes that are inherently stochastic and subject to multiple difficult to define and distinguish sources of variability. Indeed, philosophers of science that study and advocate for “model-based” measurement, as would be necessary in such circumstances, make no such stringent demands (Tal, 2020).

This thesis is concerned with one such set of processes grouped together under the category of [human] decision making. It is difficult to define decision making outside of the directly

observable: a (partially) autonomous agent is presented with one or more options after which they have to make a choice. Experimental researchers have perpetually had to grapple with the question of how to appropriately translate this concept/set of processes into the laboratory for a more precise understanding of and inference about human behavior. Usually, this involves the specification of one (and sometimes several) cognitive processes or behaviors of interest and the subsequent design/implementation of a task that can capture this process in as unequivocal a fashion as possible. This difficulty is evident in the fact that decision making and the processes that subserve decision making, however differently they may be invoked as a function of context, is generally amorphously defined. This variable definition serves a powerful purpose: it reminds us of the inherent complication of *measuring* the very thing we are interested in, even if we narrow our approach to use specific formalizations or to specific processes of interest. Rather, it should remind us.

Any such model that does meet Mari’s requirements, therefore, would be overspecified – prioritizing bias mitigation over the potential for generalization. We must therefore ensure to strike the appropriate balance as we develop tethered-yet-abstract models to explain complex human behavior. These questions are neither new nor explored in an especially novel way in this thesis. We repeat them here because these concerns merit revisiting on a fundamental scientific level.

This brings us to the matter of how to determine the *important* sources of variability and what it means for a dimension of variability to even be important in the context of coherence and consistency (Tal, 2020). For example, consider an individual walking through the grocery store in search of bread. She stands in front of the section where several loaves of carefully packaged bread lie languidly on shelves¹. From the decision maker’s perspective, she has only to figure out which loaf of bread she “wants.” However, when she puts on her decision scientist hat, she confronts a different set of questions: why is she choosing the way she does?

¹The author may have been hungry at the time of writing.

This is an example of a choice problem with several options and several different factors that the individual could consider as they make their decision. For example:

1. What did she walk past as they arrived at the bread section? (e.g. Did she walk past a special sale on jams that pair well with rye?)
2. How much does placement play a role? (e.g. Is her initial desired loaf placed on the top row? What other loaves of bread flank it?)
3. What about condition? (e.g. Does the bread look fresh? Did the individual take a long route and feel fatigued after arriving at her destination and therefore unable to explicitly evaluate the bread?)
4. How much does it matter that someone left biscuits next to the bread?
5. What about how much time an individual has to make a decision? (e.g. What if she has to decide quickly? What if she can take as long as they want?)

In this example, it is possible that in any one occasion each, several, or all questions will be of importance to the individual any time she makes this decision. However, the individual may factor into their decision some questions more systematically and regularly than others. Similarly, while it is near impossible for a researcher of decision making to be exhaustive in all sources of variability that an individual could possibly consider while making a decision, it is our impetus to think explicitly about the important ones.

Importance can therefore be motivated theoretically and empirically. In this thesis, Chapters 2 and 3 introduce and explore a source of variability we believe to be theoretically important in how we measure (and conceptualize the measurement of) human decision making under certainty (e.g. bread problem question 1). In Chapter 4, we develop and apply a modeling paradigm that models choice and response time in order to have more psychologically precise and nuanced interpretations of how good people are at memory discrimination (e.g.

bread problem question 5). In Chapter 5, we present preliminary results applying methods developed in Chapters 3 and 4 to a special population: older adults (e.g. bread problem question 3).

Chapters 2 and 3 are currently under review and have been available as preprints since February and March, 2023: <https://psyarxiv.com/zdq5v/> and <https://psyarxiv.com/a8gz3/>

Portions of Chapter 4 can be found in our 2022 International Conference on Cognitive Modeling Proceedings: Banavar, N.V., & Bornstein, A.M. Response time modeling provides stable and mechanistically interpretable measures of individual differences in behavioral pattern separation.

Portions of Chapter 5 can be found in our preprint: Banavar, N.V., & Bornstein, A.M. (2023, May 6). Multi-plasticities: Distinguishing context-specific habits from complex perseverations. <https://doi.org/10.31234/osf.io/t7vsc>

Chapter 2

Variability in Complex Constructs: Inferring Risk Preference and Temporal Discounting

We examine the prevalence and extent of variability across measurements of supposedly stable behavioral economic traits. We begin by reviewing how these traits are conceptualized in behavioral economics, and how different instruments for eliciting them lead to variability in their measurements. We then consider factors such as experiment structure, affect, and context, known to influence or correlate in some way with the inferred values of these constructs: from domain or “subject-level” influences to local influences. We introduce the idea that an important – and cognitively meaningful – source of potential variation in experimentally-inferred measures may come from temporal sequence or the influence of trial order. Finally, we discuss how some of these sources of variation may not be ultimately all be brought under experimental or analytical control, and propose that they should instead be exposed and considered for their predictive value in different settings.

2.1 Introduction

In this paper, we provide a synoptic perspective on two widely studied psychological constructs: risk preference (how individuals behave under known and immediate uncertainty) and temporal discounting (how individuals behave under future uncertainty). Increasingly, important decisions about an individual are being informed by an assessment of how that individual makes decisions under uncertainty. These can be high-stakes interventions (e.g. personalized vaccine incentives - (Andreoni et al., 2016)), but the measures they depend upon are known to be variable across contexts (Peters & Büchel, 2011) and may have limited predictive power for real-world consequences. We discuss how the different methods used to measure these constructs are subject to distinct forms of variability, each of which themselves can be valuable in different settings. Some of these measures may be more robust to time, and others may be more robust to framing effects. Those of us who study and apply risk and temporal discounting measures would therefore benefit not only from the use of appropriate methods but also from an understanding of their endemic variability.

Behavioral Economic Definitions

To understand the nature of variability in the constructions of interest, we must first define the theoretical framework under which they are measured, neoclassical behavioral economics.

Rationality. A fundamental assumption in many behavioral economic models of choice behavior under uncertainty is that ‘rational’ individuals act in order to maximize their expected utility (Von Neumann & Morgenstern, 1944). That is, assuming that they meet the axioms demarcating rationality, a utility function can be well-defined (Von Neumann & Morgenstern, 1944). Individuals can then be characterized as acting in order to maximize these utility – or satisfaction – functions. This is known as Expected Utility Theory (EUT). Other classical

theories of utility such as Random Utility Theory which assumes stochasticity (Thurstone, 1994) and Subjective Expected Utility Theory (SEUT) which extends EUT as implied by its name (Savage, 1972). Critically, economic rationality necessitates consistency: some inherent stability or structure in preference and choice behavior. Much empirical evidence and theoretical skepticism, however, has cast doubt on the value of the rational (ideal) individual, revealed preference, and how closely the behavior of said individual matches actual human behavior –famously violated in the Allais Paradox (Allais, 1990), Ellsberg Paradox (Ellsberg, 1961) and in framing effects (Kahneman & Tversky, 2013). Despite all this, it is for good reason that (S)EUT has continued to dominate the behavioral economic study of decision making under uncertainty: it is a compelling normative framework that has been used across a wide range of domains including choice behavior, public policy, and medicine (Weber, 2010). We therefore consider more carefully the formalization of decision-making under uncertainty within this framework.

Risk. Immediate uncertainty can be in the form of risk – where probability distributions are fully known – or ambiguity – where the probabilities are fully or partially unknown. For example, an individual making a choice between a 25% chance of \$20 (and 75% chance of \$0) or \$5 guaranteed, is making a choice under risk. Alternatively, if the subject is choosing between the guaranteed \$5 and a 25% chance of \$20 without knowing what the remaining 75% would give them, they are deciding under partial ambiguity – they have incomplete information about the underlying probability distribution. While conceptually intimately related, people behave differently under risky compared to ambiguous circumstances, with evidence for even more nuanced differences in clinical populations (Konova et al., 2020). As quotidian decision making rarely involves complete probabilistic information, decision making under ambiguity is of considerable ecological interest. Due to the limited scope of this paper, however, we will primarily consider decision making under risk.

Expected Utility, then, is usually characterized as a power function: Expected Utility = $p \cdot v^\alpha$,

where p represents the probability of a given reward and v represents the objective reward (e.g. dollar amount of reward). The curvature α of the expected utility function, known as risk tolerance, has very specific interpretations: if an individual is risk averse, they prefer less uncertain outcomes even if the dollar reward is lower. This is characterized by a concave utility function, with $\alpha < 1$ (individuals whose utility function is convex with $\alpha > 1$ are classified as risk seeking and $\alpha = 1$ as risk neutral). Famously, Kahneman and Tversky showed that these characterizations depend on a reference point (the domain in which decisions are made) – they held if people were making choices between rewards they could possibly win, but flipped if people were making choices between rewards they could possibly lose (i.e. risk aversion in the loss domain yields a convex utility function) (Cumulative Prospect Theory (CPT); (Kahneman & Tversky, 2013)). As there are many other ways to parametrize an individual’s decision making under risk (see *Delineating Risk*), we use risk tolerance (or α) to reference the parameter predicated on EUT/CPT (and variants) and risk preference to refer to the more general latent variable/concept. We further note the related phrase risk perception, which is often confounded with risk tolerance and preference, is generally defined as the recognition of inherent risk – the ability, given internal and external circumstances, to appropriately assess the riskiness of a situation (D. R. Hunter et al., 2002). This is contrasted with risk tolerance or preference, as they deal more specifically with an individual’s willingness to engage in risky decision-making, usually when available options are of equal expected value.

Temporal Discounting. Models of temporal discounting study the interaction between time, value, and uncertainty and aim to capture how individuals attribute differential weight to choice options closer to the present compared to the (distant) future. A rational, and therefore consistent, individual would use an exponential discounting function (Samuelson, 1937): $Utility = \nu \cdot e^{-kd}$, where ν again represents the objective amount of reward; d the delay, or how far away from the present the reward would be received and k the discount factor. This is also known as a linear exponential model as a linear utility function (with $\alpha = 1$,

interpreted as discussed in the previous paragraph) is assumed. Exponential discounting indicates that the rate at which an individual discounts is constant over time. However, decades of empirical work has demonstrated that this does not explain human behavior well, and that the linear hyperbolic discounting model is a significantly better fit: $Utility = \frac{\nu}{1+kd}$, where all variables maintain their interpretations (Thaler, 1981; Kable & Glimcher, 2007). This model also assumes that individuals are risk neutral, but critically allows time-varying discount rates: steeper discounting when the delayed option is closer in time. The discount factor can range from 0 to 1 – where a value of 0 would mean that an individual considers only the dollar amounts offered regardless of temporal distance. Interestingly, animals tend to be even more myopic in their discounting (Loewenstein et al., 2015).

Just as with risk, there are many ways to parametrize how an individual’s discounting decays (see *Delineating Temporal Discounting*). We use the term discount factor (or k) to reference parameters predicated on EUT (and variants), and the more general temporal discounting to reference the concept. Finally, we note that an individual’s predilection towards selecting the option closer in time at the cost of an objectively better (as per economic rationality) reward is also commonly referred to as their impulsivity – individuals with high k s are usually interpreted to be more impulsive. Like risk tolerance, impulsivity has also been used to mark differences in clinical populations (e.g. Addiction: (Bickel et al., 2014); Major Depressive Disorder: (Pulcu et al., 2014)). Finally, we note that impulsivity is a multidimensional construct where the dimensions themselves are yet to be agreed upon by researchers (Evenden, 1999), and therefore should not strictly be used interchangeably with temporal discounting.

Relating the two. While risk tolerance and discount factor are distinct in both their psychological and economic formalizations, they are intimately related. Indeed, they must be by definition at least within the Expected Utility frameworks, as both parameters explicitly invoke Subjective Value. The standard (exponential/hyperbolic) discounting model implicitly assumes that a given individual is risk neutral, and this imposes a systematic

bias on measures of discount factor (Lopez-Guzman et al., 2018). That such a bias exists corresponds with a well-established empirical finding that people are generally risk averse in the gain domain and risk seeking in loss (Tversky & Kahneman, 1989); they typically aren't risk neutral. They are also conceptually related as they both involve notions of uncertainty (immediate vs. temporal/delayed). Other theories that explicitly link risk preference and temporal discounting include Construal-Level Theory (Liberman & Trope, 2003) which posits that an individual's change in (subjective) value is due to changing mental representations and a differential focus on concrete vs abstract features, informing the manner in which an individual both discounts future reward and behaves under risk (Leiser et al., 2008). Dual-process models of deliberation and affect suggest that risky decision making and intertemporal choice (ITC) may also be linked in how people trade off the desirability of presented options with the cost of willpower or effort required (Loewenstein et al., 2015).

Empirical studies in animals have shown that animals are sensitive to the frequency with which they must make risky choices: one study found that rhesus monkeys that choose between risky gambles every 3s demonstrate risk seeking behavior, while others in rats and birds found evidence for risk aversion when choices were made approximately every 30s (Hayden & Platt, 2007). More concretely, Hayden and Platt found that rhesus monkeys preferred certain options over risky ones with increasing delay (larger inter-trial intervals) between choices (Hayden & Platt, 2007). That an ostensibly simple and unrelated manipulation in how animals, assuming some degree of comparative equivalency, choose between gambles in an experiment leads to behavior with distinct interpretations suggests a relationship between time, uncertainty, and choice that is evolutionarily old. This is consistent with a 2013 meta-analysis on neuroscience data sets searching for a neural value system – the two primary regions identified being the (evolutionarily newer) ventro-medial prefrontal cortex and the (older) ventral striatum (Bartra et al., 2013). There are also other evolutionarily older ways in which neural information processing and representation could relate risky decision making, intertemporal choice, and cognition more broadly (see *Trial-Trial Temporal*

Dependence). We note that despite the fact that these constructs are related across many dimensions, they are usually studied separately (i.e. risky decision-making experiments vs. intertemporal choice experiments).

Rationality Revisited. Behavioral economic parameters of interest, useful in both explanatory and predictive capacities, are multi-dimensional constructs that also manifest differently at different time scales (see *Delineating Risk, Delineating Temporal Discounting*). There have accordingly been decades of debate over whether it is appropriate to conceptualize these constructs as trait-like variables, with strong evidence of both stability and variability in choice behavior and subsequent inferred parameter values. Choice behavior, through which these constructs are regularly quantified, however, is subject to a wide range of influences (see *Delineating Risk, Delineating Temporal Discounting, Variability in Experiments*). We might consider an individual in a financial difficult circumstance, or more mundanely, trying to purchase an out of budget treat for themselves: they may be biased in an intertemporal choice experiment towards more immediate rewards, and therefore be characterized as very impulsive. In that their current desire for the immediate reward is identified, measurements and standard interpretations of temporal discounting are appropriate. However, it is an entirely different, but related question as to whether this measurement is reflective of their ‘innate’ impulsivity – if such a construct is even meaningful (reasonable test-retest reliability, at the very least, suggests that could be, e.g. (Kirby, 2009; Frey et al., 2017)). However, it is difficult to make a decisive statement on this, after all, individuals are asked explicitly to act in accordance with their true preference. In fact, if individuals were truly compliant with task instructions, there would be little dissociation between (personal) circumstantial and task-congruent influences. Similar logic holds true for individuals participating in a risky decision-making task – there are a multitude of reasons why an individual behaves the way they do in any given moment. The relationship between goals, especially clearly articulated goals, context, behavior, and parameters of interest is one that requires careful thought. This is especially the case when more “extreme” behaviors are usually

interpreted as deviations from rationality and its subsequent throng of implications.

The wide range of observed behaviors that deviate from economic rationality, then, are usually either characterized as “not as irrational as they may seem,” or as necessitating a new conceptual, veridical framework (Vlaev, 2018). Recently, Vlaev formalized a theoretical compromise, leveraging the definition of rationality most prominent in cognitive science (Anderson, 1990): it is more appropriate to consider human behavior as locally rational – that humans make (environmentally) contextualized rational inferences as opposed to universal (Vlaev, 2018). Given various constraints in human decision making (e.g. limits on information processing) and the sheer depth of computation theoretically necessary for globally consistent choice (Simon, 1990), this strikes as a more plausible casting of the human implementation of [economic] rationality. This framework therefore implies that value itself is not well defined or consistent. That is, as others (e.g. (Slovic, 2020)) have proposed, preference is constructed and not just revealed during elicitation (though note that this does not mean that preference is necessarily constructed from scratch, or independently of previous experiences, each time). Thus, Vlaev synthesizes that, using (limited) resources and privileged information, value, local comparisons, and subsequent inferences are all computed online and in a rational, sequential manner. Despite suggesting a lack of consistency in value and its predicated constructs like risk tolerance and discount factor, we do not argue that the notion of stability in these constructs should be entirely divorced from their conception. It is, after all, impossible to formulate constructs that are sufficient in depth and breadth.

The remainder of the paper is outlined as follows. First, we discuss variability in eliciting and measuring both risk preference and temporal discounting. Then we consider different ways in which we can analyze experimental data and relate these measures to preference. Finally, we consider the ideas of task-incongruent temporal dependencies (i.e. perseveration and serial dependence).

2.2 Delineating Risk

We first consider the different ways in which risk preference can be measured behaviorally.

Measurements. Experiments designed to estimate an individual’s risk preference generally fall into one of three categories: statistically dependent sequential choices (SDSC, e.g. Balloon Analogue Risk Task, n-armed bandit tasks); statistically independent ordered choices (SIOC, e.g. Holt and Laury gambles) and statistically independent single choices (SISC, e.g. lottery tasks) (Pedroni et al., 2017). Other features that can vary across tasks include how choices are displayed (numerically – with monetary values and probabilities listed, graphically – with graphical or pictorial depictions of probability, or both); choice domain (gain, loss, or mixed); incentivization (e.g. Becker-deGroot-Marschak random draw, cumulative reward); the presence or absence of feedback (that is, the immediate realization of their choice resulting in feedback informing them of their win/loss) and the amount of time an individual has to respond (Pedroni et al., 2017). While SDSC tasks often explicitly model learning and other possible temporally evolving processes and dependencies, SISC tasks focus on “in the moment” decision making and typically consist of randomized, and therefore temporally unstructured choice sets.

A further, related, distinction across these experiments is the “description – experience” gap: that people behave differently when complete information is provided about the problem (and by extension the environment) versus when they are provided incomplete information, and need to rely on experience (previous or current) (Hertwig & Erev, 2009). If we consider description – experience, risk / ambiguity and the presence or absence of feedback, we can further taxonomize these experiments. SDSC tasks are typically experiential, while SISC tend to be descriptive, with feedback acting as an important arbitrator between SDSC and SISC and between whether the individual is making decisions under ambiguity or risk. In SDSC tasks, incomplete information about the probabilistic structure of the environment (or

bandit machines, for example) is reducible – people can actively learn about and mitigate the underlying uncertainty through feedback (usually in the case of rewards won or lost after a choice). In SISC tasks, however, the reducibility of ambiguity is entirely dependent on the construction of the choice set and the presence of feedback. In fixed ‘unstructured’ (i.e. choice set does not change over the task like with staircasing) experiments without feedback, the underlying uncertainty is irreducible. The individual makes choices in the dark and with, in theory, only their preference and description of the problem to guide them. These are descriptive decisions under both risk and ambiguity, as information is explicitly presented to the individuals, with nothing to be “learnt” as lotteries presented are usually fixed (i.e. at least 25%, 50%, 75% chances of winning).

Similarly structured experiments that involve feedback, however, can allow for individuals to learn about what the underlying probability of the various gamble types presented, much like in bandit tasks, except that subsequent choices are unrelated to each other. Feedback therefore allows for individuals to “experience” the consequences of their decisions and, depending on the goal of the experiment, there is variability in how feedback is expected to influence trial and aggregate choice behavior (Barron & Erev, 2003; Brooks & Sokol-Hessner, 2020). Usually, however, the standard modeling framework of SISC experiments, especially in the context of inferring these parameters of interest, is to treat data as explicitly descriptive and not account for potential transient within-task influences or learning, however task-irrelevant they might be.

Many risk preference elicitation methods exist in the literature. Beyond behavioral experiments, there exist more subjective measures, usually in the form of Likert scales (e.g. “how risk seeking are you in general?”) or surveys, such as the Domain Specific Risk Taking (DOSPERT) Scale (Weber et al., 2002) (see (Frey et al., 2017) for a more comprehensive list). Finally, measures of an individual’s risk preference can also come from frequency measures by tabulating the occasions on which an individual engages in risky behavior, though

we note that this information is usually also collected through self-report (Frey et al., 2017) unless, for example, in a clinical setting.

Variability. That risk preference is complex is intuitive, if not patent in the many means by which it can be defined and assessed. Changes in risk preference have been observed as a function of affect/motivational state (fear increases risk aversion while anger decreases it (Kugler et al., 2012)); age (older adults are more risk averse (Tymula et al., 2012)); clinical disorders (patients with substance use disorder are more risk tolerant (Konova et al., 2020)); and sex (women are more risk averse (Croson & Gneezy, 2009)). We note that for most studies that find structured evidence of the malleability of risk preference, there are studies that find no evidence of any such sensitivity (e.g. no systematic effect of stress (Sokol-Hessner et al., 2016)).

Risk preference is also thought to be related to important variables such as income, intelligence, and education – though a recent study in a large ($N = 916$) diverse cohort of adults found that only sex and age have robust, consistent associations with risk preference (Frey et al., 2021). Importantly, this study found that the relationship between risk preference and these correlates varied as a function of how risk preference was measured, with subjective measures being more sensitive to these correlates relative to experimental measures. Earlier work by the same group sought to examine whether these ostensibly different measures of risk could be consolidated into a single latent variable R , much like g (intelligence) (Frey et al., 2017). Using 37 different risk elicitation measures in a sample of 1507 from two different countries, the authors found that they were indeed able to extract a temporally stable R that accounted for 50% of observed variation in a factor analysis. Critically, however, almost all this stability was attributed to measures elicited from surveys (e.g. DOSPERT) and frequency counts of risky behavior, also measured through surveys. Further, subjective and frequency measures had much higher temporal stability and correlations between and within themselves relative to measures elicited experimentally (ranging from SDSC to SISC tasks).

Nonetheless, recent work has demonstrated the value of temporal fluctuations, finding in a clinical setting that only week-to-week fluctuations in experimentally elicited measures of ambiguity tolerance and recent risky behavior (e.g. recent drug use) were predictive of future real-world behavior under uncertainty (Konova et al., 2020). More generally, other studies have also found relatively low correlations between experimentally induced measures of risk tolerance, including differential contextual or emotional sensitivities (Guan, 2019; Kugler et al., 2012; Pedroni et al., 2017; Radulescu et al., 2020; Sokol-Hessner et al., 2016).

As risk tolerance is well established, and by definition subjective and relative (Weber, 2004), and experiments themselves can vary widely in their construction, it is perhaps unsurprising to find such high levels variability in behavioral experiments – as self-reports might manifest more like personality traits than as functions of socioeconomic status or cognitive ability (Frey et al., 2017). Self-reports also assess risk preference at a different (global) timescale, and therefore elicit qualitatively different information. Further, as the questionnaires often ask individuals to respond hypothetically and “in general,” it would be more appropriate to characterize these measures as decision making under ambiguity, not risk. It might be even more appropriate to consider these measures as meta cognitive: that they reflect an individual’s thinking about how they think about the question vs their thinking during actual choice. Local rationality would suggest that, due to context-dependent differential information sampling, it is the former (Vlaev, 2018). Thus, we have focused in this paper on experimentally elicited measures, as there at least individuals largely make incentive-compatible (“real”) as opposed to hypothetical choices and probabilities are explicit, and therefore truly in the domain of risky decision making.

The complexity of the matter at hand, however, still does not diminish. Researchers have shown that individuals adopt different strategies depending on the structure of the experiment (Pedroni et al., 2017). More damningly, even after differences in the structure of the experiment were controlled for, Pedroni et al were unable to elicit a stable measurement

of risk preference, suggesting that an individual’s experimentally induced risk preference is likely constructed in the moment, multi-dimensional, and the product of multiple cognitive processes interacting. Overall, these variations appear to be largely a function of context, experimental structure, and the interaction of variable processes: very much in line with reasons to consider frameworks that explicitly account for contextual variability such as Vlaev’s local rationality as both plausible and appropriate.

2.3 Delineating Temporal Discounting

We note that less quantitative research has been conducted on the variability of measures relating to an individual’s temporal discounting, relative to risk.

Measurements. We can leverage the same overarching taxonomy to categorize intertemporal choice tasks as with risky decision making. In a typical behavioral-economic intertemporal choice task, individuals will choose between a smaller sooner “SS” option or a larger later “LL” option. Thus, experiments can be SDSC (e.g. titration methods when options presented depend on previous choice to arrive at ostensibly more precise estimates as in (Solway et al., 2017)), SIOC (e.g. options presented are independent of choice but have some structure, e.g. increasing LL option by \$5 each trial as in (Steinglass et al., 2017)) or SISC (most common: no built in cross-trial relationship as in (L. Hunter et al., 2018)). Unlike with risky decision making, intertemporal choice experiments about money are usually displayed only numerically – with monetary values and delay listed (e.g. (L. Hunter et al., 2018)). This ceases to strictly be the case when individuals make choices about non-monetary rewards like food or alcohol, where both pictorial representations and the physical objects they are choosing between can be presented (e.g. (Addessi et al., 2014)). Some experiments also vary choice domain (gain or loss, often in conjunction with gain/loss in risk, e.g. (Estle et al., 2006)); incentivization and the amount of time an individual has to respond (Scherbaum

et al., 2012). As there is no immediate uncertainty involved in pure intertemporal choice, feedback via choice realization as studied in risky decision making is largely inconsequential.

Similarly, while some authors have considered the description-experience gap in intertemporal choice, this is typically only examined in the relatively uncommon context of probabilistic rewards (e.g. (Dai et al., 2019)): that is when either immediate and/or delayed rewards are themselves offered probabilistically. This makes sense as there is no ostensible learning or underlying uncertainty to be reduced besides inherent temporal uncertainty which is both outside the decision-maker’s control and unable to be experienced – and thus seemingly resolved – till that moment in time. This is the case unless, for example, the experiment is situated in a virtual world where the experimenter is imperator and can control time.

Just as with risk preference, many temporal discounting elicitation methods exist in the literature: experiments, surveys (e.g. Barratt Impulsiveness Scale, (ES, 1983)), and frequency measures.

Variability. Measurements of temporal discounting or impulsivity show varying degrees of temporal stability and predictive power. They are demonstrated to vary with affect (increases with sadness and reduces with gratitude (K. Lempert & Phelps, 2016)); age (older adults are more patient (Green et al., 1994)); attentional and framing manipulations (increases with focus on delay and decreases with focus on magnitude (Leiser et al., 2008; K. Lempert & Phelps, 2016)); pathology (patients with substance use disorders are more impulsive (MacKillop et al., 2011)); prospection (decreases with emphasis on future concreteness (K. Lempert & Phelps, 2016)) and sex (women discount more steeply than men in the lab (Weafer & de Wit, 2014)). Again, however, there is much extant literature suggesting a lack of systematic relationship (e.g. no conclusive direction one way or the other for sex differences (Cross et al., 2011)).

More concretely, human discounting, like risk preference, is sensitive to domain and circum-

stance – not only do people differentially discount across goods and money, they allocate fixed resources (money) on these goods depending on their current financial situation (Ubfal, 2016). In this study conducted in rural Uganda with non-hypothetical rewards, the less income an individual had, the more money they were willing to spend on items they discount highly. Similarly, individuals with gambling use disorder discounted money more highly when in a gambling context, as opposed to a non-gambling context (Peters & Büchel, 2011). This is putative evidence for the influence of personal goals and contexts (income in the Ubfal example; physical location in the Peters & Büchel example) on choice behavior, something only speculated about earlier in this paper. Further, the test-retest reliability in measures of discount factor could also partially be dependent on reinstating the same context in which initial measurements were made: in a study, 5 week test-retest reliability of k was 0.77 (95% CI: 0.67-0.85, $n = 81$), 1 year was 0.71 (0.5-0.84, $n = 37$), and 57-weeks was 0.63 (0.41-0.77, $n = 46$) when subjects made choices between which delayed reward they preferred (Kirby, 2009). It would be interesting and might lead to better correspondence with real-world behavior to consider a within-subject design where such data were collected across multiple different contexts (e.g. in Ecological Momentary Assessment style experiments which we are sure must currently be in progress).

Unsurprisingly, inference on how an individual discounts value over time has also shown to be sensitive to the structure of the experiment. For example, Lempert and colleagues inferred different discount factors for subjects depending on how they manipulated stimuli in the experiment: people discounted significantly more steeply when there was greater variation in the delayed reward relative to the immediate reward but that the rank ordering of discount factors remained consistent regardless of experiment structure (K. M. Lempert et al., 2015). The researchers also had subjects complete various surveys measuring impulsivity / related factors and, like the Frey group, found significant correlations in temporal discounting measures derived from surveys but not between surveys and experiments. Other researchers have found similar evidence across multiple clinical populations. For example, Ledgerwood

and group found these patterns held in control subjects and pathological gamblers with and without a history of substance use disorder (Ledgerwood et al., 2009). They further found that pathological gamblers were generally more impulsive regardless of substance use history, but that the gamblers with a history of substance use were more risk tolerant. This is just one simple example to demonstrate simultaneously the clinical significance of these economic constructs and how they may (or may not) vary across populations. That the overall relationship between methods of temporal-discounting elicitation seems to hold despite pathology suggests that, unless we assume that this variability is irreducible, there may be other factors – cognitive or otherwise – that may not be considered carefully enough by the field.

Overall, we see remarkable correspondence in research studying the elicitation and sensitivity of both risk preference and temporal discounting. Regardless of the construct, there is much variability in the conceptualization and inference of parameter values. We further see that this variability tends to be greater in experimentally induced parameter inference. We next consider other observations of variability in choice behavior in experiments.

2.4 Variability in Experiments: Measurements

The idea that choice behavior and decision-making strategies in behavioral-economic experiments may change over the course of an experiment is not novel and has been explored for decades (Slovic, 2020; Vlaev, 2018). When characterizing variability in behavior, we can, in general, consider macro (e.g. domain differences) and micro levels. Framing effects, like those described by Prospect Theory where people use a single reference point to guide behavior in the gain domain compared to loss, are examples of macro influences in that they demarcate domains (Kahneman & Tversky, 2013). Such reference points are traditionally assumed to be fixed – these frames may impact how an individual behaves on aggregate

and on a given trial (but see (Koop & Johnson, 2012) for empirical observations suggesting multiple reference points). Extant literature has shown that context can also exert a macro level influence on behavior and inferred parameters (Peters & Büchel, 2011). Context – depending on how it is defined – however, is particularly precarious and can also influence decision-making at micro-levels (K. M. Lempert et al., 2015) and in non-human primates (Zimmermann et al., 2018).

We can, then, decompose micro effects into trial, within-trial, and between-trial levels. Trial level measurements include the gold standard but highly variable choice behavior, and reaction time (though note that models of reaction time themselves are within trial as they seek to describe dynamics over the course of the trial itself). Choice behavior is generally modeled in accordance with SEUT or CPT maximization as described in the first section of this paper in conjunction with a choice rule.

Within-trial measurements generally include response time models and process-tracing methods such as mouse and eye tracking (see (Schulte-Mecklenbeck et al., 2017) for more). Response times are widely modeled using a sequential sampling framework which assumes that we accumulate information in favor of/against the options presented to us in a noisy manner until we have accrued enough to make a choice (or never accrue enough to ever make a choice) (Forstmann et al., 2016). One of the most widely used frameworks, the Drift Diffusion Model, breaks down the accumulation process into four parameters in two alternative forced choice tasks: bias (predisposition towards Option A or B), drift rate (the rate at which evidence is accumulated), threshold (the amount of information needed to make a choice), and non-decision time (generally considered to be irrelevant to the decision process) (Ratcliff, 1978). These psychologically interpretable parameters that model components of deliberation have shown to correlate with discount factor ((L. Hunter et al., 2018; Konovalov & Krajbich, 2019) Konovalov & Krajbich, 2019), risk tolerance (Konovalov & Krajbich, 2019), and preference more broadly (Konovalov & Krajbich, 2019).

Process-tracing methods, then, are direct measurements of the dynamics in decision making, capturing the online formation or reversal of preference (Koop & Johnson, 2013) and the mitigation of conflict as individuals choose between two options (Stillman et al., 2020). Scientists can measure and model these dynamics in decision making by examining the path subjects take via their computer mouse: a direct and swift movement from trial start to the option selected suggests decisive choice, whereas more winding trajectories could indicate decision difficulty (conflict) and even preference reversal. Such measurements further allow arbitration between different theories of preference formation. We omit discussion of eye tracking and value-based decision making due to space constraints (see (Orquin & Loose, 2013) for a review) and instead focus on mouse tracking.

For example, Koop and Johnson demonstrated that preference reversals during risky choice inferred via mouse tracking were inconsistent with heuristic decision strategies like “take the best” (TTB) which posits that individuals focus on a particular dimension and select the choice that ranks highest on that dimension (Koop & Johnson, 2013). TTB’s incompatibility with preference reversals were demonstrated by the degree with which mouse trajectories deviated from relative linearity (i.e. moving the mouse directly to the object of choosing). Similarly, Stillman and group found in a risky decision-making task that the more similar the subjective values of choice options, the less direct and more conflicted the subjects’ trajectories were despite controlling for response time greater the conflict (Stillman et al., 2020). The authors argued further that mouse trajectories could correlate with an individual’s risk tolerance: an individual who follows a direct trajectory to the gamble as opposed to the certain option is likely more risk tolerant than someone who takes a meandering path. Stunningly, the authors found that decision conflict on single trials correlated strongly with risk tolerance, inferred in accordance with the Prospect Theory framework, and predicted behavior on the subsequent decision (Stillman et al., 2020). The authors argue that mouse-tracking dependent inference outperforms traditional behavioral measures of choice behavior and reaction time analyses as mouse-tracking might be more robust to other factors known

to affect response time and choice behavior (e.g. non-decision time). Similarly, scientists have correlated decision strategy as inferred through mouse tracking dynamics with discount factor in intertemporal choice (Reeck et al., 2017). Such analyses suggest a promising avenue to understand more about locally rational, online decision making, and especially the role of similarity between options presented on a given trial.

2.5 Trial-Trial Temporal Dependence

Any discussion on human behavior and rationality would be incomplete without a brief further comment on capacity constraints and adaptive behavior. A key signifier of ‘intelligence’ is the ability to navigate complicated environments. Animal – and artificial – behavior is however hardware constrained: there are limits to the ability and flexibility that organisms and algorithms can demonstrate. Many theories of how the human brain evolved to be able to maneuver such a complicated world given limited resources revolve around the idea of adapting to or leveraging (stationary) statistical information in the environment (Anderson, 1990). In perceptual neuroscience, this is referred to as the Efficient Coding Hypothesis where limited probabilistic neuronal representations maximize information and minimize redundancy (Barlow et al., 1961) in a context-sensitive way (Schwartz et al., 2007). Such adaptive sequential sensitivity has been demonstrated in lower-level cognition (Simoncelli & Olshausen, 2001) and, more recently, in non-human primate economic decision making (Zimmermann et al., 2018). The results from the Zimmermann paper are in particular valuable because the authors demonstrate the first evidence of [the necessity of] trial-trial temporal dependencies in canonical neuronal computations during economic choice. That is, not only does behavior in economic choice change as a function of variability in rewards (a prediction of the Efficient Coding Hypothesis), but models of neuronal computations that are consistent with efficient coding – typically specified at the intra-trial level – can only describe behavior

well if the temporal order and (local) contexts of the experiment are preserved and explicitly accounted for. Taken together with theoretical and empirical neuroscientific research on sequential sampling in the brain (e.g. (Gold & Shadlen, 2007)), this suggests that in the realm of rationality – resource or economic – context is king, and hence lends further credence to frameworks like Vlaev’s local rationality. Indeed, sequential sampling models like the Drift Diffusion Model have been monumentally successful in describing behavior alone (Forstmann et al., 2016).

There is, therefore, a strong intuition as to the normative reliance on recent history during experience in the moment – be it simply perceiving stimuli or during the decision process and subsequent choice itself. The mechanisms through which this might manifest are still fundamental open questions in the field, though there is general speculation on the (often complementary) roles of attention and working memory in propagating this temporal continuity (Kiyonaga et al., 2017). While this is usually considered to be adaptive (Kiyonaga et al., 2017), we highlight two cases in which reliance on recent history can prove to be problematic or task incongruent: environments without sequential dependencies and clinical pathology.

Much research has examined how reliance on the past can cause problems in lower-order cognition due to task-irrelevancy. Some of the earliest evidence of this comes from the absolute identification literature in the 1950s and onwards: where individuals were demonstrated to treat independently generated stimuli (i.e. presented a sequence of stimuli that were not related by time, like in SISC tasks) as if they were actually related (e.g. (Verplanck & Blough, 1958; Lockhead & King, 1983; Stewart et al., 2005)). For example, when people were asked to make judgements about line lengths or tone frequency, experimenters found robust evidence of transient framing effects: the lines or tones they had seen immediately (1 – 4 lags) before influenced their judgements on the current stimulus shown (Stewart et al., 2005). Interestingly, some experiments have shown different effects as a function of lag: more recent stimuli

tend to produce an attractor-style effect, while more distantly observed stimuli produce a contrast effect (Stewart et al., 2005). Researchers in visual perception have termed this effect, ostensibly distinct from priming, hysteresis, statistical artefacts, and learning, as serial dependence (Fischer & Whitney, 2014). This effect, consistent with the Efficient Coding Hypothesis, is also thought to be adaptive despite any inferential obstruction it may cause in such randomized tasks. We note that serial dependence has important consequences in real-world contexts too, and not just as a potential ‘contaminator’ of psychological inference. Recently, work from David Fischer’s group showed that radiologists demonstrated serial dependence while making medical judgements about simulated patient lesions (Manassi et al., 2021). More broadly, Fischer and Whitney suggest that serial dependence is characterized along three dimensions: similarity (only present when stimuli have similar features), temporality (decays over time), and spatiality (strongest when stimuli presented in the same location) (Fischer & Whitney, 2014). The authors also identify attention as a fundamental player (Fischer & Whitney, 2014). It is still, however, an open question as to whether this type of between-trial effect extends to higher-order (behavioral-economic) decision making.

In SISC (randomized) tasks, the standard experimental structure in intertemporal choice and risky decision making, stimuli are expectation controlled and thus presented in a randomized fashion. That is, successive stimuli will possess varying degrees of similarity to each other. For example, on trial $t - 1$ an individual chooses between \$5 today and \$45 in 80 days and on trial t chooses between \$4 today and \$48 in 70 days. Here the immediate and delayed rewards are similar in value, as is the delay of the rewards. Further, as choice options are often displayed in similar spatial locations (though we note that there is usually randomization at least in terms of the side of the screen – left or right – each option is presented), and decisions are made in a sequence, the criteria for plausible serial dependence according to Fischer and Whitney appear to be met.

It is entirely conceivable and ostensibly efficient for computations made during trial $t -$

1 to be (partially) cached and reused or referenced on trial t as a function of similarity (Dasgupta et al., 2018), amongst other things, thus affecting response times and possibly choice behavior. Such influences, however transient they may be, may provide us information as to the cognitive health of an individual (see below) and may aggregate to the point of affecting inference on our parameters of interest, especially if they are not accounted for in statistical analyses. Indeed, some of the concerns raised in the *Delineating* sections earlier with regards to noisiness in experimentally-induced parameter inference, could be due in part to such spillover. These spillover effects may also be consistent with cognitive theories of intertemporal and risky choice. As intertemporal choice involves uncertainty about the future, Peters and Büchel (amongst others) suggest that people’s choices are guided by the deliberative process of prospection – they imagine what their future may look like some days out and use the outcome of that simulation to guide their choice (Peters & Büchel, 2011). Recent work has also shown a relationship between how model-based an individual is and the way they discount the future: people who spend more time considering future rewards in temporal-discounting tasks are also more likely to plan ahead in sequential reinforcement learning tasks (L. Hunter et al., 2018). Further, scientists have hypothesized that the manner in which people choose also changes as a function of delay: Construal Level Theory posits that representations of the future are more abstract (e.g. lower statistical precision) than representations of the present (Leiser et al., 2008) and that people tend to consider more “primary” attributes (e.g. healthiness, “should” behaviors) when thinking of the future and more “secondary” attributes (e.g. tastiness, “want” behaviors) for the present (Rogers & Bazerman, 2008).

We can therefore infer that thinking carefully about anything – in this case the future – can be resource intensive. Thus, in the example above, the individual has already imagined what their life might look like 80 days into the future on the previous trial. Barring some specific event that they expect to meaningfully shape their experience within the 10-day difference, it is likely that their future 70 days out will be similar and thus they could avoid

computational redundancy by reusing (part of) the simulation generated on the previous trial. Indeed, if representations of the future are in actuality more uncertain and requiring the recruitment of higher-order cognitive processes, there is even more reason to support the reuse of previous computations to guide current inference and choice to minimize computationally expensive operations. While decision making under risk may not involve the simulating the future, individuals still need to resolve the immediate uncertainty and complex choice options presented in the form of probabilistic gambles to guide their choice. Thus, computations incurred over the course of the experiment may likewise be carried over from trial to trial, also possibly as a function of (dis)similarity.

Our recent work introduces a novel statistical framework that suggests choice behavior, response times, and risk tolerance/discount factors themselves are indeed influenced by recent history as defined by previous stimuli and choices made (N. V. Banavar & Bornstein, 2023). We term this dependence computational perseveration to distinguish its higher-order nature (involving complex mental calculations) from serial dependence. We find specific effects of computational perseveration in choice behavior, while reaction time parameters showed more widespread sensitivity. However, our results suggest further complexity in the nature of this higher-order serial dependence as we also found evidence for a contrast-like effect: in the risky decision-making task, choice behavior was influenced by previous stimuli when the previous choice was easy, and the current was difficult. This is the opposite of what would be expected given Fischer & Whitney’s criteria (Fischer & Whitney, 2014). Critically, our analyses have shown that the majority of subjects in an Intertemporal Choice task and in a Risk/Ambiguity gambling task show evidence of computational perseveration. Further, we demonstrate that for several subjects, sequential-effect-adjusted parameters for risk and ambiguity tolerance change sign, and therefore, psychological interpretation. For example, someone who was previously identified as ‘ambiguity seeking’ based on their non-sequential-effect adjusted ambiguity tolerance parameter would now be identified as ‘ambiguity averse.’ We argue therefore that while computational perseveration may not be the sole source of

variability in risk tolerance/discount factor inference in experiments, there is theoretical and empirical impetus for us to consider explicitly the influence of temporal context in how we define, measure, and infer these constructs. We finally note that computational perseveration is likely present in SDSC and SIOC tasks too, but due to the potentially confounding nature of structured experiments and learning, we omit further consideration of this topic in this paper.

We believe that this work has deep theoretical and empirical implications. Our analyses suggest that these sequential effects are not noisy artefacts but are instead the consequence of a systematic influence of trial properties on components of the decision process. This suggests a potential need for the theoretical reconceptualization of *experimentally-inferred* parameters as *explicitly dynamic* and sensitive to (highly) local contexts and not *exclusively* a static and psychologically interpretable end (N. V. Banavar & Bornstein, 2023). Our method also allows scientists to analyze a novel dimension of information about the decision maker (i.e. degree of trial-trial dependencies) without having to collect any new measures, as both choice behavior and response times are standardly recorded in experiments. This additional information could have use beyond the purely methodological – it could result in meaningful cognitive and clinical implications.

To underscore the idea that short-term temporal dependencies provide cognitively meaningful information and other research directions beyond parameter calibration (as it may be tempting to infer from the previous paragraph), we consider in brief a complementary, yet distinct, line of work in clinical psychology and neurology. Decades of evidence in these fields has shown a differential reliance on recent history as a function of aging and neurodegenerative pathology (Sandson & Albert, 1984; Goldberg, 1986; Van Patten et al., 2015). Here the abnormal, often over- and task-incongruent reliance on the past, relative to healthy individuals, is termed perseveration. In particular, there exists a three-dimensional hierarchy of perseveration with primary dimensions of content, disorder, and temporal profile. Content

references the material itself that is repeated (ranging from lower-order motor to higher-order semantic/verbal repetitions); disorder references the various ways in which outcome measures might differentially relate to the clinical progression of neural degeneration (e.g. frontal lobe vs basal ganglia damage) and temporal profile, which delineates the varying timeframes along which perseveration can manifest (e.g. perseverate information from seconds ago, minutes ago, or even tasks ago) (Sandson & Albert, 1984; Goldberg, 1986; Serpell et al., 2009; Van Patten et al., 2015). Like with serial dependence in visual perception, the upper limits of the content hierarchy are unknown, and a future line of research examining the presence or absence of computational perseveration – the degree to which there is dependence on the recent past – in aging and disease during complex decision making may lead to a novel marker of cognitive decline.

2.6 Conclusion

Preferences are by definition subjective. Decades of research into risk preference and temporal discounting have conclusively shown that these concepts, however they may be defined, instantiated, or measured, are variable. One extensive form of variability comes from the multiple well-established methods to elicit these measurements – often either in an experiment or by completing surveys. Further, individuals (and subsequently inferred parameters) demonstrate sensitivity to domains, context (recent history in both choice and stimulus, environmental uncertainty), and adaptation (e.g. shifting reference points, preference reversals). These sensitivities have been demonstrated in healthy individuals and clinical populations, with often meaningful differences between groups. In the growing field of computational psychiatry there is much research focused on linking measures of risk preference and temporal discounting to maladaptive behavior. While there has been much success on this front, understanding and appropriately characterizing these concepts in health and disease is critical.

In this paper, we have reviewed some of the different ways in which these concepts have been characterized and operationalized and have proposed another source of variability that we believe deserves further scrutiny: the explicit influence of recent history on choice behavior, response times, and subsequently inferred values.

We suggest that such trial-level sequential influences are adaptive and consistent with ideas of contextual or local rationality. Ample evidence in the psychophysics and perceptual decision-making literature (amongst others) demonstrates that even when all pains are taken to minimize sequential dependencies within an experiment, the seriality of our temporal experience [and neural processing] plays a profound, arguably causal and adaptive role in shaping behavior. People’s fundamental conceptualizations about parameters and constructs are largely shaped by the functional forms and methods used to describe and infer them – compare classical Bernoulli Utility to Random Utility models, or evidence accumulation models with and without noisy accumulation of evidence, for example. By incorporating trial order (and recent history more generally) into the modeling of risk preferences and discount factors themselves, we hope that the field will move more concretely towards embracing these concepts as inherently, and therefore necessarily, contextual. This could lead towards better reconciling myriad behavioral observations and moving towards a more veridical notion of human rationality.

Chapter 3

Independent not Irrelevant: Trial order causes systematic misestimation of economic choice traits in humans

In fields spanning policy, medicine, and finance, it is increasingly common to guide individual interventions on measures of a person's decision-making characteristics. As the practical application of these measures grows, so have efforts to improve their robustness to ostensibly irrelevant contextual factors such as time of day or ephemeral motivational state. Here, we examine whether such instruments exhibit a fundamental context-dependence: namely, the order in which decision problems are presented. In three datasets evaluating decision-making under different forms of uncertainty, we find systematic, meaningful effects of trial order, which in many cases qualitatively change the measurement's psychological interpretation (e.g. from risk-seeking to risk-averse). We further examine how trial properties modulate this phenomenon and provide an augmented modeling framework to reliably characterize and correct for these effects.

3.1 Introduction

Increasingly, important decisions about an individual are being informed by an assessment of how that individual makes decisions under uncertainty. For example, asset managers measure their clients' risk tolerance when deciding how to allocate their investment portfolios (Kumar & Persaud, 2002), clinicians use the same property to evaluate their substance-using patients' likelihood to relapse (Kwako et al., 2016), and aid workers use measures of intertemporal choice preferences to personalize vaccination incentives (Andreoni et al., 2016). The success of these high-stakes interventions depends on identifying stable traits that will be robust across time and setting. For instance, an intervention with long-term consequences should ideally not depend on whether the measurement was taken on Monday morning versus Wednesday afternoon, or if a local sports team has just won a championship. Indeed, temporal and situational variability in these measures has been widely demonstrated (K. Lempert & Phelps, 2016). Critically, this variability has in some cases been related to directly observable endogenous mechanisms (Lazzaro et al., 2016), and to predict clinically-significant behavioral outcomes (Konova et al., 2020), supporting the construct validity of these measures and further emphasizing the need to distinguish relevant from irrelevant sources of variability.

One approach to isolating meaningful variability has been to incorporate insights from decision neuroscience by modeling the time it takes an individual to make a decision (Clithero, 2018; Konovalov & Krajbich, 2019). This approach produces more robust estimates of individual decision characteristics because it reveals aspects of the decision-making process that are obscured when examining choice alone – e.g. long decision times may indicate near-indifference between two options. Indeed, scientists have demonstrated that sequential sampling models of response times (RT) not only make predictions that capture widely observed effects in RTs during economic choice (Clithero, 2018), but that these models can also be used to derive descriptive theories of choice such as Random Utility (Webb, 2019).

Further, researchers have shown that models that jointly model choice and RT can better describe data across multiple experiments (Peters & D’Esposito, 2020) and that subject RTs from one experiment can predict behavior on another (Konovalov & Krajbich, 2019).

Most decisions are made under at least some degree of uncertainty and under constraints such as limited time or partial information. Choice behavior under uncertainty has therefore been formalized across multiple disciplines with behavioral economics establishing some of the most widely used theoretical and empirical frameworks. Key behavioral economic and psychological parameters of interest include *risk tolerance* (α) which captures how people behave when they have complete information about the underlying probability distribution(s) of reward, *ambiguity tolerance* (β) which captures how people behave when they have partial or no information about probabilities, and *discount factor* (k) which captures how people trade off time and reward. Critically, the number inferred when estimating these parameters is tightly linked to a psychological interpretation – especially in the case of risk and ambiguity.

When making decisions under uncertainty and with limited time, humans and animals must balance efficiency with completeness. One mechanism through which humans may do so is through the efficient coding of valuation information (Zimmermann et al., 2018). The efficient coding hypothesis, originally formalized in perceptual neuroscience, says that resource-constrained organisms leverage environmental structure to maximize information and minimize redundancy (Barlow et al., 1961). A corollary of this hypothesis is that stimulus information is encoded relatively (i.e. what has changed now compared to the previous moment?). This suggests that individuals should be sensitive to trial order over the course of an experiment: the decision they were presented with (and the choice they made on the previous trial) could impact how individuals perceive the current choice problem and the subsequent choice they make. Such behavior would be incongruent with standard behavioral economic tasks which deploy a randomized structure paired with explicit instructions

to participants that they should “treat each trial independently and as if it were the only one that counts.” These sorts of task-incongruent temporal dependencies have been widely established in the visual perception literature where such behavior is termed serial dependence (Fischer & Whitney, 2014). These researchers specify that such a dependency is separate from priming and hysteresis, likely not driven by higher-order processes (Fischer & Whitney, 2014), and is adaptive (Kiyonaga et al., 2017).

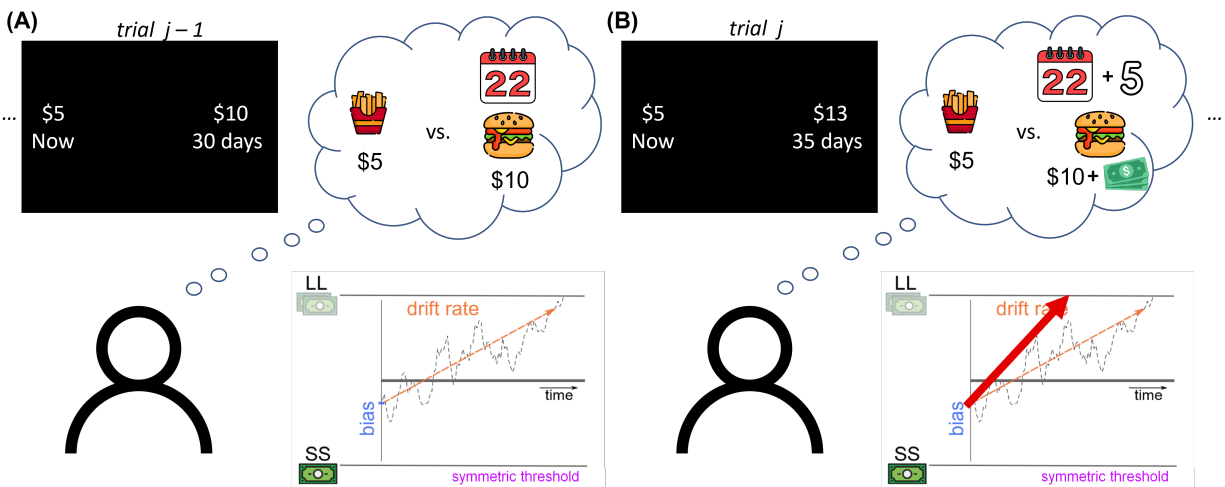


Figure 3.1: **Example of how trial order may influence decision making (choice and RT modeling parameters) in ITC.** The subject participates in the experiment on November 22nd. (A) On trial $j - 1$, they have to choose between \$5 today and \$10 in 30 days. They simulate their future 30 days from now, December 22nd, and imagine what they could buy. This is a difficult choice for the subject and deliberation is captured here by a slow rate of evidence accumulation. The subject ultimately decides to take the \$10 and wait. (B) On the next trial j , the choice is between \$5 today and \$13 in 35 days. Standard models assume that the previous trial should have no influence on the current problem. The spillover hypothesis predicts that a subject could perceive and evaluate the current decision relative to the preceding choice, effectively reusing the outcome of their previous simulation instead of starting afresh: here demonstrated by a steeper drift rate and then a repeated choice.

Taken together, it becomes plausible that one source of variability in inferring key behavioral economic parameters may come from trial order. Further, observed choice, RTs and sequential sampling parameters may also be meaningfully impacted by trial order: while trial-trial variability is accounted for in such models, rarely are influences of the recent past explicitly modeled (e.g. Figure 3.1). Such sequential dependencies are typically and unsur-

prisingly observed in “sequential” tasks like multi-armed bandits where learning rates over the course of an experiment are modeled (e.g. (Bornstein et al., 2017)). We emphasize that the phenomenon of interest here is a higher-order serial dependence or informational spillover, something that is usually not explicitly modeled in these higher-order tasks regardless of whether they are “sequential” or not. In this paper, we develop and deploy a framework that allow us to examine temporal dependencies as a function of decision problem properties, such as reward value, probability, and delay.

In this paper, we demonstrate that trial order, in particular relative differences in successive trial properties, affects decision-making under uncertainty. We show, across three experiments (total $n = 656$), that sequential effects modulate behavioral economic parameters estimated both by jointly modeling response time (RT) and choice behavior, and choice behavior alone (as is standard practice). Experiments 1 and 2 re-analyze previously collected data (as in (N. Banavar et al., 2021), but with different model specifications). We further designed and collected data for Experiment 3 to build on these analyses so that we could more closely examine the specific factors that give rise to sequential effects and how they can contribute to changes in parameter estimates. We demonstrate both high degrees of individual differences and high degrees of systematicity in these effects (i.e. consistency of the sign of sequential effects parameters in the task). Critically, in the third experiment, we show that accounting for such sequential effects not only results in different numerical parameter estimates, but also different psychological interpretations. These differences in interpretation go in both directions (i.e. some previously classified “ambiguity averse” subjects are reclassified as “ambiguity seeking” and vice versa). Taken together, these results suggest that widely-used measures of decision-making traits are mischaracterizing a sizable fraction of individuals, and that experiment randomization alone does not correct for this bias. We conclude with a discussion of how this systematic effect should be accounted for before such measures are used to guide meaningful interventions into individual lives.

3.2 Methods

3.2.1 Experiments and Data

None of the experiments presented involved the use of feedback where choice outcomes were realized over the course of the experiment (i.e. after each trial). In Experiments 2 (ITC) and 3 (AMB), subjects received feedback confirming only their choice selection (i.e. left or right option selected).

Experiment 1: Risky Choice (RISK). We model $n = 56$ subjects who made 80 choices between two lotteries in the gain and loss domain conditions (rewards range: Gain: \$1–\$100, Loss: $-\$99-\0 , probability range: 1%–99%, Figure 3.2A). Gain and loss gambles were not intermixed (i.e. participants made their choices in two conditions: 40 in the gain domain, condition 1, and then 40 choices in loss). All data were collected previously on Amazon Mechanical Turk (for details, see (Guan, 2019)). The experiment was fully randomized with no experimentally-designed trial-level dependencies. Stimuli were displayed numerically (rewards and probabilities) and graphically (probabilities were also presented as pie charts). Subjects had an unlimited amount of time to make their choices, which they indicated via mouse click. In this experiment, subjects were instructed to maximize rewards in the gain domain, and minimize losses in the loss domain.

Experiment 2: Intertemporal Choice (ITC). We model $n = 482$ adult subjects who made a sequence of 102 binary choices between a same-day monetary reward (SS: “smaller sooner”, range: \$1 – \$85) and a larger delayed reward (LL: “larger later”, range: \$10 – \$95; delay range: 4 – 180 days, Figure 3.2B). All data were collected previously in person (for details, see (L. Hunter et al., 2018)). The experiment was fully randomized with no experimentally-designed trial-level dependencies. All stimuli were displayed numerically and counterbalanced so that the SS and LL options occurred equally often on either side of the

computer screen. Subjects had 6s after stimulus onset to make a choice. Each choice was followed by 0.5s of feedback confirming the option selected and then a variable inter-trial-interval (ITI) between 3-5s. In this experiment, subjects were instructed to act in accordance with their genuine preference between the two choice options (i.e. there was no “correct” answer unlike with Experiment 1 (RISK).) The task was also incentive compatible: a single trial was selected at random, realized, and paid out at the end of the experiment as a bonus.

Experiment 3: Risky and Ambiguous Choice (AMB). We model $n = 118$ adult subjects who made a sequence of 196 binary choices between a certain reward (range: \$3 – \$9.5) and a lottery (range: \$5 – \$24), in 4 blocks (Figure 3.2C). The amount a subject could win by choosing the lottery was almost always larger than the certain reward, except during 16 catch trials (4 per block). All data were collected on Amazon Mechanical Turk via psiturk (Gureckis et al., 2016). Lotteries were either *risky* (1/7 of trials) where the full probability distribution was presented graphically (win probabilities: 25%, 50%, 75%) or *ambiguous* (6/7) where partial information was presented (ambiguity levels: 15%, 40%, 60%, 85%). In this experiment, [win] probabilities for any given gamble are represented as the number of red and blue chips out of a total of 100 red and blue chips. Ambiguity levels indicate the degree to which information about win probabilities is limited. This correlates with the size of the grey bar in Figure 3.2C, which shows an example trial with ambiguity level 40%: subjects know that *at least* 30 chips are red and *at least* 30 chips are blue, but they do not have any information about the remaining 40 chips. While we maintained a general “non-sequential” structure, we ensured that 50% of successive trials increased in ambiguity, and 50% decreased. A risk trial followed by an ambiguous trial would be considered as an increase in ambiguity, as risky trials are unambiguous with respect to the probability of reward. Likewise, an ambiguous trial followed by a risk trial would be considered a decrease in ambiguity. Further, we controlled for median risk/ambiguity levels, lottery reward, and fixed reward across blocks. As with ITC, the stimulus options occurred equally often on either side of the computer screen. Subjects had up to 3s after stimulus onset to make a

choice. Each choice was followed by 0.5s of feedback confirming the option selected and then a variable inter-trial-interval (ITI) between 0.5-2s. The task was also incentive compatible: a single trial was selected at random, realized, and paid out at the end of the experiment as a bonus. Participants on Amazon MTurk were bonused 10% of their winnings to be consistent with pay rates on the platform.



Figure 3.2: **Example trials for all three tasks.** (A) E1: Example RISK trial in the Loss domain where a subject has an unlimited amount of time from stimulus onset to choose between either the gamble on the Left or the Right. The Expected Value maximizing (and correct) choice is the gamble on the Left (figure from (Guan, 2019)). (B) E2: Example ITC trial where the subject has up to 6s to chose between \$9 today or \$30 in 42 days (figure from (L. Hunter et al., 2018)). (C) E3: Example AMB ambiguous trial where the subject has up to 3 seconds to make a choice between a certain reward of \$3 and a chance to win \$11 by playing the lottery. The Expected Value maximizing (but not necessarily “correct”) choice is the lottery reward.

We model both choice behavior and response times (RT) in the ITC and AMB tasks. We only model choice behavior in the RISK task, as RT data was not available. For ITC and AMB, we excluded any responses that were made in less than 300ms. We also excluded any missed trials and the trial immediately after them from the following analyses, in addition to the first trial in each block. This is because our primary analysis focuses on one-trial-back effects. Finally, we excluded any subjects that missed more than 25% of trials.

3.2.2 Models: Choice Behavior

Baseline (“Non-Sequential”) Models

For tasks involving immediate uncertainty (RISK and AMB), we pair a logistic choice rule with models consistent with the Subjective Expected Utility Maximization Framework. We implement these models in a hierarchical Bayesian fashion using JAGS (Plummer et al., 2003) to better capture individual differences (Lee, 2018). Unless otherwise stated, all parameters are hierarchical Normals, with hyperprior specifications of mean $\mu \sim Normal(0, 1)$ and standard deviation $\sigma \sim Uniform(0.01, 4)$. A parameter X is thus hierarchically distributed: $X \sim Normal(\mu_X, \sigma_X^2)$.

To infer risk and ambiguity values for the RISK and AMB tasks, we model Subjective Value as follows:

$$\begin{aligned}
 \textit{Subjective Value Lottery}_{Risk} &= p \cdot v^\alpha & (3.1) \\
 \textit{Subjective Value Lottery}_{Ambig} &= [p - \beta(A/2)] \cdot v^\alpha & \textit{Subjective Value Fixed} = v^\alpha
 \end{aligned}$$

Here, p is the objective probability of a reward (risk level, or $p = 0.5$ on ambiguous trials per (Levy et al., 2010)). A represents the degree of ambiguity on the trial (ambiguity level, $A = 0$ on risk trials and for all of the RISK task). Finally, v represents the monetary reward associated with that lottery. The key behavioral economic parameters of interest, then, are α which is risk tolerance, and β which is ambiguity tolerance. In this formulation, both parameters are subject-specific. We use hyperprior $\mu_\alpha \sim Gamma(2, 1)$ for risk tolerance, with mode = 1 (risk neutrality) and $\mu_\beta \sim Normal(0, 1)$ for ambiguity tolerance, with mode = 0 (ambiguity neutral). For the RISK task, we further allow for the curvature of the utility function to differ as a function of domain c (i.e. infer $\alpha_{(i,c)}$), where $c = 1$ for gain domain and $c = 2$ for loss.

To infer temporal discounting for the intertemporal choice task, we model Subjective Value using a hyperbolic discounting function:

$$\text{Subjective Value Future} = \frac{v}{1 + k \cdot D} \quad (3.2)$$

Here, v again represents the monetary reward associated with that lottery and D represents the delay with the future (LL) reward is offered. The key behavioral economic parameter of interest is k which is the individual’s discount factor. We use hyperprior $\mu_k \sim \text{Beta}(1, 1)$ to be “uninformative.”

Our logistic choice rule, which has the same basic parametrization across all three tasks, for subject i on trial j is as follows, where $\theta_{A,B(ij)}$ is the probability of choosing Option A :

$$\theta_{A,B(ij)} = \frac{1}{1 + \exp(\gamma_i + \phi_i \cdot SVD_{ij} + \epsilon_{ij})} \quad y_{A,B(ij)} \sim \text{Bernoulli}(\theta_{A,B(ij)}) \quad (3.3)$$

Here, SVD_{ij} represents the difference in the Subjective Value between the two options presented on any given trial. γ_i represents the shift, or bias, in a decision. ϕ_i represents response variability, and we use hyperprior $\mu_\phi \sim \text{Gamma}(2, 1)$, where the mode corresponds to probability matching. Finally, ϵ_{ij} represents effects of simple choice – or motor – perseveration (repeat previous choice made). All parameters allow for variability at the individual and, in the case of the RISK task, domain (gain or loss) level. We pair these prior specifications with a $\text{Bernoulli}(\theta_{A,B(ij)})$ likelihood, as no two stimuli are presented together more than once.

Sequential Effects

Intuitively, we might imagine that there would be more (less) of an effect on a given parameter on sequential trials that present the subject with similar (different) values for the decision problem regardless of task structure or goals: e.g., if on RISK trial $j - 1$, a subject decides

between a 81% chance of winning \$41 or a 55% chance of winning \$39, and the next trial j asks the subject to choose between a 85% chance of winning \$45 or a 55% chance of winning \$37, there might be little need to re-deliberate, which could thus yield an effect on either choice or response time (refer to Figure 3.1 for example in ITC). This may be the case even if subjects are told to treat decisions independently, as they typically are in these “non-sequential” behavioral economic tasks, and trials themselves are randomized.

This task-incongruent transient reliance on recent history (henceforth also referred to as perseveration) can manifest in one of many ways, ranging from “lower-order” (e.g. motor) to “higher-order” (e.g. carrying over frontal cortex dependent computations) (Peters & Büchel, 2011; Levy et al., 2010). We explicitly test for three qualitatively different types of perseverations (Table 3.1 *Lower*): motor or choice (i.e. repeat the same choice previously made), perceptual (i.e. on logistic bias or variability), and computational (i.e. influencing people’s risk/ambiguity/impulsivity preferences). We define these perseverations as driven explicitly by cross-trial *differences* in the trial properties listed in Table 3.1 *Upper*, as these are the normative properties involved in choice computations. As we further expect these influences to be transient due to the complex and temporally constrained (ITC, AMB) choice environment, we restrict our analyses to one-trial-back (i.e. differences in current and immediately preceding trial properties).

Specifically, we use Indicator Variables to subset increases or decreases in specific sequences of trials, resulting in a 8-fold tiling of trial property space for all three experiments (e.g. for ITC: increase in delay, decrease in delay, increase in value, decrease in value, increase in value and delay etc.) We then augment our baseline models by allowing these trial properties to exert linear additive influences on the parameters of interest. For example, if we consider ITC trials that increase in value from trial $j-1$ to trial j :

$$\gamma'_{0,ij} = \gamma_{0,i} + \eta_i \cdot 1([V_{a,j} - V_{b,j}] > [V_{a,j-1} - V_{b,j-1}]) \quad (3.4)$$

Property Type	RISK	ITC	AMB
Uncertainty Reward	Gamble Entropy (H)	Delay (D) value	Risk/Amb levels (RA)
Normative Interaction	Expected Value (EV) EV x H	value x D	EV EV x RA

Parameter	
Logistic Bias	ϕ
Logistic Slope	γ
Choice Perseveration	ϵ
Risk Tolerance	α
Ambiguity Tolerance	β
Discount Factor	k

Table 3.1: **Trial properties considered for sequential effects and the choice parameters on which we test for these effects.** *Upper* Differences in trial properties considered as potential drivers of sequential effects. $H = -\sum p \log(p)$ is the Shannon Entropy of a gamble. *Lower* Parameters we simultaneously test for sequential effects by allowing them to vary trial-trial as a function of relative differences in properties as described in Table 1 a. We test the first three parameters in all three tasks. The final three are task-dependent (RISK: α ; ITC: k ; AMB: α, β).

Thus, in Equation 3.4, $\gamma_{0,i}$ becomes the sequential-effect-adjusted logistic bias for individual i and the indicator variable is 1 if there is an increase in value difference from trial $j - 1$ to trial j . The posterior value of η_i tells us how much an individual is weighting relative changes in trial properties (i.e. increase or decrease). We define η_i hierarchically: $\eta_i \sim \text{Normal}(\mu_\eta, \sigma_\eta)$.

Our fundamental analysis is structured around a hypothesis test on η_i : $H_0 : \eta_i = 0$ vs. $H_a : \eta_i \neq 0$. We quantify the strength of evidence in favor of the null and alternative by using the Savage-Dickey ratio to estimate the Bayes Factor (BF). This ratio compares prior and posterior density at any point in parameter space (i.e. $\eta_i = 0$). As others have done, we interpret any values of $BF > 3$ as evidence in favor of our alternative hypothesis (Lee & Wagenmakers, 2014).

3.2.3 Models: Response Times

Recall that we only have response time data for ITC and AMB.

Baseline Models

We implement a hierarchical Bayesian drift diffusion model (DDM) to model response times using the Wiener module (Wabersich & Vandekerckhove, 2014) in JAGS (Plummer et al., 2003) for both the ITC and AMB tasks. That is, for subject i and trial j , we model observed response time as Wiener first passage time (*wfpt*) distributed:

$$RT_{ij} \sim wfpt(\alpha_i, \tau_i, \beta_i, \delta_{ij})$$

Here, α_i represents the subject-level threshold or boundary separation, τ_i is the subject-level non-decision time (processes ostensibly unrelated to the value-based decision process), β_i is the subject-level bias ($\beta_i < 0.5$ bias towards immediate option in ITC and towards the fixed option in AMB), and δ_{ij} is the subject-and-trial-level drift rate (the rate of evidence accumulation). We model all these parameters as hierarchical Normals in order to better capture individual differences (Lee, 2018). For α_i, τ_i , and β_i , we use the same prior and hyperprior specifications for both tasks, referencing (Wabersich & Vandekerckhove, 2014) for mean hyperpriors and using 'noninformative' priors for the standard deviation:

$$\begin{aligned} \mu_\alpha &\sim Uniform(0.001, 3) & \mu_\tau &\sim Uniform(0, 0.6) & \mu_\beta &\sim Uniform(0.01, 0.99) \\ \sigma_\alpha, \sigma_\tau, \sigma_\beta &\sim Uniform(0.01, 4) \end{aligned}$$

Similar to previous work (Peters & D'Esposito, 2020), we take a cognitive psychometrics approach to modeling the drift rate. Critically, however, we allow the drift rate to be driven by *objective* trial properties (i.e. not Subjective Value or even differences in SV)

and normative combinations of these trial properties (e.g. Expected Value). This is because explicitly relating untransformed trial properties to elements of the decision process is critical for our question of interest. Incorporating transformed trial properties like Subjective Value might perpetuate the very biases we seek to mitigate as they would be inferred without accounting for potential effects of trial order. We keep the broad functional relationship between trial properties as dictated by behavioral economic models of choice behavior (e.g. allowing an inverse relationship between the drift rate and delay for ITC). We also normalize all stimulus properties such that they fall between 0 and 1. Then, for subject i and trial j :

$$\delta_{ITC,ij} = \beta_{0,i} + \beta_{1,i} \cdot (value_{LL,ij} - value_{SS,ij}) + \beta_{2,i} \cdot delay_{ij}^{-1} \quad (3.5)$$

$$\delta_{AMB,ij} = \beta_{0,i} + \beta_{1,i} \cdot (EV\ Diff) + \beta_{2,i} \cdot A_{ij} \quad (3.6)$$

In Equation 3.6, *EV Diff* represents the difference in Expected Value ($EV_{Lottery} - EV_{fixed}$) between the two options presented. $EV_{Lottery} = p \cdot v$, where p is the objective probability of reward (i.e. risk level) and v retains its interpretation of objective reward. For an ambiguous trial, $p = 0.5$ as in (Levy et al., 2010). The Expected Value of a certain reward is simply v . On ambiguous trials, A is the function of the lottery that is occluded by the grey bar as seen in Figure 3.2C. Recall that on risky trials, $A=0$. Finally, we allow all drift rate decomposition parameters β s to be hierarchical Standard Normals.

Sequential Effects

We allow all sequential effect parameters to be hierarchical Standard Normals. We simultaneously assess the influence of relative trial properties on all *drift rate decomposition* (i.e.

β s) and *bias* parameters (same properties as in choice, Table 3.1A). For example:

$$\beta'_{0,ij} = \beta_{0,i} + \eta_i \cdot 1([V_{a,j} - V_{b,j}] > [V_{a,j-1} - V_{b,j-1}]) \quad (3.7)$$

Thus, in Equation 3.7, $\beta_{0,i}$ becomes the sequential-effect-adjusted drift rate “regression” intercept for individual i and the indicator variable is 1 if there is an increase in value difference from trial $j - 1$ to trial j . Just as in Choice Behavior (section 3.2.2), we test whether the sequential effect parameters (i.e. η_i) is non-zero using the Savage-Dickey ratio to approximate the Bayes Factor (BF). We interpret any $BF > 3$ as evidence in favor of sequential effects.

3.3 Results

We analyze data from two previous experiments (Experiment 1 (RISK): Risky decision-making task in gain and loss domains $n=56$; Experiment 2 (ITC): Intertemporal choice task $n=482$) and one new experiment (Experiment 3 (AMB): Risky and ambiguous decision-making task in gain domain $n=118$) (Figure 3.2). We explicitly incorporate sequences of trial properties into standard choice and response time models using hierarchical Bayesian modeling, considering one “sequence type” at a time. Critically, this approach allowed us to estimate the reliability of sequential effects at an individual-subject level.

3.3.1 Sequential Effects in Choice Behavior

We first assessed whether choices exhibited trial property-dependent sequential effects at a broad, “model-free,” level. Our primary assessment of model free signatures of sequential effects in the RISK and AMB tasks involved comparing whether the proportion of times a

subject selected a choice differed across trial sequence types. The trial properties of interest include cross-trial differences in Expected Value and entropy (see table 1A for complete list). Specifically, we look at pairs of trials as differences of differences: for example, an increase in Expected Value difference between successive trials means that the choice options on the current trial are more distinct than the choice options on the previous trial. This could correspond to a current trial being relatively easier than the previous. This interpretation is consistent with the experimental data on a whole.

In RISK, subjects chose the EV maximizing option significantly more often not just when a trial was relatively easy with respect to the rest of the experimental choice set, but also when there was a relative increase in EV from the previous trial to the next (EV increase 71%, vs. EV decrease 64%, Wilcoxon Rank Sum $V = 1276.5$, $p < 0.001$), but no significant difference when comparing proportions of EV maximizing choices when trials increase or decrease in gamble entropy (H increase 67%, H decrease 68%, $V = 705$, $p = 0.45$, Figure 3.3A). In AMB, we find a similar pattern: subjects chose the lottery option significantly more often when there was a relative increase in EV from the previous trial to the next vs. when there was a relative decrease (EV increase 50%, EV decrease 47%, Wilcoxon Rank Sum $V = 3693$, $p < 0.001$), but not when considering only increases or decreases in gamble ambiguity (Amb increase 48%, Amb decrease 50%, $V = 2441$, $p = 0.43$, Figure 3.3B). We find no differences in choice proportions as a function of trial properties (value, delay increase/decrease) in the ITC task. This simple analysis suggests that, in at least two of the three types of choices evaluated, people are not just sensitive to trial order but more specifically to the relative differences in normative calculations - the relative ease or difficulty of the current choice problem.

For our model-based analyses, we augment logistic choice rules to include such relative cross-trial differences on multiple parameters as described in Methods. That is, we allow for trial order to potentially manifest as a variety of qualitatively different types of “perseverations”

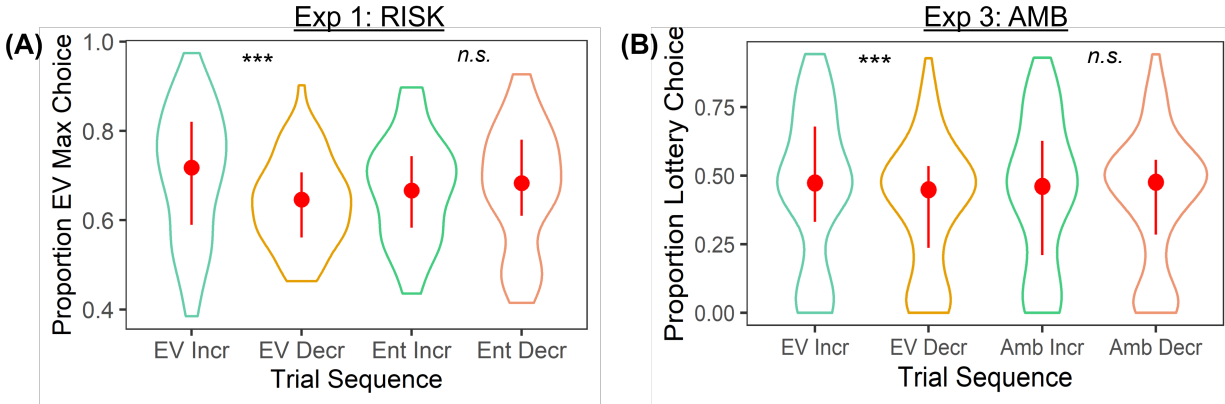


Figure 3.3: **Participants in RISK and AMB tasks make distinct choices depending on trial sequence.** (A) In the RISK task, they are more likely to make the EV maximizing choice between two gambles when the previous trial was more difficult (i.e. choices less distinct). (B) In AMB, they are more likely to pick the lottery option as opposed to the fixed when the previous choice was easier (i.e. less difference between options). Conversely, there was no significant difference between EV maximizing choices when considering relative increases or decreases in trial entropy (A) or trial ambiguity (B). Overlaying the violin plots are the median and IQR.

– from motor, to perceptual, to cognitive. In RISK, we observed reliable sequential effects on logistic slope and risk tolerance for 7% of individuals. Critically, and consistent with these effects being cognitively specific, these individuals only had non-zero sequential effects for specific sequences of trials: when a trial with a high difference in Expected Value between the two options was followed by a trial with a low difference in EV – “easy” then “difficult” in sequence. For example, for a specific subject, an initial risk tolerance of 1.043 updated to 1.205 when adjusted for this sensitivity. Importantly, the magnitude of risk tolerance was not the only changing factor: the interpretation of the individual’s risk tolerance changed from risk neutral to risk averse in the loss domain. For the other two experiments, unlike our RISK analysis, our hypothesis test of the posterior values of the sequential effect terms did not result in any evidence in favor of trial-trial sensitivity. However, we note that we also did not find strong enough evidence in favor of the null hypothesis.

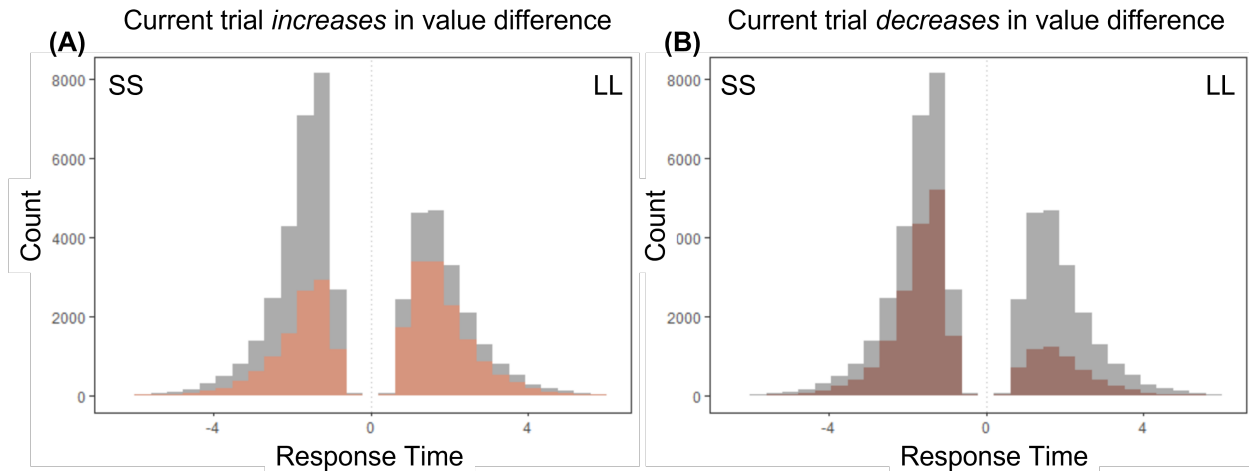


Figure 3.4: **Choice and RT in ITC change as a function of trial properties.** In the intertemporal choice task, subjects took a median of 1.64s when they chose the Smaller Sooner (SS) option, and 1.71s when they chose the Larger Later (LL) option. Grey histograms are RTs for all subjects, all trials. Negative RTs correspond to making the SS choice, while positive correspond to LL. (A) When the current trial increased in value difference from the previous trial, subjects took a median of 1.64s to make a SS choice and 1.68s to make a LL choice (overlaid orange histograms). (B) However, when the current trial decreased in value difference from the previous trial, subjects took a median of 1.35s to make a SS choice and 1.77s to make a LL choice (overlaid brown histograms).

3.3.2 Choice and Response Time

We find “model-free” evidence suggesting that the amount of time subjects take to make their decision can vary as a function of relative differences in trial properties (Figure 3.4). Subjects are faster to go with the SS option when the current trial has a lower value difference between SS and LL than the previous trial. This is intuitive as a smaller difference reduces the difficulty of the computation and for small enough differences, may negate the influence of delay duration.

Drift Rate Decompositions

ITC	Mean(CI)	AMB	Mean(CI)
β_0	-0.66(-1.26, -0.06)	β_0	-0.90(-1.12, -0.68)
β_1	3.45(1.19, 5.70)	β_1	1.51(1.04, 1.94)
β_2	2.34(-2.15, 6.84)	β_2	0.03(0.001, 0.1)
α	2.47(1.70, 3.24)	α	1.84(0.96, 2.71)
τ	0.76(0.36, 1.16)	τ	0.50(0.03, 1.03)
<i>bias</i>	0.51(0.39, 0.63)	<i>bias</i>	0.50(0.4, 0.6)

Table 3.2: **ITC + AMB: Drift rate decompositions capture meaningful variance in both ITC and AMB.** Each cell shows aggregate posterior means (95% Credible Intervals) for Drift Rate decompositions and other DDM parameters. Bolded parameters are ones we also test for sequential effects.

We find that for the ITC task, subjects tend to accumulate evidence more quickly when the value difference increases (β_1), all else held constant. Similarly, subjects tend to accumulate evidence more quickly when the delay decreases (β_2) – recall that we parameterize delay as $delay_{ij}^{-1}$ – all held constant. Both make sense intuitively, as larger value differences might push individuals towards selecting the LL option, and delayed rewards offered in the far future may not be worth the wait.

With the AMB task, the average subject’s drift rate increases as the Expected Value difference between choice options increases (β_1), all else held constant. Subjects seem nominally sensitive to the degree of Ambiguity during evidence accumulation (β_2). We also highlight the credible interval ranges to confirm that there are considerable individual differences, as we might expect. We note that we considered many different combinations of trial properties as drivers of drift rate and present these model results (Equation 3.5, 3.6) as they best fit the data ($DIC = 19,335$ vs $DIC > 20,000$ for other models).

ITC	Proportion	AMB	Proportion
<i>value</i> ↑	1	<i>value</i> ↑	0.90
<i>value</i> ↓	0.998	<i>value</i> ↓	0.94
<i>delay</i> ↑	1	<i>amb</i> ↑	0.91
<i>delay</i> ↓	1	<i>amb</i> ↓	0
<i>v. ↑ d. ↑</i>	1	<i>v. ↑ a. ↑</i>	0.97
<i>v. ↑ d. ↓</i>	1	<i>v. ↑ a. ↓</i>	0.10
<i>v. ↓ d. ↑</i>	1	<i>v. ↓ a. ↑</i>	0.92
<i>v. ↓ d. ↓</i>	0.84	<i>v. ↓ a. ↓</i>	0

Table 3.3: **ITC + AMB: Most participants are sensitive to trial sequences.** The majority of subjects show sensitivity to trial sequences in both ITC and AMB tasks. Each cell shows the proportion of subjects that demonstrated sequential effects ($BF > 3$) on any one of the drift rate decomposition parameters or bias. Each row represents specific successive trial properties (e.g. *value* ↑ subsets successive trials that increased in value difference as noted in Equation 3.4). The top four rows can be thought of as “main effects” of specific trial properties and the bottom four “interactions.”

We interpret the threshold and non-decision time parameters being greater for ITC to be a reflection of differences in task structure. Recall that subjects have up to 6s to respond in the ITC task, whereas they only have up to 3s in the AMB task. Finally, aggregate posterior means suggests that subjects are generally unbiased in both tasks, which is in contrast with their choice behavior: in ITC subjects chose the smaller sooner reward more frequently and in AMB chose the certain reward more frequently.

Sequential Effects

As described above, response time models provide more reliable estimates of individual-subject decision processes. We therefore analyzed trial-type-specific effects using an augmented drift-diffusion model in place of the logistic choice rule. We performed these analyses for ITC and AMB as RT data was unavailable for RISK. This approach allowed us to identify

sequential effects at the level of both bias (pre-choice inclination) and deliberation (evaluative processing of choice properties). We also consider models that include the Intertrial Intervals as trial properties of interest and find, interestingly, that there is no meaningful influence of ITI on our models of choice and RT.

We observed reliable trial-property-driven sequential effects on the DDM decomposition (Equations 3.5, 3.6) and bias terms for both ITC and AMB tasks. For almost every possible combination of stimulus sequences, 100% of subjects showed evidence of sensitivity to previous stimuli in the ITC task (Table 3.3, Table S1). Effects in the AMB task are more specific to the combination of trial properties considered – in particular, individuals seem to be particularly sensitive to relative cross-trial differences in Expected Value (i.e. the choice becomes easier or harder) (Table 3.3, Figure S1, Table S2). Interestingly, while in the ITC data we find effects on both the drift rate and bias terms, in AMB we only find evidence of sequential effects on the drift rate terms. This suggests potential differences in both the cognitive process that drive spillover but also, more evidently, the choice problem at hand (evaluating future uncertainty vs. present is distinct from evaluating immediate uncertainty in the form of gambles).

To explore how sequential effects may change as a function of experiment length, we fit our models to subsets of the AMB data. As the AMB task was already divided into blocks, we considered subsets of block 1 alone, blocks 1 and 2, and blocks 1, 2, and 3. We found, unsurprisingly, that greatest proportion of sequential effects were recovered when considering the entirety of the experiment (i.e. all four blocks). Further, we found that the trial properties subjects were sensitive to changed over the course of the experiment, resulting in a more global sensitivity in the final block of the experiment.

	Block 1	Block 1+2	Block 1+2+3	All Blocks
<i>value</i> ↑	0	0	0	0.9
<i>value</i> ↓	0	0	0.31	0.94
<i>amb</i> ↑	.07	0.18	0.17	0.91
<i>amb</i> ↓	0.45	0	0	0
<i>v.</i> ↑ <i>a.</i> ↑	0.68	0	0.30	0.97
<i>v.</i> ↑ <i>a.</i> ↓	0.06	0.14	0.03	0.1
<i>v.</i> ↓ <i>a.</i> ↑	0	0.78	1	0.92
<i>v.</i> ↓ <i>a.</i> ↓	0	0	0	0

Table 3.4: **E3: Evolution of sequential effects over task.** We find that the proportion of subjects that show sensitivity to sequential effects changes over the course of the task, with the majority of recovered sensitivities crossing a statistically meaningful threshold presenting when considering all trials in the experiment. Overall, subjects tend to be sensitive increases in ambiguity

Changes in parameter estimates.

The reader, especially if they are interested in applications of behavioral economics, may wonder why accounting for sequential effects should matter beyond finding varied individual differences. First, magnitude information is valuable especially for researchers interested in temporal fluctuations in ambiguity/risk attitudes. We further argue that adjusting for these effects can meaningfully change inferences on the parameters we care about. We thus examined how accounting for sequential effects can reveal qualitative changes in the characterization of individuals. We fit effect terms capturing spillover across all “main” (Table 3.3) trial sequences at the same time (as opposed to one-at-a-time), generate 1000 datasets from these model fits, and redeploy a logistic choice rule to infer risk and ambiguity tolerance as in the standard approach. That is, for the AMB task, Equation 3.7 is instead written as follows (see Table 3.1a for properties for all trials):

$$\beta'_{0,ij} = \beta_{0,i} + \eta_{i,01} \cdot \text{EV Increase} + \eta_{i,02} \cdot \text{EV Decrease} + \eta_{i,03} \cdot \text{Amb Increase} + \eta_{i,04} \cdot \text{Amb Decrease} \quad (3.8)$$

ITC	β_0	β_1	β_2	<i>bias</i>
<i>value</i> \uparrow	0.34(0.02)	-2.72(0.07)	*	-0.06(0.004)
<i>value</i> \downarrow	-0.50(0.02)	3.67(0.10)	-2.12(0.38)	0.05(0.004)
<i>delay</i> \uparrow	-0.35(0.01)	-0.14(0.04)	9.95(0.4)	0.03(0.002)
<i>delay</i> \downarrow	0.32(0.01)	0.19(0.04)	-7.77(0.29)	-0.03(0.004)
<i>v.</i> \uparrow <i>d.</i> \uparrow	0.08(0.02)	-1.06(0.05)	6.06(1.00)	*
<i>v.</i> \uparrow <i>d.</i> \downarrow	0.45(0.01)	-0.63(0.04)	-1.35(0.23)	-0.06(0.004)
<i>v.</i> \downarrow <i>d.</i> \uparrow	*	3.14(0.12)	-5.78(0.31)	*
<i>v.</i> \downarrow <i>d.</i> \downarrow	*	1.32(0.22)	-2.36(0.38)	*

Table 3.5: **ITC: Widespread sequential effects across drift rate and bias terms.** Participants in ITC show sequential effects across all trial sequence types and all four RT parameters (see Equation 3.5). All results reported in a cell are posterior sequential effect group means (*mean (standard deviation)*) and have a BF > 3 if they are non-zero. If instead there is a H_0 then we find evidence (BF > 3) in favor of the null. Finally, if a cell contains an asterisk (*) then the data does not contain enough evidence to favor either the null or alternative hypothesis.

We refer to these models as “stacked” as they include all potential “main effects” of trial properties. We report the models that contained only the main effects for two reasons: 1) these were the primary drivers of sequential effects in our above analysis for both tasks – especially AMB – and 2) we wanted to avoid “over-parametrization” by including the interaction terms. As this parametrization adds 12 more variables to the model, we define the sequential-effect terms (η s) as non-hierarchical Standard Normals: $\eta_{i,xx} \sim \text{Normal}(0, 1)$ to aid model convergence. All models reported converged according to standard metrics.

Differences between the original parameter estimates and those from the simulated data would strongly suggest that cross-trial temporal dependencies are capturing something fundamental in human decision making under uncertainty. We report ratios between simulated data parameter estimates and the baseline model. If $0 < |\text{Ratio}| < 1$, then the sequential effect adjusted parameter estimates are numerically smaller than than the initially inferred, and $|\text{Ratio}| > 1$ the converse. The median ratio change for risk tolerance across all subjects is 0.857 ($IQR = 0.716$), suggesting that models without sequential effects tend to overestimate an individual’s risk tolerance. Likewise, the absolute value of median ratio change for ambiguity across all subjects is 0.574 ($IQR = 1.13$).

AMB	β_0	β_1	β_2	bias
<i>value</i> \uparrow	H_0	-0.38(0.12)	H_0	H_0
<i>value</i> \downarrow	*	0.45(0.11)	*	H_0
<i>amb</i> \uparrow	H_0	*	-0.29(0.07)	H_0
<i>amb</i> \downarrow	H_0	H_0	0.32(0.09)	H_0
<i>v.</i> \uparrow <i>a.</i> \uparrow	H_0	*	H_0	H_0
<i>v</i> \uparrow <i>a.</i> \downarrow	H_0	*	0.59(0.13)	H_0
<i>v.</i> \downarrow <i>a.</i> \uparrow	H_0	*	*	H_0
<i>v.</i> \downarrow <i>a.</i> \downarrow	*	0.5(0.16)	H_0	H_0

Table 3.6: **AMB: Sequential effects are restricted to drift rate decomposition parameters.** Sequential effects are more selective compared to ITC. Specifically, they are related to online evidence accumulation and are only on the terms that include trial property information (see Equation 3.6). Cell interpretations are as in Table 3.5.

Critically, we find evidence not just of differences in sequential-effect-adjusted parameter estimates (Figure 3.6A,C) but also widespread qualitative changes in interpretation (Figure 3.6B,D). For ambiguity tolerance, this change moves in both directions (i.e. some subjects are newly classified as ambiguity averse or ambiguity tolerant), but for risk tolerance, all reclassified subjects move from risk seeking to risk averse.

3.4 Discussion

Measures of individual decision-making traits are becoming increasingly used in applied settings. However, there are many sources of variability that can bias these measures. In this paper, we present a quantitative behavioral analysis of one of these potential sources of bias: trial-trial dependencies in experiments without feedback. We build off a neuroeconomics literature that demonstrates the importance of trial order in how neurons represent value signals (Zimmermann et al., 2018), and visual perception studies (Fischer & Whitney, 2014) that establish serial dependence. By explicitly incorporating trial-order information into behavioral models we find: a) differential sensitivity across experiments that involve different types of stimulus-level uncertainty, b) greater prevalence across subjects when jointly

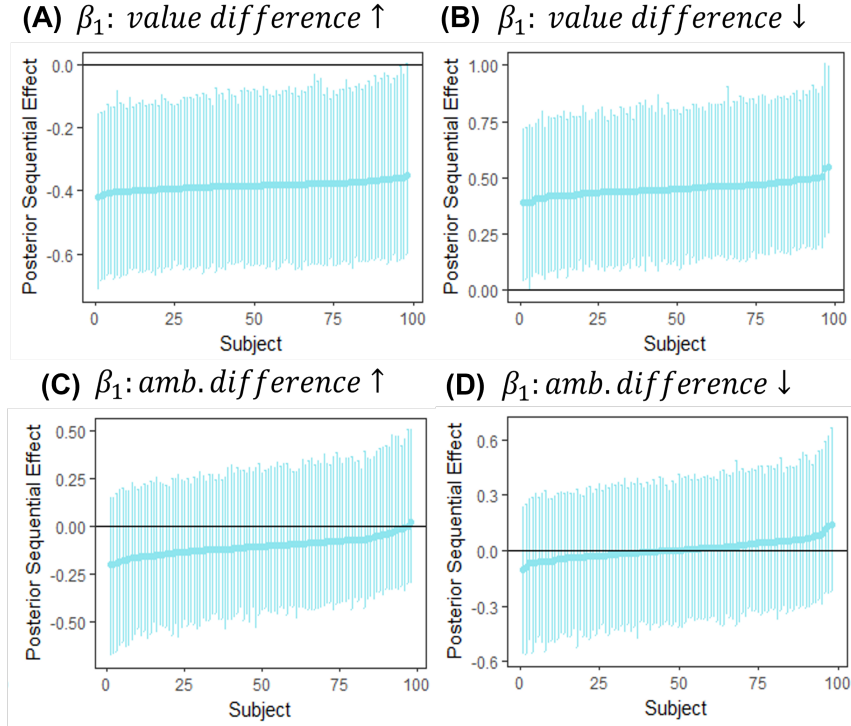


Figure 3.5: **Participants in the AMB task show sequential dependencies on the drift rate EV term: “main effects” of value difference but not ambiguity difference.** Sorted posterior 95% Credible Intervals of sequential effects on the drift rate Expected Value term β_1 , Equation 8, when successive trials (A) increase in value difference, (B) decrease in value difference, (C) increase in ambiguity difference, and (D) decrease ambiguity difference as summarized in table S4.

modeling choice and response time, c) evidence accumulation is impacted by trial-order across tasks, and d) individual-level posterior inferences on key parameters of interest can be meaningfully changed by accounting for these effects.

Our analyses highlight that sequential effects in both choice behavior and reaction time are a function of individual differences, with non-trivial changes in parameter magnitude and interpretation. This work has deep theoretical and empirical implications. Firstly, our analyses suggest that these sequential effects are not noisy artefacts but are instead the consequence of a systematic influence of trial properties on components of the decision process. This suggests a potential need for the theoretical re-conceptualization of *experimentally-inferred* parameters as explicitly dynamic and sensitive to (highly) local contexts and not *solely* a

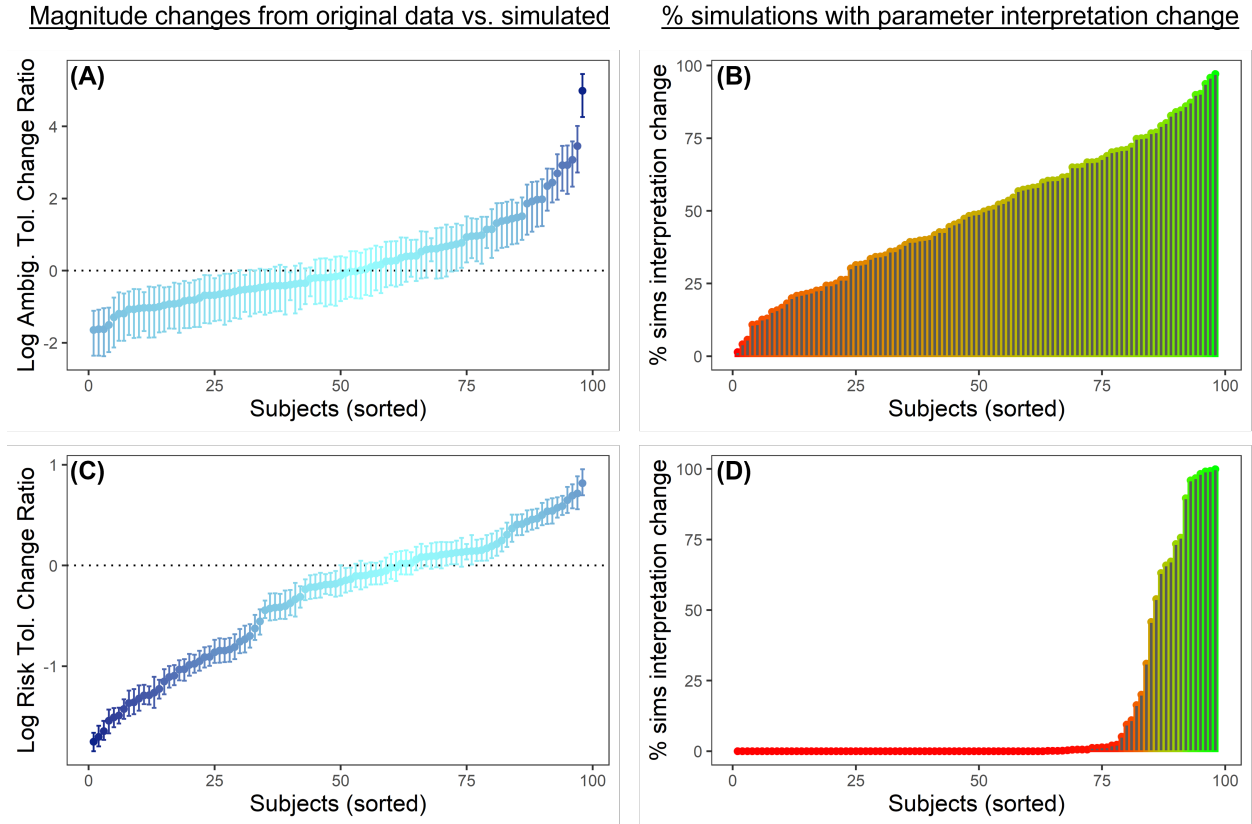


Figure 3.6: **AMB task: Magnitude and interpretation changes in (A, B) ambiguity and (C, D) risk tolerance.** *Left:* Ratio of (A) ambiguity tolerance and (C) risk tolerance estimates: $\log(\text{simulated choice set}/\text{observed data})$. We plot median ratios and IQRs from 1000 simulated choice sets for 98 subjects. *Right:* The percentage of simulation-fit parameters that change interpretations in (B) ambiguity and (D) risk attitudes when compared to parameter fits in the original data. Subjects re-sorted by effect size in each plot.

static and psychologically interpretable end. Secondly, much work has shown that parameters inferred from experiments tend to correspond poorly with real-world behavior and with other tasks that purportedly measure the same construct (e.g. sequential risk tasks vs non-sequential) (Frey et al., 2017) and the approach demonstrated here may be one way to more closely reconcile these discrepancies. Thirdly, there is much valuable information that can be gained from the joint modeling of choice and response times, and scientists should aim to collect RTs whenever possible (Clithero, 2018; Konovalov & Krajbich, 2019; Krajbich et al., 2012). Finally, the widespread nature of our results could suggest that susceptibility to sequential effects in complex choice under uncertainty could be a by-product of the rational

use of limited resources (Dasgupta et al., 2018).

It is tempting to interpret these sequential effects as “corrupting” the choice process and parameter inference, despite possible rational justifications (Dasgupta et al., 2018; Daw et al., 2008). While we show that our posterior parameter estimates do indeed differ when we account for relative differences in trial properties, we suggest three empirical alternatives to how we can move forward with this information. Firstly, as suggested in the previous paragraph, we do explicitly consider the degree of sequential dependency to be informative and report two parameters with every task: the parameter the experiment was built around (e.g. risk tolerance) *and* some parameter or summary that reflects sequential sensitivity. Secondly, experimenters could consider interleaving distractor tasks to clear the working memory buffer that may generate sequential effects. Thirdly, in a similar vein, experimenters may consider merging relatively unrelated experiments together so that each subsequent trial may come from any one of n tasks. While some may argue that this strategy is folly to the same idea behind how experiments are currently designed – that randomized structure does more than what we in this paper demonstrate it does – this is why we suggest unrelated but relevant (i.e. non-distractor) tasks. In this paper, we remain agnostic about the circumstances and mechanisms through which sequential dependencies are generated. However, Table 3.4 suggests that subjects may be differentially sensitive to trial properties as they gain more experience in the experiment. Initially, we recover trial-trial sensitivities when ambiguity decreases and when both value and ambiguity increase. A reduction in ambiguity for the current trial suggests that decision ease could play an important role. Likewise, an increase in both value and ambiguity suggests that subjects must consider both attributes (uncertainty and reward) and that decision difficulty matters too. The proportion of subjects sensitive to these trial combinations changes over experience. This suggests that sequential dependencies are not uniformly present throughout the experiment. Future work is necessary to more carefully tease apart the interaction between experiment duration, sequential effects, and parameter inference for risk, ambiguity, and discount factor.

While there are a great many open datasets available for both ITC and Risk/Ambiguity tasks, few contain RT information, and even fewer adhere to the strict “non-sequential” scope of experiments under consideration in this paper (e.g. staircasing tasks are not “non-sequential” as subsequent trials are explicitly determined by previous trials and choices.) Another limitation in these analyses is that we consider only “non-sequential” experiments. The inclusion of tasks which require learning - and the testing of whether there is “learning” in these environments adds computational complications that are out of the scope of the current paper.

Results from the AMB study suggest that differences in reward magnitudes (in addition to increases in ambiguity) play an important role in these effects: people seem to be evaluating current rewards relative to what they had immediately seen before. One possibility as to why people may be sensitive to relative value differences might have to do with the logarithmic and noisy encoding of magnitudes (Khaw et al., 2021). A logarithmic internal representation of rewards presented in the experiment would have smaller rewards more closely represented on the “number line” (allowing for greater discrimination) and larger rewards to be disproportionately further apart. A natural consequence of using such a number line would mean that relative differences between successive sets of choice options, especially for options on opposing ends of the line, would be perceived to be greater than their absolute magnitudes. This is akin to a transient [numerical] contrast effect. In fact, this line of work by Khaw and colleagues argues that hallmarks of human behavior under uncertainty – like risk aversion in the gain domain – can actually be better explained by these (Bayes optimal) noisy logarithmic coding models relative to standard power utility models (Khaw et al., 2021).

This is not to suggest that relative differences in how uncertainty in these experiments is represented internally do not also contribute to the generation of sequential effects. In this paper, we have considered three qualitatively different types of uncertainty: immediate with

complete probability information (risk), immediate with incomplete probability information (ambiguity), and temporal with 100% probability of reward (temporal discounting). The cognitive processes invoked in managing these types of uncertainty, especially in sequence, may vary as a function of task – in particular risk/ambiguity vs. temporal discounting. Indeed, popular theories of intertemporal choice involve simulating the future (Peters & Büchel, 2011). This is a potentially a resource-intensive process that conflates uncertainty and concreteness – simulations further out in the future are also less likely to be concrete (Leiser et al., 2008). This may lead to the reuse or indeed reification of a simulated future that is n or greater days from the present. While such a future may not directly lead to one-trial-back sequential effects as patently as logarithmic representations of numbers may, it may form some sort of reference point that can influence perceptions of relative increases or decreases in delay. It is also worth noting that in this ITC experiment, as is standard, delay is represented only in a numerical form (i.e. the delayed reward will be offered n days in the future, see Figure 3.2). This is in contrast to RISK and AMB experiments, where there are both numerical and graphical representations of uncertainty. Thus it may be reasonable to interpret the ubiquity of sequential effects in ITC relative to AMB as possibly also driven by the purely numerical presentation of uncertainty (suggesting that uncertainty could be coded on a similar noisy logarithmic scale) and/or the re-use of computationally expensive simulated futures. Conversely, we may interpret the relative lack of sequential effects driven by differences in ambiguity in the AMB task as partially a consequence of scale anchoring due to the graphical presentation of risk/ambiguity. More consequentially, the experiment design included 50% trial sequences that increased in ambiguity and 50% trial sequences that decreased in ambiguity. If sequential effects are a type of adaptation effect, they may be washed out by this experiment design (the only factor we explicitly manipulated to be 50-50). Finally, unlike with ITC, there are no clear theories as to what cognitive processes are invoked during the decision process, lending further credence to the potentially important role of graphical presentation – which itself can vary over experiments – of uncertainty.

We are further agnostic about the processes involved in propagating these sequential dependencies. Researchers studying serial dependence in visual perception argue that attention plays a critical role (Fischer & Whitney, 2014) in addition, possibly, to working memory (Kiyonaga et al., 2017). The relationship between attention and value-based decision making has been well studied (Krajbich, 2019), with some researchers arguing that attention plays a causal role in the formation of value. Further, researchers have also demonstrated evidence for “last fixation bias,” where the first fixation on a given trial is shaped by what the individual was last looking at on the previous trial, in value-based decision making (Krajbich et al., 2012). In related but distinct work, researchers have used eye tracking (the most commonly used proxy for measuring attention) to demonstrate individual differences in strategies for an ITC task (Khaw et al., 2018). In particular, Khaw and colleagues demonstrate not only that search strategy is a predictor of an individual’s discount rates, but also that search strategy can be shaped by tweaks in experimental design (Khaw et al., 2018). Taken together, attentional processes may indeed play a critical role in generating, or propagating, these experiment-design-dependent but task-incongruent spillover effects. Less studied, though equally plausible, is the relationship between working memory and economic decision making under uncertainty. Researchers have suggested that persistent activity in cortex, a typical signature of working memory, also supports value-based decision making (Curtis & Lee, 2010). As attention and working memory putatively operate on different time scales, it is entirely possible that both processes are involved in the phenomenon studied in this paper. Further research, and perhaps multi-modal data is necessary to be able to tease apart differential contributions of either process.

Finally, we return to the potential clinical implications of this research. Scientists have established for decades that there are exist meaningful differences between “healthy” controls and clinical populations in working memory and attention (e.g. aging (T. Salthouse, 1994; Commodari & Guarnera, 2008) and schizophrenia (Forbes et al., 2009; Luck & Gold, 2008)). We note that the degradation of working memory in aging appears to be more established,

and less selective, than that of attention and aging. This opens up an interesting line of research in both health and disease: how to study the presence of sequential effects, recoverable through the joint modeling of choice and RT, to make more nuanced subject level inferences about health status, especially in disorders that demonstrate degradations in the processes that putatively support sequential effects?

Overall, in this paper we have demonstrated that the near-ubiquitous assumption in modeling economic choice, that choices made in sequence can be treated as independently made due to experiment randomization, is false. We show that this is the case through the joint modeling of choice and RT and show that our model puts forth meaningfully different parameter estimates than standard choice models. This work takes an important initial step in quantifying the effect of a heretofore underexplored source of variability in the inference of risk tolerance, ambiguity tolerance, and discount factor.

Chapter 4

A response time model of memory discrimination.

4.1 Introduction

How do individuals encode objects in memory, and how does the distinctiveness of encoding affect behavioral expressions of recognition? These functions are thought to be supported by a process known as *pattern separation*, whereby similar sensory or latent input patterns are projected into higher-dimensional space to create highly distinct patterns that support later discrimination among fine degrees of difference (S. M. Stark et al., 2019). Traditionally, this process has been attributed to the hippocampus, a critical brain structure for learning and memory (Long et al., 2016; Marr, 1971; S. M. Stark et al., 2019). Computational models

We thank Dr. Craig E. Stark for providing data for Experiments 1 – 4 and Dr. Christopher Wahlheim for Experiment 5. We also thank Dr. Joachim Vandekerckhove, Dr. Michael D. Lee, and the Cognitive Modeling Lab for helpful discussions.

predict that the more distinct object representations are (i.e. the “better” an individual is at pattern separating), the better an individual will be able to discriminate between objects that were seen previously and those that weren’t. In particular, people who are better at pattern separating should be less susceptible to interference when novel objects are similar to the previously seen objects. The ability to create distinct representation is a necessary component in episodic memory – if the decisions we make are guided by our previous experiences, we need a mechanism in place to be able to distinguish these experiences. The hippocampus also performs a complementary process to pattern separation known as pattern completion: this is where incomplete or more gisty memories are “filled in.” It is necessarily the case that episodic memory requires both pattern completion and pattern separation, but this paper will focus on the latter.

The premier experimental task to capture a behavioral measure of pattern separation is the Mnemonic Similarity Task (MST) (S. M. Stark et al., 2019). The MST is a modified object-recognition task that is split into two distinct phases: study and test. During study, participants are given an incidental encoding task where they are presented with images of objects which they need to classify as belonging “indoors” or “outdoors.” Then, during test, participants are again presented with images. However, they are presented three different types of objects, each with the same frequency: *Repeats*, which are exactly the same objects they saw during study, *Lures*, which are objects similar to but not exactly the same as the study images, and *Foils*, which are objects that have never been seen before in the context of the experiment. The lure objects also vary on how similar they are to the object shown in the study phase, ranging from difficulty 1 (most) to 5 (least). There are multiple variants of the MST, varying from study design (two phase vs continuous), number of responses (3: Repeat/Lure/Foil or 2: Repeat/Foil), to stimulus sets. In this paper, we consider only the “standard” 3AFC two phase version of the MST. Subjects across experiments were not necessarily shown the same stimulus sets, but all stimulus sets were matched in their difficulty.

The primary measure of memory discriminability used in the MST is the Lure Discrimination Index (LDI) (S. M. Stark et al., 2019). Since the adoption of the MST as the premier task to assess behavioral pattern separation, the LDI has been parameterized in several ways. Today, the most common formulation used is:

$$LDI = P(\text{Lure Response} \mid \text{Lure Trial}) - P(\text{Lure Response} \mid \text{Foil Trial}) \quad (4.1)$$

Equation 4.1 can be thought of as the “hit rate” on lure trials corrected for the “bias” of incorrectly saying the foil is a lure (but correctly identifying the higher-order category of haven’t-seen-before-in-the-experiment). This LDI has been shown to vary with age and across a wide range of clinical measures (S. M. Stark et al., 2019). However, while there is rich evidence for the external validity of the LDI (S. M. Stark et al., 2019), there has also been extensive debate about whether the LDI is process pure, as much as any cognitive parameter can be. This in parts is related to the decades-old debate about the roles of recollection and familiarity in recognition memory (Yassa & Stark, 2011). Further, across the various parameterizations of the LDI, there is necessarily information loss. Any formulation of the LDI considers only one response type (out of three) and two trial types (out of three). If the key question of interest has to do with discrimination across various degrees of fidelity between old and new objects, it may make sense to use a measure that captures information about all three response types. Similar logic holds to related measures such as Recognition Score (which is like the LDI but instead focuses on the *Repeat* response hits and false alarms). In this paper, we propose a joint model of choice and response time that provides us with more psychologically nuanced parameters. This model further allows us to decompose the LDI in order to begin to consider – through behavioral and neural data – how process pure it may be.

4.2 Methods

4.2.1 Experiments and Data

We analyze data from 5 experiments collected by several researchers at different universities. We summarize them in the Table 4.1.

Experiments 1–4: Stark et al. The first set of experiments we analyze are from Stark and colleagues (C. E. Stark et al., 2023). In this paper, Stark and colleagues contrasted several variants of the MST to assess reliability and efficacy of measures. We consider a subset of these experiments that include the “full” or baseline version of the MST (number of test trials = 192, number of responses = 3). These experiments were collected across different individuals, and the only other point of variation were the stimulus sets used. However, all stimuli were matched for difficulty across each experiment. Here, we include analyses from 4 experiments in the paper.

Experiments 5: Wahlheim et al. We include data collected by Dr. Christopher Wahlheim and colleagues (Wahlheim et al., 2022). This is again the baseline version of the MST but consists of a lifespan sample (age 18 - 80, number of test trials = 108, number of responses = 3). Later, we consider the relationship between the LDI, LBA parameters, age, and resting state functional connectivity. All data was retrieved from the OSF repository provided by the authors: <https://osf.io/f6vg8/>

4.2.2 Response Time Model: The Linear Ballistic Accumulator

We adapt the Linear Ballistic Accumulator (LBA) (Brown & Heathcote, 2008) to model choices and response times in the MST. As with all sequential sampling models, the core process explained by the LBA is as follows: the stimulus is presented at the beginning of the

Experiment	Number of Subjects	Number of Trials	Source
E1	$n = 53$	192	Stark-1
E2	$n = 46$	192	Stark-2
E3	$n = 81$	192	Stark-3
E4	$n = 53$	192	Stark-4
E5	$n = 72$	108	Wahlheim-rsfc

Table 4.1: **MST experiments modeled.** The *Source* column indicates the researchers who originally collected the data. Experiment 5 has the additional suffix “rsfc” to indicate that the researchers also collected resting state functional connectivity measures. E5 also included a lifespan sample.

trial. Then, after some “non-decision-related” processing (e.g. identifying the trial image as a duck, or even towards the end of the trial after the decision has been made internally but needs to be executed through motor movement), individuals start to accumulate evidence – sampling both the trial image and in theory memory – till they have enough evidence to make a choice of either Repeat, Lure, or Foil. The LBA is a simple sequential sampling model that has the benefit of accommodating n -AFC experiments: we can fit n accumulators for n response types. It further assumes that evidence is independently and noiselessly accumulated for each response type.

There are four main parameters of the LBA, all common to most models of sequential sampling: the drift rate, or the rate of evidence accumulation/signal strength, the boundary b , or the amount of information needed for a response to be made, the non-decision time (NDT) τ , or the amount of time for non-decision-relevant processes, and the upper limit of the start point – the bias towards making a particular response.

For those more familiar with Signal Detection Theory (SDT), you may interpret the drift rate as similar to d' and starting point bias as similar to criterion. The difference in this case, is that instead of taking one sample from your representation (as is necessarily the case in SDT), response time models assume you take *multiple* samples [over time]. This is what gives rise and meaning to the other two parameters – boundary and NDT – which cannot be captured in SDT models.

As shown in Figure 4.1, we allow the drift rate and start point to vary per subject *and* per accumulator, while allowing the boundary and NDT to only vary across subjects. As the three responses in the MST are distinct, it stands to reason that the evidence accumulated in favor of each response should be different. Similarly, the bias or predisposition to making one response over another should also intuitively vary as a function of response type. Otherwise, a fixed bias might suggest that subjects have uniform tendencies to respond Repeat, Lure, and Foil. As we discuss in the Results section, this turns out not to be the case, with most subjects, across experiments, tending to respond Repeat disproportionately more often.

The LBA assumes that the drift rates are drawn from some Normal distribution and are sampled on each trial: Drift Rate \sim Normal($v_{i,r}$, $s_{i,r}$) for subject i and response type r . In this paper, when we report values associated with the drift rate, we are talking specifically about the *mean* drift rate ($v_{i,r}$). We fix the standard deviation of the drift rate ($s_{i,r}$) to be 1 across all subjects and accumulators for model identifiability purposes.

The LBA assumes that starting points are uniformly distributed, also sampled on each trial, and are numerically lower than the boundary: Start Point \sim Uniform[0, $A_{i,r}$] $T(0, B_i)$ for subject i and response type r .

To further keep the model identifiable, we impose the same constraint on the mean of the drift rates and start point upper boundaries for each subject: $\sum_{r=1}^3 v_{i,r} = 1$, $\sum_{r=1}^3 A_{i,r} = 1$. We fit a Bayesian implementation of the LBA in RStan (Team, 2023). The sum to one constraint is operationalized by allowing the drift rate mean and start point upper bound to be simplex types. We use the following relatively uninformative priors for the LBA parameters:

$$\begin{aligned} \text{Boundary} &\sim \text{Normal}(0.5, 1) & \text{NDT} &\sim \text{Normal}(0.5, 1) \\ \text{Drift Rate Mean} &\sim \text{Normal}(0.5, 0.5) & \text{Start Point Upper Bound} &\sim \text{Normal}(0.5, 0.5) \end{aligned}$$

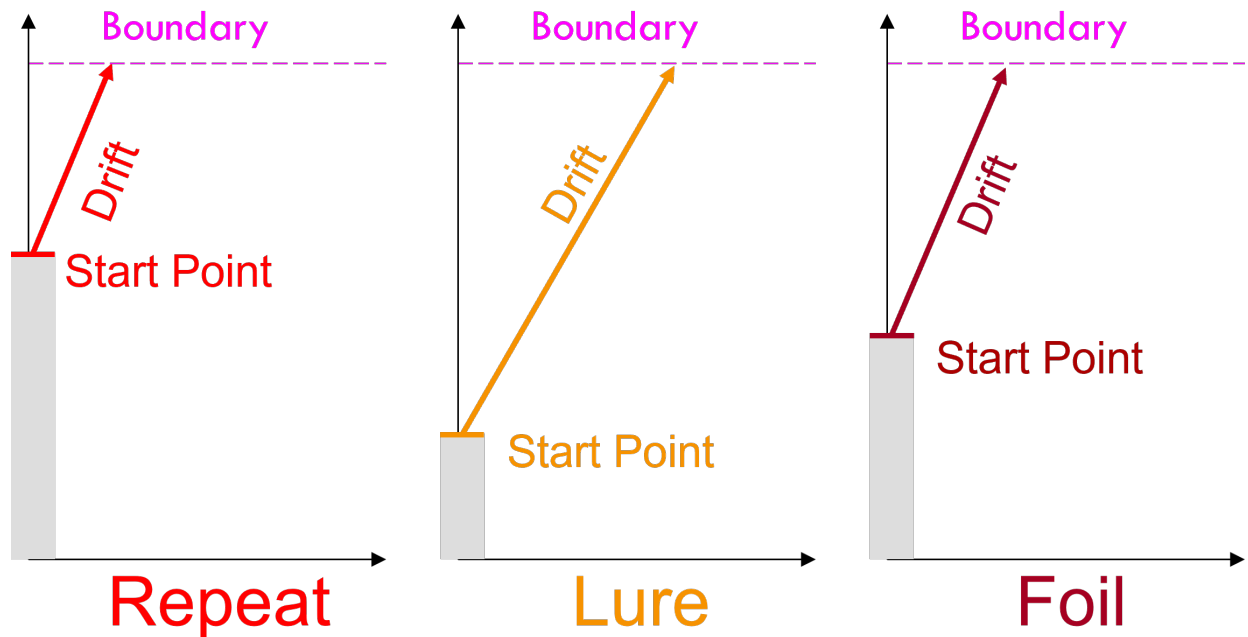


Figure 4.1: **Schematic of the Linear Ballistic Accumulator.** As the standard MST is 3AFC, we allow for three accumulators for each response type. We further allow the drift rate (rate of evidence accumulation) and the upper limit of the starting point (tendency to make a type of response) to vary for each subject and accumulator. Boundary (amount of evidence needed to make a response) and non-decision time (non-decision-relevant processes, not pictured) vary only at the subject level.

4.3 Results

4.3.1 Model Free.

We see that across accuracy, median RT, and response proportion, experiments are comparable (Tables 4.2, 4.3, 4.4). There is greater variability in median LDI, but none that is statistically significant (Table 4.5).

Another point of interest concerns the evolution of choice and response time over the course of the experiment. Previous research on memory and response time suggests that choices that are easier (more accessible in memory) should be faster, and choices that are more difficult should take longer. In the MST, the repeated stimuli are typically considered the

Experiment	Accuracy: Repeat	Accuracy: Lure	Accuracy: Foil
E1	0.86(0.18)	0.4(0.32)	0.77(0.3)
E2	0.87(0.09)	0.41(0.29)	0.78(0.18)
E3	0.84(0.15)	0.54(0.23)	0.80(0.13)
E4	0.84(0.18)	0.44(0.22)	0.82(0.13)
E5	0.92(0.12)	0.31(0.27)	0.86(0.16)

Table 4.2: **Accuracy across experiments by accumulator.** We show median(IQR) accuracy for each response type. Unsurprisingly across all experiments, subjects are more accurate when identifying repeat and foil stimuli vs. lure stimuli. However, lure stimuli accuracy is also the most variable.

Experiment	Med RT: Repeat	Med RT: Lure	Med RT: Foil
E1	1.17(0.28)	1.34(0.32)	1.31(0.28)
E2	1.21(0.22)	1.34(0.24)	1.26(0.23)
E3	1.25(0.23)	1.38(0.29)	1.35(0.31)
E4	1.72(0.36)	1.99(0.42)	1.80(0.36)
E5	1.05(0.21)	1.30(0.34)	1.11(0.26)

Table 4.3: **Median Response Time across experiments by accumulator.** We show median(IQR) RT for each response type. We find that, on average, participants tend to take longer when they make a Lure response especially when compared to when they make a Repeat response

Experiment	Prop: Repeat	Prop: Lure	Prop: Foil
E1	0.44(0.19)	0.26(0.13)	0.30(0.09)
E2	0.47(0.12)	0.23(0.12)	0.30(0.06)
E3	0.41(0.12)	0.28(0.12)	0.31(0.06)
E4	0.43(0.11)	0.24(0.09)	0.33(0.05)
E5	0.51(0.11)	0.19(0.13)	0.32(0.06)

Table 4.4: **Choice Proportions across experiments by accumulator.** We show median(IQR) choice proportions for each response type. We see that across experiments, participants to most often classify stimuli as Repeat (recall that the true proportion of repeated stimuli presented during test is 0.33).

Experiment	Lure Discrimination Index
E1	0.18(0.40)
E2	0.21(0.33)
E3	0.34(0.28)
E4	0.28(0.25)
E5	0.17(0.30)

Table 4.5: **Lure Discrimination Indices for each experiment.** We show median(IQR) LDI, which is the current standard metric for summarizing choice behavior in this experiment. We find no statistical differences in median LDI across experiments due to how much the LDI varies within each experiment.

easiest to identify (indeed these are the stimuli that have already been seen before in the context of the experiment). Indeed, as we demonstrate in Figure 4.2, this appears to be the case, though there is variability across experiments. When we collapse across Experiments 1 – 4, with the exception of the very fastest responses, we find that subjects tended to label a stimulus as a Repeat most often when making a quick decision ($RT < 1.38s$). For slower decisions ($RT > 1.38s$), the highest frequency response was Foil. However, in Experiment 5, we found that the majority of responses over time were Repeats (with the exception of 4 later RT bins out of 15 ($RT > 1.29s$) which had a majority response of Foil), Figure 4.2. Interestingly, and perhaps unsurprising given the low frequency with which people made Lure responses (Table 4.4), there is no RT bin where Lures are the most often chosen response type.

4.3.2 Model Based.

All results reported are from models that pass all metrics of convergence. We note that while we do not use a hierarchical implementation of the LBA, we find via descriptive adequacy (posterior predictive visual and summary statistic) checks that the model successfully captures individual differences.

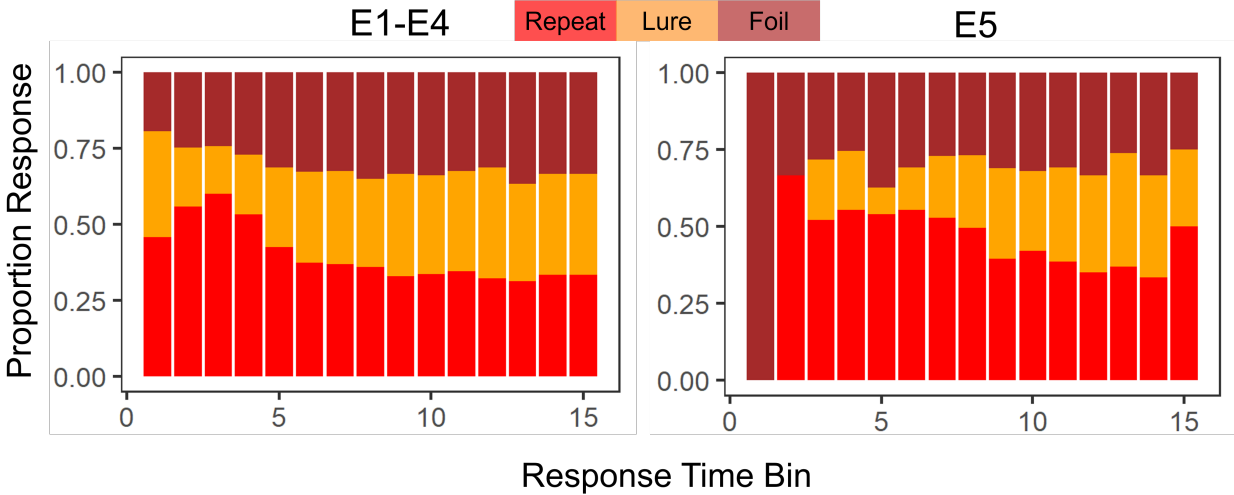


Figure 4.2: **Choice proportions change as a function of response time.** We find, in general, that faster choices tend to be *Repeats*.

Posterior Summaries

We find that the LBA parameters across the the experiments show generally the same patterns: start point upper limits and drift rates for Lure responses tend to be the lowest, whereas start point upper limits and drift rates for Repeat responses tends to be the highest, Figure 4.3. This is consistent with our expectations given the patterns observable in the response times themselves. We also believe that this suggests reasonable recovery of information by our model: while there can be great heterogeneity across individuals, with the exception of E5, which contains a subset of older adults, there is no reason to expect qualitative differences in model fits across experiments. After all, the experiment itself (again except for E6) also does not change. This is particularly of interest given how much variability there is in the LDI across experiments (see Table 4.5).

Relating the LDI to LBA parameters

A key goal of this work is to try to compare the relationships between the LDI and our model parameters. Of particular interest is the relationship between LDI and drift rate,

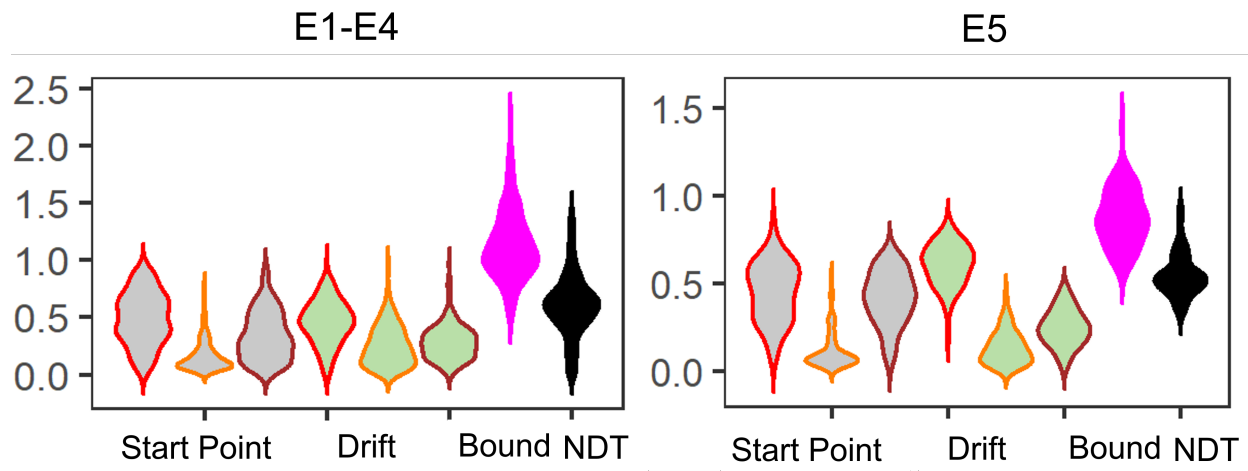


Figure 4.3: **Model posteriors for Experiments 1-3, Experiment 5** We find overall that LBA posteriors follow qualitatively the same patterns across experiments. The left most three grey violins are the start point upper bound (color coded by response type: Repeat, Lure, Foil). Next are the drift rates, similarly color coded. Finally, we have the boundary and non-decision time.

and LDI and start point upper bound. This relates to the question briefly considered in the introduction: is the LDI capturing a signal of recognition memory? How much of it is conflated by other processes? If the LDI correlates *only* with the drift rate, which is the LBA’s measure of signal strength, it suggests that the LDI may indeed be largely a measure of how distinct people’s internal representations are. The more distinct the internal representation, the stronger the internal signal during evidence accumulation. Conversely, if the LDI correlates *only* with the start point upper bound, which is the LBA’s measure of a tendency to make a particular response, it suggests that the LDI may be largely capturing something else, perhaps closer to familiarity or other more “gisty” concepts. However, it is rare for any one cognitive process to work in isolation, and indeed we find that the LDI correlates with *both* the accumulator drift rates and the accumulator starting point upper bounds in Experiments 1 – 4 (Figure 4.4). We collapse data across all four experiments as they were not designed to be meaningfully different across from each other. This results in a total sample size of 233 subjects.

Across experiments, we found strong statistically significant ($p < 0.05$) correlations between

the LDI and drift rates for each accumulator (Kendall’s tau: $\text{corr}(\text{LDI}, \text{Repeat drift rate}) = -0.35$, $\text{corr}(\text{LDI}, \text{Lure drift rate}) = 0.29$, $\text{corr}(\text{LDI}, \text{Foil drift rate}) = 0.19$). Similarly, we found statistically significant ($p < 0.05$) correlations between the LDI and start point upper bounds for the Repeat and Foil accumulators (Kendall’s tau: $\text{corr}(\text{LDI}, \text{Repeat start point upper bound}) = -0.11$, $\text{corr}(\text{LDI}, \text{Foil start point upper bound}) = 0.16$). We corrected for multiple comparisons using the Bonferroni-Holm procedure. We note that the sign differences in the correlations between Repeat vs. Lure and Foil are as expected. The LDI is explicitly calculating a “signal” of how well an individual discriminates between items that haven’t in their totality been seen before in the context of the experiment. This is the complementary process to recognizing old items. Finally, we find a positive correlation between LDI and Non-Decision Time ($\tau = 0.17$): the better the memory discrimination, the longer the non-decision relevant processing.

The degree to which these correlational relationships hold across experiments is, however, variable. In E5, we find that the Repeat and Lure accumulator drift rates correlate significantly in the same way as in Figure 4.4, but that there is no significant linear relationship between the LDI and Foil accumulator drift rate, Figure 4.5. Conversely, we find exactly the same statistical patterns in the relationship between LDI and the starting point upper bounds. We also find no relationship between NDT and LDI.

Comparing Correlation Strengths. To formally compare correlation strengths, we used bootstrapping to resample the data and calculate Kendall’s τ s and the differences between each pair of τ s (e.g. $\tau_A - \tau_b$). In particular, we wanted to test whether the correlations between the LDI and Drift Rates were stronger than the correlations between the LDI and Start Points. We then examined whether the bootstrapped 95% confidence interval distributions of the differences between each pair of correlations included zero. If they did not include zero, we interpreted this as evidence as a rejection of the null (no difference between the correlations).

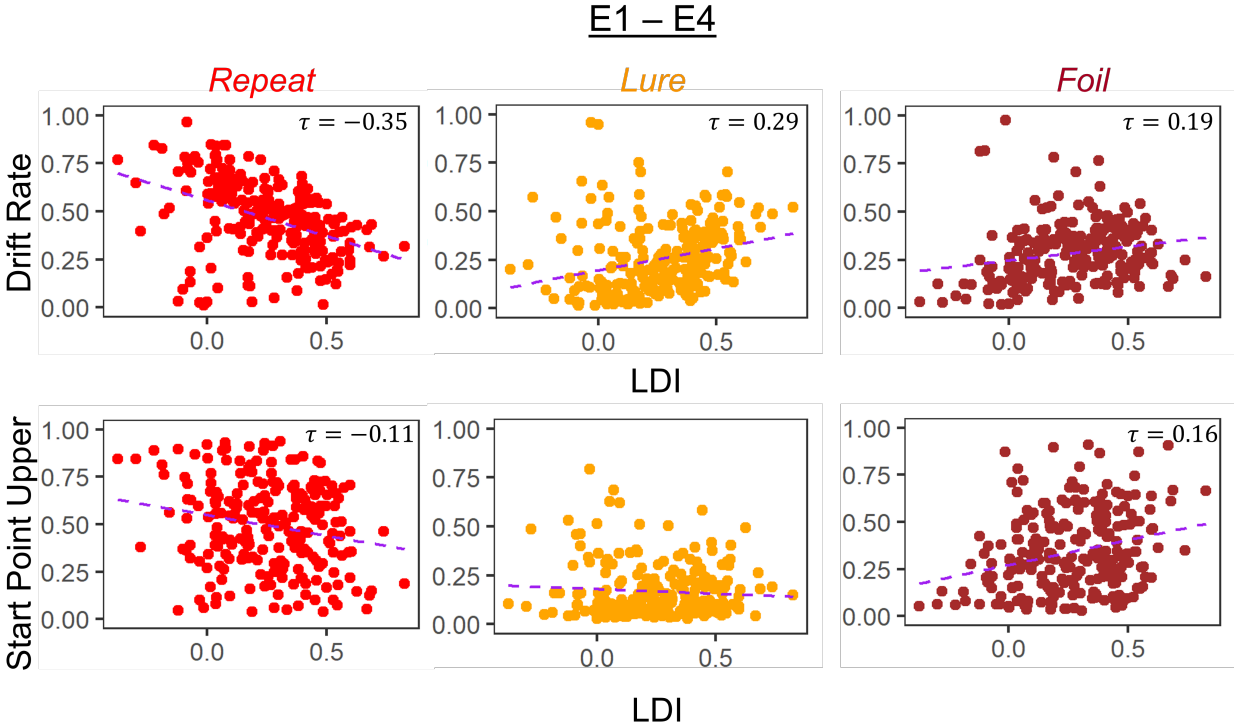


Figure 4.4: **LDI correlates with drift rate *and* start point upper bounds in E:1-4.** We collapse across all 4 experiments ($n = 233$) and correlate mean drift rate and start point upper bound with LDI. We find statistically significant correlations between the LDI and drift rates for all accumulators. We also find significant correlations between the LDI and Repeat Foil accumulator start point upper bounds. Correlations shown in plots are statistically significant ($p < 0.01$) after adjusting for multiple comparisons.

For Experiments 1 – 4, we found that the 95% CIs for the correlation difference between LDI-Drift Rate and LDI-Start Point for the Repeat and Lure accumulators did not contain 0 (Repeat accumulator correlation difference (0.13, 0.35), Lure accumulator correlation difference (0.16, 0.36), Figure 4.6). However, this was not the case for the Foil accumulator (Foil accumulator correlation difference (-0.10, 0.15)). For Experiment 5, however, we found that none of the 95% CIs excluded zero (Repeat accumulator correlation difference (-0.29, 0.21), Lure accumulator correlation difference (-0.05, 0.37), Foil accumulator correlation difference (-0.39, 0.09), Figure 4.6). Formally, this suggests that we cannot reject the null of no difference in correlation strength between LDI and respective accumulator drift rate/start point upper boundary, however, we point out that for the Lure and Repeat accumulators, the

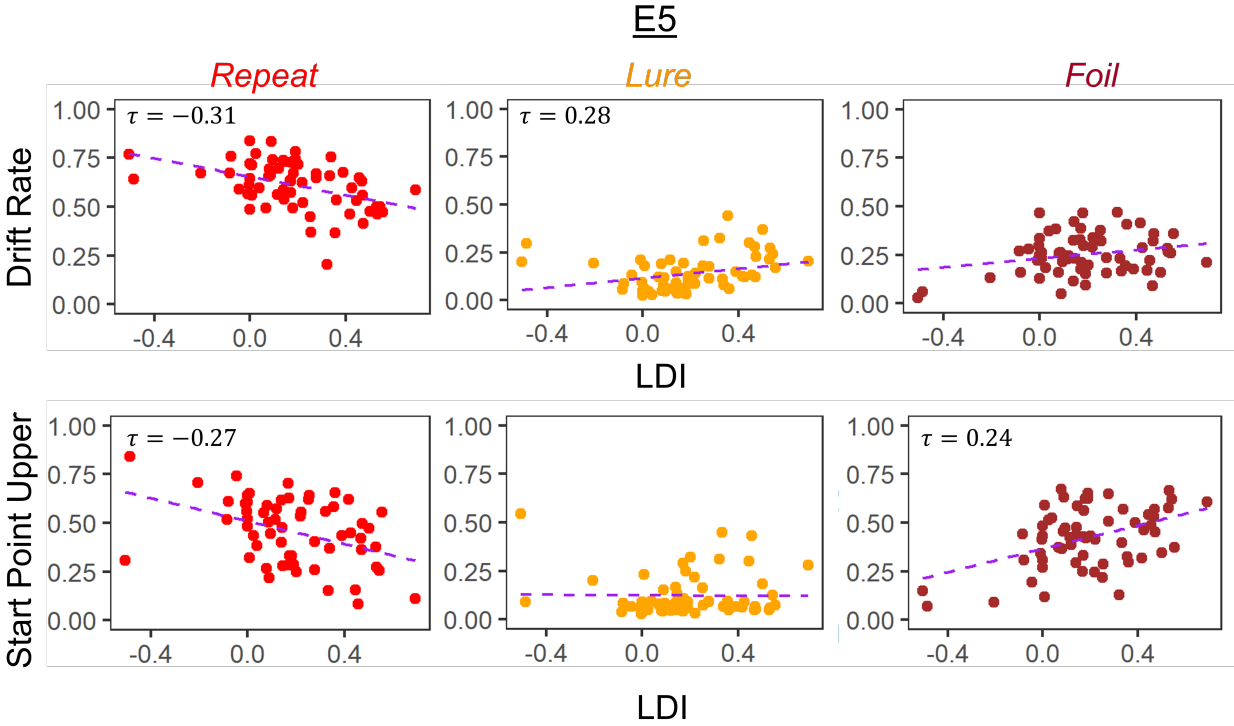


Figure 4.5: **LDI correlates with drift rate *and* start point upper bounds in lifespan sample.** In a dataset comprised of older and younger adults, we find similar qualitative relationships between the LDI and accumulator drift rates/start point upper bounds. Correlations shown in plots are statistically significant ($p < 0.01$) after adjusting for multiple comparisons.

95% CIs only just include 0. Overall, this analysis suggests that, particularly for the Lure response type, the LDI seems to be more of a measure of signal strength than response bias.

The LDI, LBA Parameters, and Other Variables

Another way to consider how valuable our model is to consider measures outside the LDI.

In experiment 5, which was collected in conjunction with resting state functional connectivity measures, we consider two primary questions: (1) is there a differential relationship between age and LDI/LBA parameters and (2) what is the relationship between LDI/LBA parameters and resting state functional connectivity in the hippocampus.

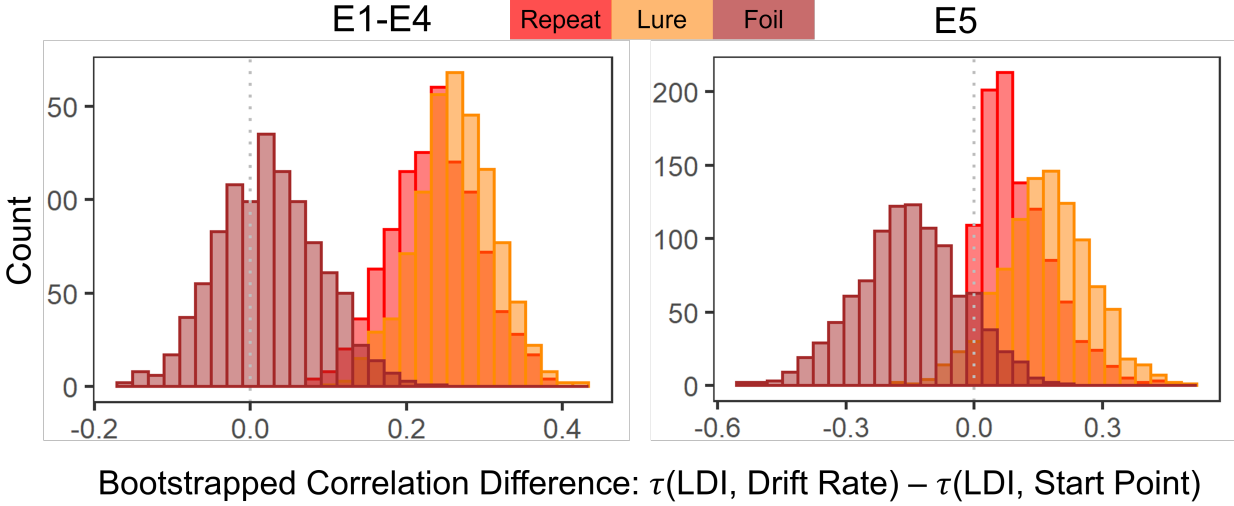


Figure 4.6: **E1-E4: LDI-Drift Rate correlations for Repeat and Lure accumulators are stronger than LDI-Start Point correlations. E5: No significant correlation differences.** We plot histograms showing bootstrapped correlation differences between LDI and drift rate, and LDI and start point (left: Stark 1-4, right: E5: Wahlheim et al). Repeat and Lure accumulator 95% CIs in E1 – 4 do not contain 0, suggesting we can reject the null of no correlation difference.

Age. The “young” age group had a median age of 21 ($IQR = 3$) and a sample size of 34. The “old” age group had a median age of 69.5 ($IQR = 8.5$) and a sample size of 28. Younger adults had a median accuracy of 0.69($IQR = 0.11$) compared to 0.66($IQR = 0.08$). Younger adults were also faster than older adults (median RT young = 1.03(0.16) vs median RT old = 1.19(0.16); $W = 779.5$, $p < 0.01$ – we report the results of a Wilcoxon Rank Sum Test due to non-normality in the distribution of untransformed RTs.).

Consistent with what the authors report in their original paper (Wahlheim et al., 2022), we find significant differences in lure discrimination as a function of age group. Younger adults have a significantly higher LDI than older adults (median LDI young = 0.27(0.26) vs median LDI old = 0.09(0.18); $W = 693.5$, $p < 0.01$).

For the LBA parameters, we only find statistically significant differences in non-decision time, with older adults taking longer with non-decision-related processes than younger adults (median NDT young = 0.49(0.08) vs median NDT old = 0.60(0.16); $W = 794$, $p < 0.01$).

As we show in our earlier analysis, the LDI is meaningfully correlated with both drift rates and start point upper boundaries. Therefore, while the LDI is different as a function of age group, it is not necessarily surprising that the relevant LBA parameters do not differ across age groups: they are *components* of the LDI.

Interestingly, when we compare correlations between the LDI and LBA parameters as a function of age group, no statistically significant correlations remain after adjusting for multiple comparisons, Table 4.6, Table 4.7. However we note that several correlations are trending, in particular the LDI and Repeat accumulator drift rate in older adults.

Parameters Correlated	Correlation	p-Value
LDI-Drift <i>Repeat</i>	-0.26	0.03
LDI-Drift <i>Lure</i>	0.28	0.02
LDI-Drift <i>Foil</i>	0.07	0.57
LDI-Start Point <i>Repeat</i>	-0.18	0.14
LDI-Start Point <i>Lure</i>	-0.04	0.72
LDI-Start Point <i>Foil</i>	0.13	0.27
LDI-Boundary	-0.02	0.86
LDI-NDT	0.05	0.68

Table 4.6: **Younger adult LDI-LBA correlations follow the same qualitative patterns as in E1-E4.** None of the correlations are statistically significant, however, after adjusting for multiple comparisons.

Parameters Correlated	Correlation	p-Value
LDI-Drift <i>Repeat</i>	-0.36	0.0076
LDI-Drift <i>Lure</i>	0.16	0.24
LDI-Drift <i>Foil</i>	0.20	0.14
LDI-Start Point <i>Repeat</i>	-0.22	0.10
LDI-Start Point <i>Lure</i>	0.11	0.40
LDI-Start Point <i>Foil</i>	0.20	0.13
LDI-Boundary	-0.06	0.68
LDI-NDT	0.34	0.01

Table 4.7: **Older adult LDI-LBA correlations follow the same qualitative patterns as in E1-E4.** None of the correlations are statistically significant, however, after adjusting for multiple comparisons.

Hippocampal Resting State Functional Connectivity. We then consider the more explicit question of external validity. In their original paper, Wahlheim and colleagues demonstrated that resting state functional connectivity (RSFC) in the Default Mode Network predicts LDI (Wahlheim et al., 2022). Here, we examine the relationship between resting state functional connectivity in the hippocampus (8 regions: left/right medial head, left/right lateral head, left/right body, left/right tail) and the LDI/LBA parameters. We correlate LDI and LBA parameters with a matrix of Fisher’s z transformed correlation coefficients of hippocampal connectivity (procured from OSF and preprocessed by (Wahlheim et al., 2022)).

As the MST is designed to capture a process attributed to the hippocampus, and we see a behavioral relationship between the LDI and some LBA parameters, we expected both the LDI and LBA parameters to correlate with hippocampal RSFC.

Recall from the previous section that in this dataset (not accounting for the differently aged subgroups), we found that LDI correlated with both mean drift rate (Repeat, Lure)

and start point upper bound (Repeat, Foil). We therefore wanted to examine whether similar patterns held with resting state function connectivity: how does LDI correlate with hippocampal RSFC, how do LBA parameters correlate with hippocampal RSFC, and what is the intersection and nodes of divergence between the two?

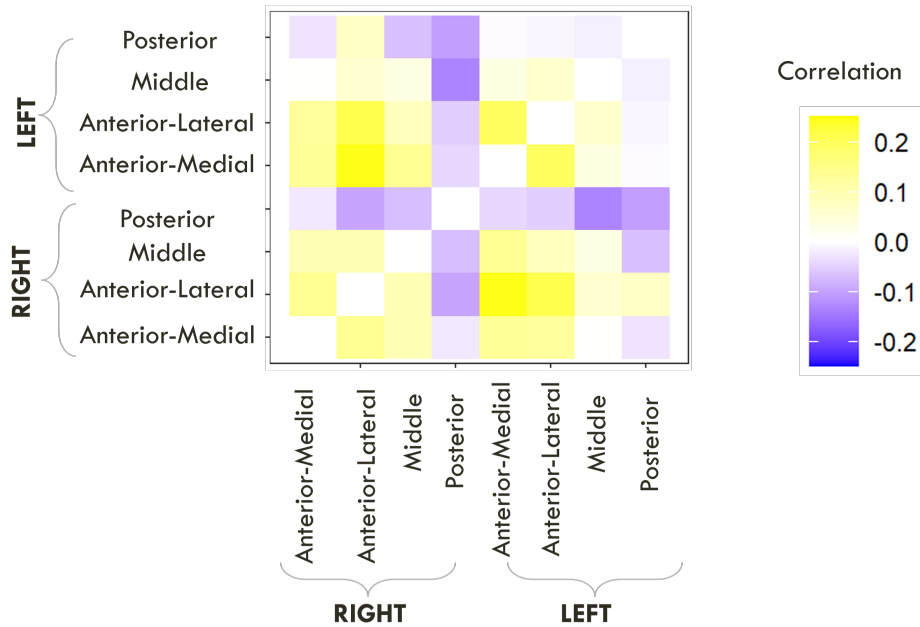


Figure 4.7: **E5: LDI tends to correlate negatively with posterior hippocampal RSFC and positively with other regions.** We show a symmetric correlation matrix plot where each square represents the correlation between the row-column hippocampal resting state functional connectivity and the LDI. Cells colored darker yellow show stronger positive correlations, and darker purple stronger negative correlations.

We first consider raw correlations between hippocampal RSFC and our behavioral parameters of interest. In Figure 4.7, we find that the LDI correlates negatively with posterior hippocampal RSFC and positively with other regions, perhaps in line with the representational specificity gradient within the hippocampus. In Figure 4.8, we show that Repeat and Lure accumulator drift rates differentially correlate with hippocampal RSFC: with Repeat drifts correlating negatively and Lure drifts correlating mostly positively. Finally, in Figure 4.9, we find less clear directional patterns between accumulator start points and hippocampal RSFC. Further, absolute values of the correlations suggest that start point – RSFC correlations may be weaker than drift rate – RSFC.

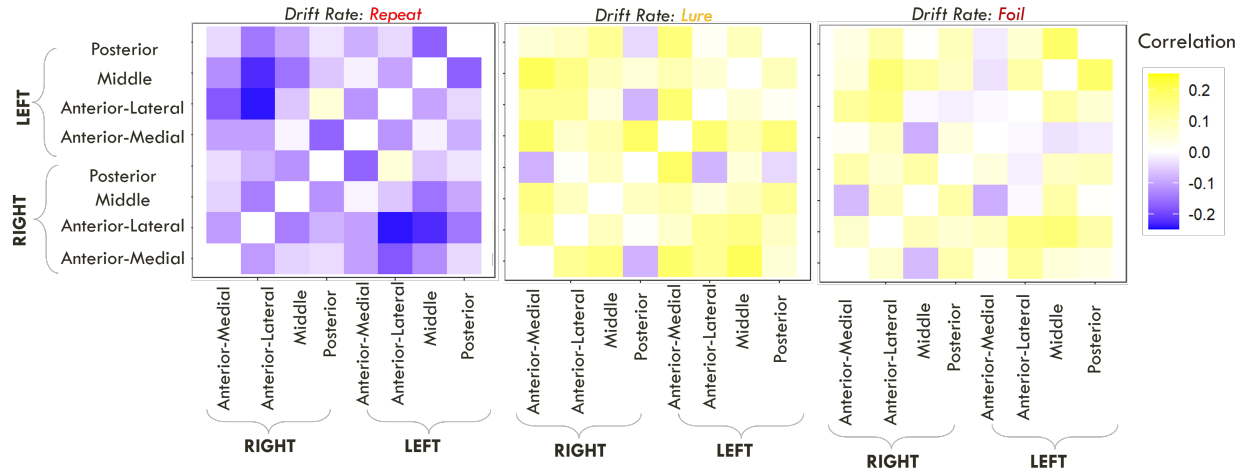


Figure 4.8: **E5: Lure accumulator drift rates mostly correlate positively with hippocampal RSFC while Repeat accumulator drift rates correlate negatively.** We show a symmetric correlation matrix plot where each square represents the correlation between the row-column hippocampal resting state functional connectivity and the respective accumulator drift rate. The Repeat and Lure accumulator – RSFC correlations also seem to be stronger than the Foil accumulator – RSFC correlations.

When we threshold correlations, we find that the LDI, drift rates for all accumulators, upper boundary for the Repeat accumulator start point, and non-decision time all show statistically significant ($p < 0.05$) Kendall τ correlations with various RSFC regions in the hippocampus (Table 4.8). Of particular interest is that *only* the drift rate for the Repeat and Lure accumulator correlated with the same RSFC regions as the LDI (and in directions consistent with previous independent datasets, Experiments 1 – 4: negatively correlated with Repeat accumulator and positively correlated with Lure). We further highlight that each parameter correlated with hippocampal RSFC correlates *in the same direction* – for example, Repeat accumulator drift rates always negatively correlated with RSFC and but Lure accumulator drift rates are all positively correlated. This suggests that the correlations we recover are not necessarily spurious (in which case we may expect variability in the directionality of the correlations – contrast for example, Figure 4.8 vs. Figure 4.9). The thresholded analysis, then, may further suggest that the LDI, while behaviorally correlated with both our measures of signal strength and response bias, may indeed relatively “process

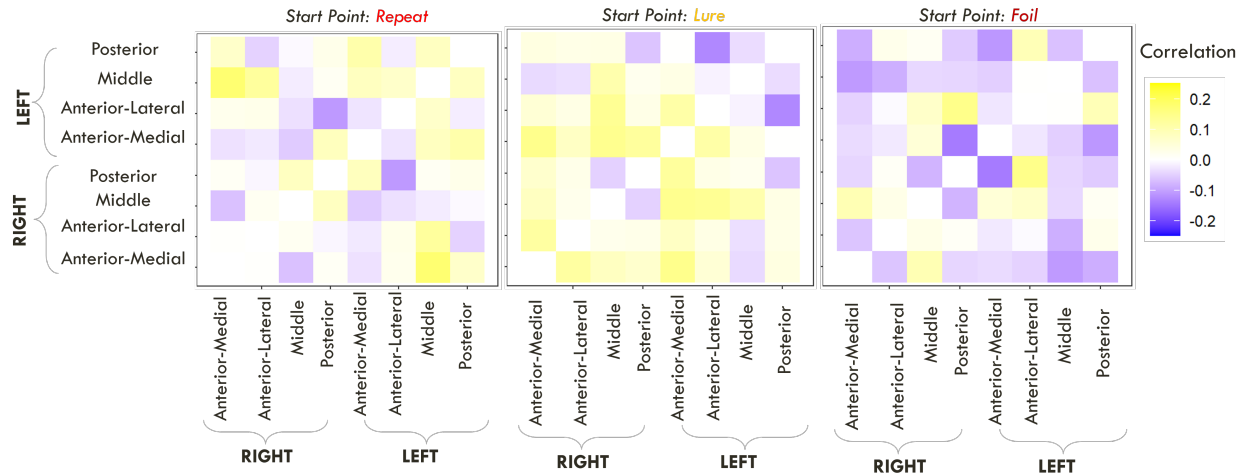


Figure 4.9: **E5: Start Point** – hippocampal RSFC correlations appear to be weaker than Drift Rate – hippocampal RSFC correlations. We show a symmetric correlation matrix plot where each square represents the correlation between the row-column hippocampal resting state functional connectivity and the respective accumulator drift rate. We also see greater variability in the qualitative patterns in the start point – RSFC correlations (i.e. evidence of both negative and positive correlations).

pure” (i.e. signal strength vs. response bias).

4.4 Discussion

In this chapter, we have introduced a way to model choice and response time in the Mnemonic Similarity Task. We adapt a version of the Linear Ballistic Accumulator to model each response distinctly: Repeat, Lure, and Foil. A primary contribution of this work is to introduce psychologically interpretable parameters, allowing us to separate signal strength (i.e. drift rate) from other processes (e.g. response bias).

We also demonstrate that the LBA parameters relate systematically with the previous gold standard choice-based measure, the LDI. Specifically, that the LDI correlates *both* signal strength and response bias: it is capturing facets of recognition memory, familiarity, and other adaptive behaviors that evolve over the course of an experiment. The drift rate is a

Parameter	Regions	Correlation
LDI	Anterior lateral LH - Anterior lateral RH	$\tau = 0.21$
	Anterior medial LH - Anterior lateral LH	$\tau = 0.20$
Drift: Repeat	Anterior lateral LH - Anterior lateral RH	$\tau = -0.24$
	Anterior lateral RH - Middle LH	$\tau = -0.23$
	Middle LH - Posterior LH	$\tau = -0.17$
Drift: Lure	Anterior medial LH - Anterior lateral LH	$\tau = 0.18$
	Anterior medial RH - Middle LH	$\tau = 0.20$
Drift: Foil	Anterior lateral RH - Middle LH	$\tau = 0.17$
	Middle LH - Posterior LH	$\tau = 0.18$
Start Point: Repeat	Anterior medial RH - Middle LH	$\tau = 0.18$
Start Point: Lure	-	-
Start Point: Foil	-	-
Boundary	-	-
Non-decision Time	Anterior medial LH - Anterior lateral LH	$\tau = -0.19$
	Anterior lateral LH - Anterior lateral RH	$\tau = -0.21$

Table 4.8: **Significant Hippocampus resting state functional connectivity and LDI/LBA parameters: *Only drift rates correlate with same RSFC regions as LDI.*** Each row delineates which parameter correlated significantly ($p < 0.05$) with RSFC between the two listed hippocampus subregions. LH and RH are shortform for Left Hemisphere and Right Hemisphere respectively.

measure of signal strength, or deliberation. In this paper we use it as a proxy for pattern separation as after all, if the signal being captured by the MST is a behavioral measure of pattern separation, this should be what individuals deliberate over. This may suggest, therefore, that the LDI does indeed index some measure of pattern separation (via signal strength as recovered by the accumulator drift rates) but is also capturing other processes that naturally occur during the decision making process. Importantly, while we found variability across experiments in which accumulator drift rates and start point upper bounds correlated with the LDI, our secondary analysis quantifying the difference in correlation strengths showed that the LDI may indeed be capturing more “signal” than response bias.

To further explore what insights our modeling approach could provide, we considered a dataset that measured several other measures of individual differences – in particular age (known to correlate with LDI (S. M. Stark et al., 2019)) and hippocampal resting state functional connectivity. Across age groups, while we found significant differences in LDI,

we found no differences between the putative LDI decompositions drift rate and start point upper bounds. With hippocampal RSFC, we found that, of all parameters we model, the accumulator drift rates correlate the most with hippocampal RSFC. Critically, we found that *only* the drift rates correlate with the same connectivity regions as the LDI: all in the anterior hippocampus. While the hippocampus has several functions and there may be other brain regions whose functional connectivity may correlate with both the LDI and start point upper bound, we find these results to be an encouraging step towards addressing how “process pure” the LDI may be.

Our findings may enhance the application of MST in several ways. First, the use of sequential sampling models can allow researchers to extract trial-by-trial timeseries reflecting putative underlying computations that drive behavior, which should support analysis of more precisely defined functional neuroimaging measures (Long et al., 2016). Secondly, the robust statistical frameworks often used to fit these sorts of models may allow further refinement of the approach, producing even more stable trait-level estimates by, e.g., incorporating informative priors and models of contaminant behavior, and integrating trial-wise neural measures to simultaneously test mechanistic hypotheses and improve model fit to behavior (Turner et al., 2019). Finally, we draw general attention to how response times can provide meaningful information about an individual’s memory discrimination – regardless of whether RT is explicitly modeled and perhaps especially when considering vulnerable or clinical populations.

Future work involves comparing the stability of LBA measures vs. the LDI – could LBA parameters be used to more finely predict the same sorts of outcomes currently predicted by LDI? We also hope to examine resting state functional connectivity in other brain regions in order to make more clear statements about how process pure the LDI is, and how clearly the LBA drift rate measures a signal of pattern separation.

Chapter 5

Putting it together: Motivating sequential effects in older adults and preliminary analyses in economic choice and memory discrimination.

5.1 Introduction

Background and Motivation

In Chapters 2 and 3 of this thesis, and gestured at in Chapter 4, is the notion of resource constraints. In Chapter 2, we motivate the idea of sequential effects even in “non-sequential” environments by introducing theoretical frameworks – neuroscientific and behavioral – that argue that processing new information in a relative fashion is an adaptive feature of humans (and other animals). Relativity is adaptive because of our finite capacity to store, process,

and retrieve information (or experience, more broadly).

In Chapter 3, we empirically demonstrate the importance of considering relative differences in trial properties while modeling economic decision making under uncertainty: indeed, we are only able to recover sequential effects when we code properties relatively (i.e. is the current trial ambiguous *more* than the previous and *not* how ambiguous was the previous trial?) We speculate that it is, due to resource constraints, *efficient* to reuse simulations of the past or reference recently made calculations (and indeed that this reuse may be a consequence of how we encode numbers in the first place). In this chapter, we instead turn more explicitly to a consideration of resource constraints.

It is well established that as humans age, their various capacities – particularly those dependent on the prefrontal cortex – diminish (T. A. Salthouse, 2009). This – outside of pathology – is often explained by the natural decay of the brain and the consequent notion that older adults have fewer resources than younger adults. Older adults (OAs) are therefore unable to process and store information the way they may previously have been able to. A key behavior theoretically associated with such reduction in capacity is perseveration: goal-incongruent repetition (internally or as directly observed in behavior). We have previously argued that the sequential dependencies we empirically demonstrate in Chapter 3 are also a type of perseveration. However, as we use the term, we extend the notion of perseveration to mean a goal-incongruent reference to the past, not just a repetition, as it is typically defined.

As we touch on in a paragraph in Chapter 2, perseverative behaviors can manifest differently across various stages of development and normal cognitive aging and age-related cognitive decline. Perseverative behaviors can be broadly observed across content (semantic, perceptual, response), age/disorder (pathology can play a causal role in how perseveration manifests – individuals with Parkinsons, for example, tend to exhibit more response perseveration), temporal profile, and other forms (e.g. emotional).

We briefly expand on the temporal profile of perseverations. A seminal paper by Sandson and Albert highlighted three different temporal profiles of perseveration: stuck-in-set, recurrent, and continuous perseveration (Sandson & Albert, 1984). *Stuck-in-set* perseveration refers to when individuals deploy previously appropriate rules. For example, a child that focuses on a previously rewarding but now not rewarding stimulus behavior to guide choice behavior, or a patient when previously drawing a human face, began to incorporate human features onto the current drawing of a cat (Sandson & Albert, 1984). *Recurrent* perseveration is defined as the repetition of responses: from saying the same word twice in a row during free recall, to repeating the same word over a larger temporal lag (Fischer-Baum et al., 2016). Finally, *continuous* perseveration, the most “extreme” is typically only present in pathological cases: the unbroken repetition of some action (e.g. drawing increasingly numerous petals of a flower (Sandson & Albert, 1984).

Several studies have examined how different types of perseverative behaviors change over time, providing insights into the underlying cognitive and neural mechanisms. Interestingly, perseverative errors in “healthy aging” tend to follow a U shaped trajectory across the lifespan: initially present in early childhood (as early as 8 months old (Carroll et al., 2016), stabilizing to a minimal level over young adulthood, to finally presenting again in older adults (Foldi et al., 2003). In older adults, this pattern has also been formalized in dementia patients as “last in – first out,” where processes/neural circuits that develop later are more vulnerable to neurodegeneration (Scherder et al., 2011).

We note that perseveration in young children is one of the rare instances in extant literature where it can be considered to be a sign of growth: in some experiments children of very young ages (5 month old, 2 years old) demonstrated random responding, slightly older children perseverated (8 month old, 3 years old), and even older children overcame perseveration (12 months old, 4 years old) (Carroll et al., 2016). Thus, perseveration here demonstrates the ability to *maintain* some previously task-optimal or relevant information: be it as simple as

grasping in the previously correct direction (recurrent perseveration) or as complicated as learning a rule (stuck in set perseveration).

Normal cognitive aging is associated with certain declines in cognitive functions, such as processing speed, working memory, and executive functioning (T. A. Salthouse, 2009). These declines can lead to an increase in perseverative behaviors, particularly in tasks that require cognitive flexibility and inhibitory control (Andrés et al., 2008; Hasher & Zacks, 1988). For example, older adults may exhibit increased difficulty in switching between tasks or inhibiting irrelevant information, which can result in perseveration (May et al., 1999). Fascinatingly, extant research suggests that the reverse pattern of perseveration that exists in early development appears to hold as people age/undergo increasing neurogeneration. At some point in an adult's lifespan, individuals begin to perseverate in a more prevalent fashion than they previously did (i.e. they are unable to overcome perseveration as they previously may have been able to). Older adults without impairments are demonstrated to show more response (recurrent) perseveration than their younger counterparts (Foldi et al., 2003). Individuals with individuals with the most severe stages of Alzheimers tend to respond randomly (Westfall & Lee, 2021).

In summary, the literature suggests that different types of perseverative behaviors can change with normal development and cognitive aging and age-related cognitive decline. While both development and normal cognitive aging may lead to a general increase in perseveration due to the establishment of / declines in cognitive flexibility and inhibitory control respectively, age-related cognitive decline can result in more pronounced and diverse perseverative behaviors, reflecting the deterioration of various cognitive domains and neural systems.

Following the notion of “last in – first out,” we propose that the sequential effects we motivate and demonstrate in earlier chapters of the dissertation belong in this perseveration hierarchy, as a “last in” behavior. While the relative encoding of new content is itself a fundamental part of how the brain processes information, the degree to which people show

these task-incongruent sequential effects in complex economic choices may be a useful marker of neurological health.

Hypotheses

For the analyses in this chapter, therefore, we wanted to examine the presence/absence of sequential effects in older adults. To be consistent with the literature reviewed above, we hypothesized that sequential effects – themselves a type of perseveration – may be more ubiquitously present in older adults than already present in younger adults. By ubiquitously, we speculate not merely on the *proportion* of subjects that show strong enough evidence in favor of sequential effects (e.g. Table 3.3, Chapter 3) which is already high in young adults. Instead our hypothesis focus on the *prevalence* of sequential effects across the terms we test. In the Risk and Ambiguity task in Chapter 3, all recovered sequential effects were only on the drift rate decomposition terms. This suggests that the rate of evidence accumulation, and not the bias/tendency to respond one way or another (which itself can be an inherent tendency that could evolve over the course of the experiment) is the primary element of the decision process that may be affected by sequential effects. Further, the recovered sequential effects were also *only* on the drift rate terms capturing the Expected Value difference between choice options and the Ambiguity level. We speculate therefore that, in the aging population, the distribution of sequential effects will be different: *all* computations should be more intensive (i.e. sequential effects will not necessarily be restricted to the two drift rate components) in addition to an increased tendency to perseverate (i.e. we may even expect sequential effects in the bias term).

We also hypothesized that there may be a relationship between the prevalence of perseveration and performance on the mnemonic similarity task (MST). The MST, introduced in the previous chapter, provides a behavioral measure of the distinctiveness of people’s internal

representations. This could therefore be a candidate mechanism of sequential effects in aging. We predict that poorer performance on the MST (as indexed by the choice only measure – by our response time model introduced in the previous chapter and the standard behavioral measure *LDI*) will track with sequential effects in the ambiguity task. More concretely, we hypothesize that people who perform worse on the MST have an greater tendency to classify things as more similar to each other than they actually are: successive trials with larger difference in value may in effect perceive the value difference to be smaller than it is. Indeed, we might expect that performance on the MST is negatively correlated with the presence and/or prevalence of sequential effects on the *bias* term in the Ambiguity Task.

In this chapter, we deploy the two methodological frameworks developed in previous chapters in healthy aging to explore the relationship between sequential effects and the fidelity of an individual’s mental representations of images.

5.2 Methods

Here we consider data from 15 older adults¹ who completed an ambiguous decision-making and mnemonic similarity task (described below) as part of a larger set of experiments in the lab (median age = 72, 13 female) . All data collection was completed in person and no subjects were excluded from analyses. Unless otherwise noted, all the following analyses are Bayesian and therefore are not strictly subject to the concerns about statistical power as it is traditionally conceptualized despite the small sample size.

¹Data collection is ongoing.

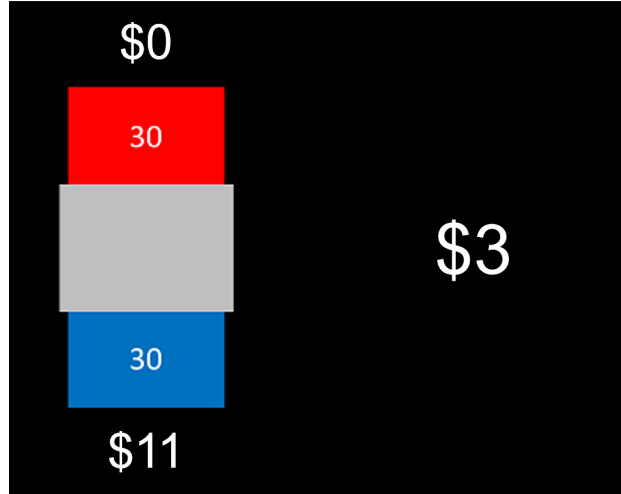


Figure 5.1: **Example trial of new AMB task for Older Adults.** The visual changes we make include delineating the left and right options more clearly (instead of having the lottery placed in the center as in Chapter 3) and making the font sizes generally larger.

5.2.1 Task Descriptions and Analyses

Ambiguity Task (AMB)

We use a task similar to that described in Experiment 3 in Chapter 2 of the dissertation, with some notable differences:

- We changed the visual display to make it more accessible to older adults, as (Tymula et al., 2012) use in their paper sampling people from ages 12 to 90 (see Figure 5.1).
- We increase the response window time from 3 seconds to 5.5 seconds to allow the OAs enough time to process and respond to each trial.
- We shorten the task from 4 blocks to 3 so as to reduce cognitive fatigue for a total of 135 total valid (i.e. non-catch) trials. We still include 4 catch trials per block.
- We include only ambiguous trials in the experiment. This also means that some trials are now repeated. From a choice modeling perspective, we therefore use a *Binomial*

likelihood and model success counts (in this case, a success is defined as choosing the lottery) as opposed to the previous Bernoulli likelihood.

- The experiment is now fully randomized. As before, we ensure that each block has the same median fixed reward, lottery reward, and ambiguity level. However, we do not ensure that 50% successive trials in each block increase in ambiguity and 50% decrease (as in Chapter 2). We made this change because we wanted to examine how the previous pseudo-randomized structure was impacting sequential effects of trial ambiguity.

Mnemonic Similarity Task (MST)

Here we use the “baseline” version of the MST as described in detail in Chapter 3. Briefly, the MST is a modified object recognition task that is split into an incidental encoding (study) phase and a subsequent test phase. Participants are shown 192 test trials where they chose whether stimuli presented had been seen before during the study phase (repeat), were similar to what they saw before but not exactly the same (lure), or had never been seen before in the experiment (foil).

5.3 Results

5.3.1 Ambiguity

Baseline

As in Chapter 2, we find that our drift rate decomposition captures meaningful variation, with β_1 , the coefficient for the Expected Value Difference between the fixed and lottery

AMB	Median(IQR)
β_0	-0.62(0.51)
β_1 (EV Diff)	1.99(2.04)
β_2 (Amb Level)	-0.52(1.32)
α	2.36(0.59)
τ	0.83(0.32)
<i>bias</i>	0.47(0.08)

Table 5.1: **AMB: Like in Chapter 2, drift rate decompositions capture meaningful variance in AMB.** Each cell shows aggregate posterior medians (Interquartile Range) for Drift Rate decompositions and other DDM parameters. Bolded parameters are ones we also test for sequential effects.

options being the largest: all else held constant, the greater the EV Difference between two options, the more quickly subjects accumulate evidence. Conversely, but consistent with our expectations, higher levels of ambiguity tended to slow down the rate of evidence accumulation, all else held constant. Keeping in mind our small sample size, we note that the slowing down of the drift rate as a function of ambiguity level is greater in the older adults as opposed to young adults in Chapter 2 ($\beta = 0.03$). While several researchers have demonstrated a positive relationship between age and risk aversion, e.g. (Grubb et al., 2016), less has been established about the relationship between age and ambiguity aversion e.g. (Raio et al., 2022). The median threshold and non-decision time (τ) parameters, are higher than those of the young adults that completed Experiment 3 in Chapter 2 (mean threshold = 1.84, mean NDT = 0.5). We do not interpret these differences to be a consequence of aging, though we may expect older adults to exercise more response caution and generally take longer to process and execute choice on any given trial (after all, the older adults also have longer to respond in this version of the experiment).

Sequential Effects

Contrary to our expectations, we find fewer sequential effects in the OAs than we did in the young adults overall, Table 5.2. We find again that that individuals seem to be sensitive

Trial Property	Proportion
<i>value</i> ↑	0.80
<i>value</i> ↓	1
<i>amb</i> ↑	0.07
<i>amb</i> ↓	0
<i>v.</i> ↑ <i>a.</i> ↑	0.07
<i>v.</i> ↑ <i>a.</i> ↓	0.07
<i>v.</i> ↓ <i>a.</i> ↑	0.66
<i>v.</i> ↓ <i>a.</i> ↓	0.31

Table 5.2: **Older adults are sensitive to most trial combinations.** We find that older adults are most sensitive to relative differences in reward *values* across trials.

to relative differences in value (both directions) but interestingly that relative increases or decreases in ambiguity do not play much of a role. However, consistent with our expectations, we do find that across these trial properties, there is greater variability in *which* drift rate or bias terms show evidence for recent history dependence. Indeed, we recover most effects on drift rate intercept and EV Difference terms and, for one subject, also on the bias term.

We further repeat our analysis in Chapter 3 to see the impact of accounting for sequential effects in our inference of risk and ambiguity tolerance in older adults. We simulate 100 choice sets from the posterior drift diffusion fit parameters and then refit them using our baseline logistic choice model with the binomial likelihood. We then calculate the (log of the) median ratio change for each individual for both risk and ambiguity tolerance. We further assess the *qualitative* impact of accounting for sequential effects by seeing how often sequential-effect-adjusted fits change interpretation (i.e. from risk seeking to risk averse). Like in Chapter 3, we find evidence that suggests even in this small dataset, that accounting and adjusting for sequential effects can produce qualitatively meaningful changes, Figure 5.2. Specifically, we find that for $\sim 27\%$ of subjects, their ambiguity tolerance parameter changes interpretation more than 90% of the time. It is similarly the case for one subject's risk tolerance, changing interpretation 100% of the time after adjusting for sequential effects.

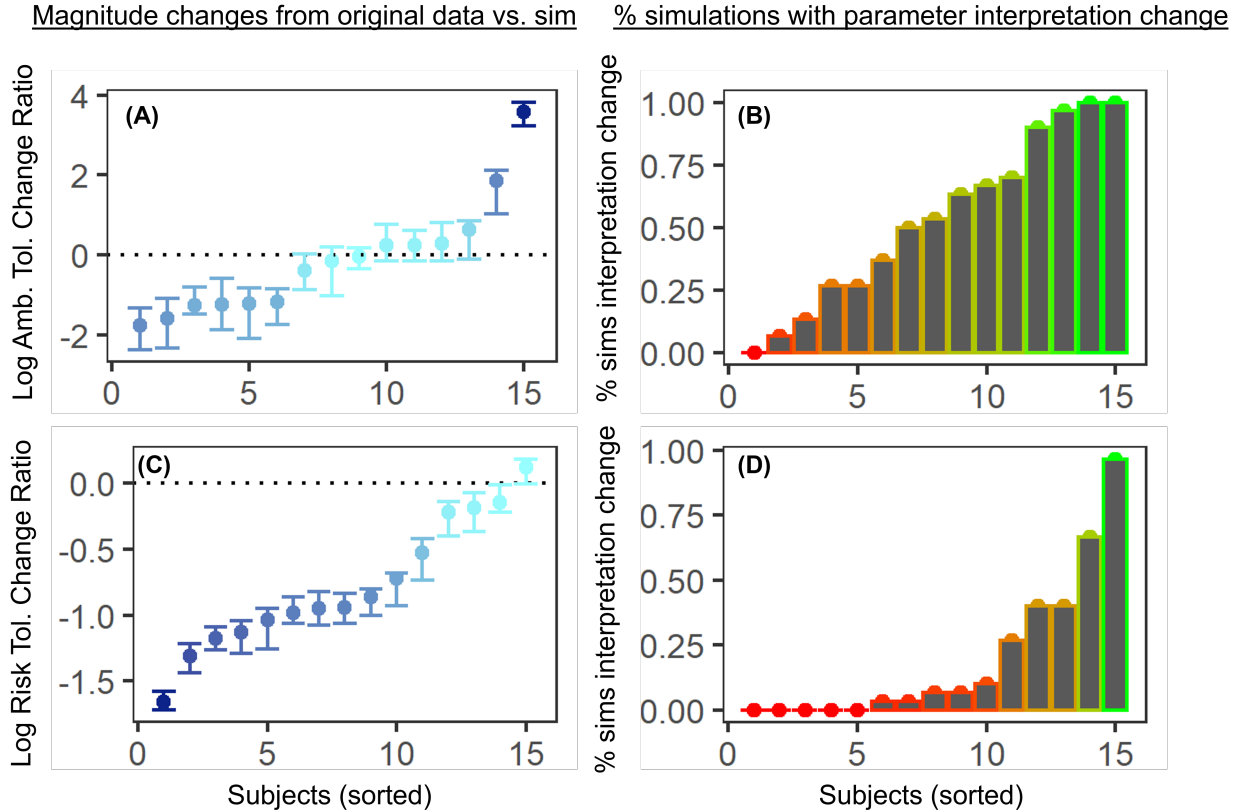


Figure 5.2: **AMB Task: Magnitude and interpretation changes in (A,B) ambiguity and (C,D) risk tolerance.** Ratio of (A) ambiguity tolerance and (C) risk tolerance estimates: $\log(\text{simulated choice set}/\text{observed data})$. We plot median ratios and IQRs from 1000 simulated choice sets. *Right:* The percentage of simulation-fit parameters that change interpretations in (B) ambiguity and (D) risk attitudes when compared to parameter fits in the original data. Subjects re-sorted by effect size in each plot (i.e. differently sorted in A/C vs. B/D).

5.3.2 Mnemonic Similarity Task

Participants had a median accuracy of 0.68(0.09) and were significantly more accurate when responding that a stimulus was a Repeat or a Foil vs. when a stimulus was a Lure (median accuracy Repeat = 0.89(0.1), Foil = 0.81(0.11), Lure = 0.34(0.24); Wilcoxon Ranked Sign Paired test $V = 120$, $p < 0.01$ for both comparisons; $V = 30$, $p = 0.14$ comparing median accuracy Repeat vs median accuracy Foil). Participants had a median RT of 1.36(0.15) with the median RT being the lowest for Repeat responses 1.25(0.11), then Foil responses 1.33(0.22), and 1.75(0.19) ($V_{Repeat,Lure;Foil,Lure} = 120$, $V_{Repeat,Foil} = 101$; all $p < 0.02$), Fig-

ure 5.3. Finally, subjects tended to most often indicate that a stimulus was a Repeat with median proportion = 0.50(0.04), then Foil with median proportion = 0.31(0.07), followed by Lure with median proportion = 0.18(0.09) ($V_{Repeat,Lure;Repeat,Foil} = 120$, $V_{Lure,Foil} = 112$; all $p < 0.02$, Figure 5.3).

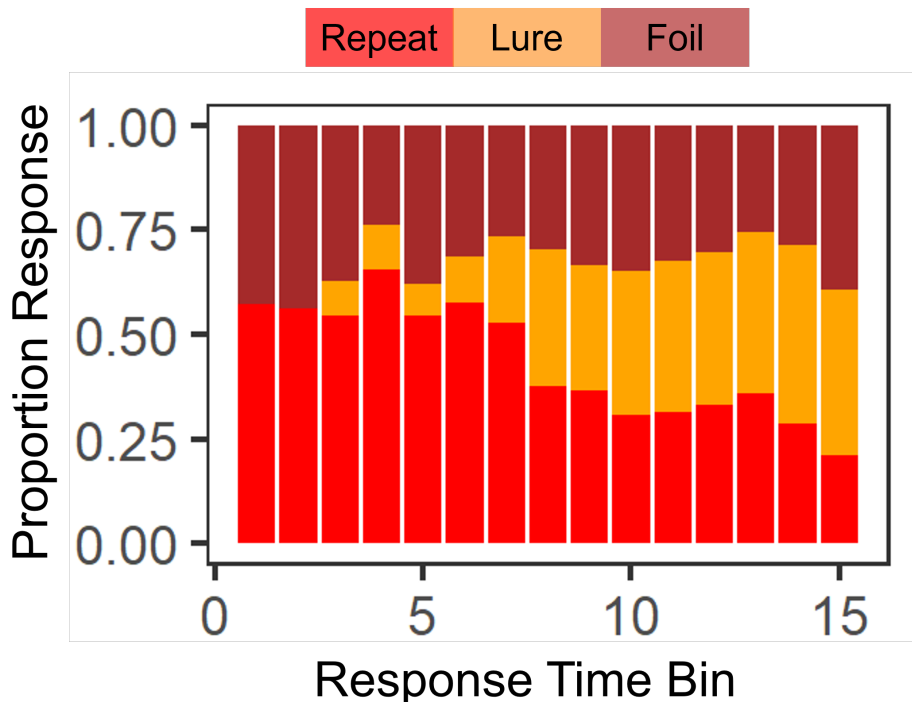


Figure 5.3: **Choice proportions change as a function of response time.** In this dataset, we find that fastest responses tend to be Repeat and the slowest tend to be Lures – indeed participants make no Lure responses in the fastest 2 RT bins.

Posterior Summaries and LDI-LBA Correlations

Interestingly, while there are several differences across all the MST data we have fit, we find approximately the same qualitative patterns in the LBA parameters (Figure 5.4), but not in how the LDI and the LBA parameters correlate (Figure 5.5). Start point upper bounds and drift rates are substantially lower for Lure responses. This makes sense given the meager proportion of responses where the subjects even indicated that a stimulus was a Lure. Of note is also that mean NDT is fairly high at 0.59 seconds, something that has been observed

in other experiments with older populations (Schuch, 2016).

We see an interesting pattern of results when we consider the relationship between the LDI and LBA parameters, Figure 5.5. We see a positive correlation between the LDI and Lure accumulator drift rate ($\tau = 0.39, p < 0.05$), as expected. We see no statistically significant relationships between the LDI and start points. Surprisingly, we further find a *negative* correlation between the LDI and Foil accumulator drift rate ($\tau = -0.41, p < 0.05$). This is contrary to our expectations, which are that the LDI correlates positively (if at all) with the Foil accumulator drift rate. This is due to the fact that the LDI captures the ability to discriminate between objects that haven't been seen before in the context of the experiment (this includes both lures and foils). Finally, we also do not see a relationship between the LDI and Repeat accumulator drift rate. We do not speculate further as to whether these differences are due to a substantive reason or due to the sample size of 15.

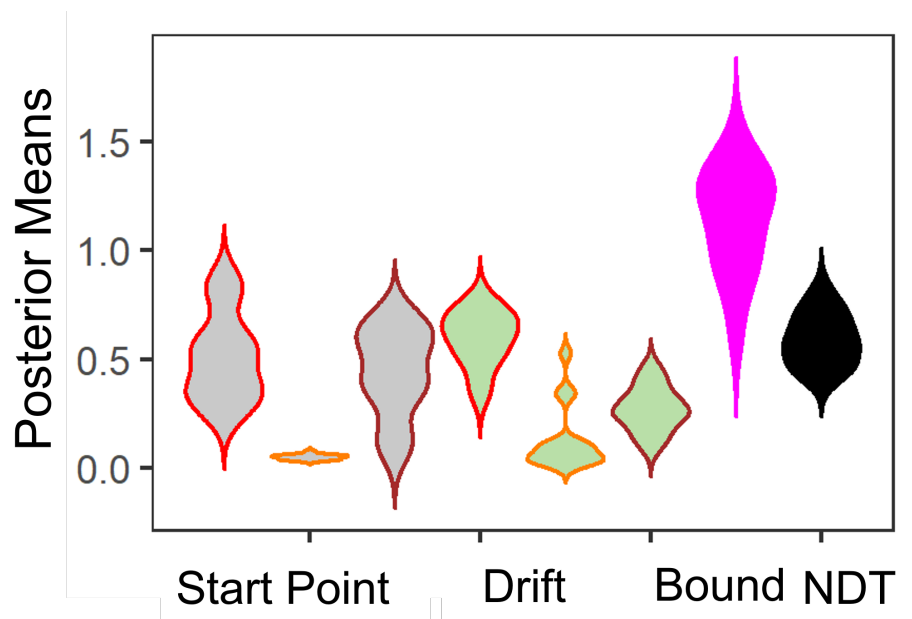


Figure 5.4: **LBA model posteriors.** We find that the Lure accumulator start point upper bound and drift rate tend to be lower than the equivalent parameters for the Repeat and Foil accumulators, consistent with expectations.

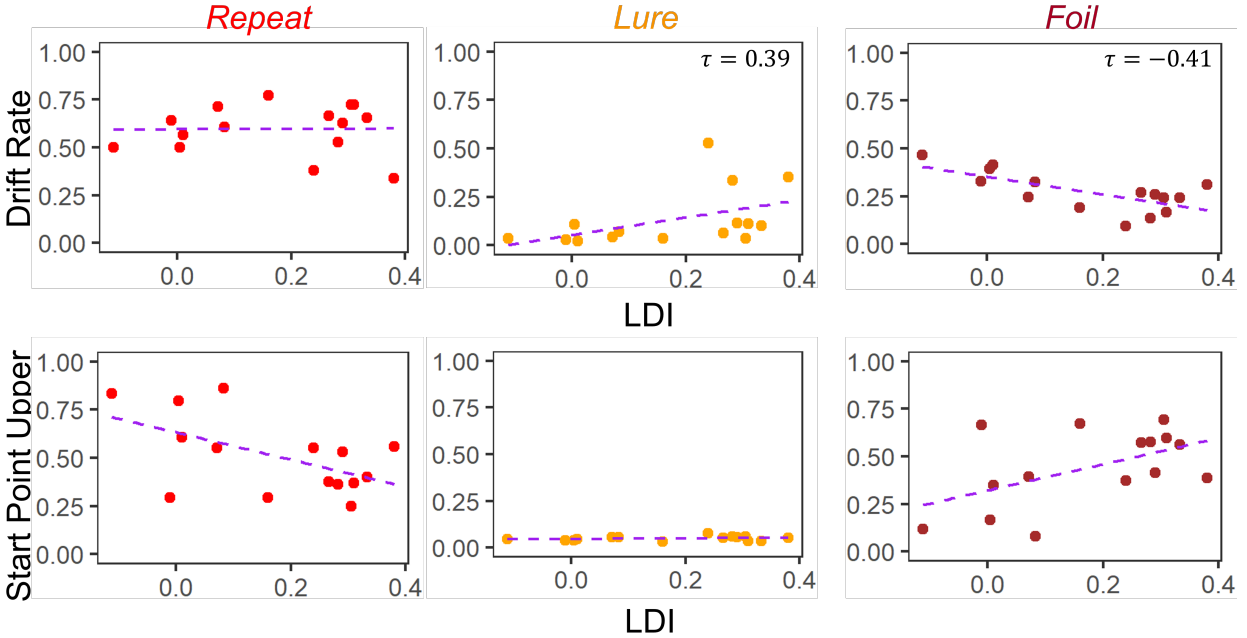


Figure 5.5: **LDI only correlates with Lure and Foil accumulators.** In the data we have collected thus far, we find different qualitative relationships between LDI and the LBA parameters of interest. Correlations shown in plots are statistically significant ($p < 0.01$) after adjusting for multiple comparisons.

5.3.3 Decision making under uncertainty and MST

Finally, we consider the relationship between sequential effects in economic choice and the distinctiveness of individual representations as captured by the MST. We contrast how risk and ambiguity tolerances (with and without sequential effect adjustment) correlate with LBA/MST parameters. As our hypotheses concern sequential effects in the AMB task and behavior in the MST, we expect that any relationship between these independently measured variables will be between the sequential-effect-adjusted risk/ambiguity parameters and MST parameters compared to the AMB parameters fit using the “baseline” model.

When we correlate the original risk/ambiguity tolerance parameters and MST/LBA parameters (including the LDI) we find no statistically significant relationship across all the comparisons. Interestingly, when using the sequential effect-adjusted risk/ambiguity tolerance parameters, we find two correlations that are statistically meaningful: adjusted risk

tolerance correlates positively with start point for Repeat accumulator and negatively with start point for Foil accumulator. Overall, risk tolerance estimates shrink when adjusting for sequential effects – suggesting that even within these generally risk averse individuals, the less risk averse may have a higher tendency to make Repeat responses and a lower tendency to make Foil responses.

5.4 Discussion

In this chapter, we motivate why sequential effects warrant careful consideration in aging: they are a type of perseveration that could have clinical implications. We demonstrate that, even in a sample size of 15, we recover substantive and *specific* sequential effects in older adults when we jointly model choice and response time, particularly when considering cross-trial differences in value. We also show qualitative differences in our inferences of each individual’s risk and ambiguity tolerance: several subjects are re-classified as the opposite category under which they were originally classified.

Separately, we model these same subjects’ behavior in a memory discrimination experiment, the MST. We find overall comparable patterns to our previous analyses in Chapter 4 when considering raw behavior and in our posterior model parameters. When comparing the relationship between the choice-based LDI and LBA parameters, we only find a relationship between the Lure accumulator drift rate and the LDI. This tells us that even in an aging population, people with higher LDIs deliberate more rapidly (higher drift rate) over whether they classify a stimulus as a Lure.

Finally, we consider the relationship between the parameters fit across these two experiments. We find initial evidence that suggests a relationship between risk tolerance (after adjusting for sequential effects) and starting point upper bounds in the MST. At the very least, this

differential relationship (with and without sequential effect adjustment) provides further evidence that sequential effects could indeed be capturing meaningful variability in human behavior under uncertainty. The relationship itself between risk tolerance and start point might be because these economic choice parameters also capture “tendencies” of sort. We speculate that we see a relationship only between risk and LBA parameters, and *not* ambiguity and LBA parameters because the magnitude of the changes in sequential-effect-adjusted risk tolerances is generally greater than that of the changes in sequential-effect-adjusted ambiguity tolerances, Figure 5.2.

Overall, we demonstrate that our methods developed earlier in the dissertation are easy to apply and can provide results even in specific (i.e. not typical control) populations and small sample sizes. Further data is needed in order to speculate as to whether the cross-task relationships hold in older adults.

Chapter 6

Conclusion

In this thesis, we have demonstrated how response time models can be easily adapted to incorporate psychologically important considerations of and novel inferences about human decision making. Once sources of variability have been identified and operationalized within the decision-making and experimental context, we can readily test how important these sources are.

In economic choice under immediate and temporal uncertainty, we explore the question of how making successive choices in an experiment actually impacts behavior during the experiment itself. We do this by incorporating recent history, in particular how *different* the current trial is to the previous, into our models of choice and response time. We demonstrate that the majority of participants across 4 experiments show evidence of sensitivity to these relative trial differences. We further demonstrate that this is not just “meaningful” in the context of explaining previously uncharacterized noise, but it is also “meaningful” in that it can change our qualitative and quantitative inferences about the parameters the experiments themselves are designed to make inferences about. We begin to motivate and show how this framework can provide cognitively and potentially neurologically meaningful information as

humans age.

In memory discrimination, we adapt a version of the Linear Ballistic Accumulator (LBA) (Brown & Heathcote, 2008) to introduce a joint model of choice and response time to a literature that has largely considered modeling this process by using only choice data. We demonstrate that our model recovers consistent qualitative patterns in the data across several independent datasets. We also show how our model based parameters allow an interpretable decomposition of the previous choice-only gold standard. Finally, we explore the external validity of our model parameters and the choice-based-measure by considering additional variables like age and resting state functional connectivity.

Together, we hope to have demonstrated the value of articulating and subsequently modeling sources of variability in our constructs of interest that may previously not have been sufficiently explored. Again, as we state in the introduction, this is not to suggest that *all* sources of underexplored variability are “meaningful” but, that as decision scientists studying multi-dimensional constructs, this is an exercise worth careful consideration.

We conclude this dissertation by highlighting that, while the primary contributions of this thesis are model-based, researchers who do not use models can still empirically engage with this work. Sequential sampling models of response times are particularly valuable in that they provide formal frameworks and interpretable parameters, however, they are not strictly necessary in order to gain insight from response times. As we show in the previous three chapters, looking at the raw response time data itself can provide meaningful information. In particular, in the economic decision-making tasks, our primary question had to do with trial sequences and behavior: looking at individual differences in response time distributions for the trial properties of interest (relative increases/decreases in difficulty and/or uncertainty) *prior* to any modeling already suggested that this could be a source of variability worth considering. Similarly, in the memory discrimination experiments, we demonstrated how choice proportions changed as a function of response time. Median response times (RTs) as a

function of response type could also provide valuable information especially when considered in conjunction with the LDI and our modeling parameters. We may, for example, expect there to be qualitatively similar relationships between LDI and median RTs as we find between LDI and the LBA parameters. In a subset of the experiments we model, we find that median RTs for each response type do indeed correlate with the drift rates for each response type but *not* with the LDI. This suggests that the RTs, even prior to being decomposed by our formal model, can tell us valuable and perhaps more “process pure” information about an individual’s memory discrimination ability. We therefore encourage researchers to engage with response time distributions and summary statistics as a useful means to not only explore the value of the contributions of this thesis, but more importantly, to further motivate and test their own hypotheses.

References

- Addessi, E., Bellagamba, F., Delfino, A., De Petrillo, F., Focaroli, V., Macchitella, L., . . . others (2014). Waiting by mistake: Symbolic representation of rewards modulates intertemporal choice in capuchin monkeys, preschool children and adult humans. *Cognition*, *130*(3), 428–441.
- Allais, M. (1990). Allais paradox. In *Utility and probability* (pp. 3–9). Springer.
- Anderson, J. R. (1990). *The adaptive character of thought*. Psychology Press. doi: <https://doi.org/10.4324/9780203771730>
- Andreoni, J., Callen, M., Hussain, K., Khan, M., & Sprenger, C. (2016). Using preference estimates to customize incentives: an application to polio vaccination drives in pakistan. *Journal of the European Economic Association*.
- Andrés, S., Lázaro, L., Salamero, M., Boget, T., Penadés, R., & Castro-Fornieles, J. (2008). Changes in cognitive dysfunction in children and adolescents with obsessive-compulsive disorder after treatment. *Journal of Psychiatric Research*, *42*(6), 507–514.
- Banavar, N., Lee, M., & Bornstein, A. M. (2021). Sequential effects in non-sequential tasks. In *Proceedings of the 19th international conference on cognitive modeling*.
- Banavar, N. V., & Bornstein, A. (2023, Feb). *Independent, not irrelevant: Trial order causes systematic misestimation of economic choice traits*. PsyArXiv. Retrieved from psyarxiv.com/a8gz3 doi: 10.31234/osf.io/a8gz3
- Barlow, H. B., et al. (1961). Possible principles underlying the transformation of sensory messages. *Sensory communication*, *1*(01), 217–233.
- Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of behavioral decision making*, *16*(3), 215–233.
- Bartra, O., McGuire, J. T., & Kable, J. W. (2013). The valuation system: a coordinate-based meta-analysis of bold fmri experiments examining neural correlates of subjective value. *Neuroimage*, *76*, 412–427.
- Bickel, W. K., Koffarnus, M. N., Moody, L., & Wilson, A. G. (2014). The behavioral- and neuro-economic process of temporal discounting: A candidate behavioral marker of addiction. *Neuropharmacology*, *76*, 518–527.

- Bornstein, A., Khaw, M., Shohamy, D., & Daw, N. (2017). Reminders of past choices bias decisions for reward in humans. *Nature Communications*, *8*(1), 1–9.
- Brooks, H., & Sokol-Hessner, P. (2020). Quantifying the immediate computational effects of preceding outcomes on subsequent risky choices. *Scientific reports*, *10*(1), 1–10.
- Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive psychology*, *57*(3), 153–178.
- Carroll, D. J., Blakey, E., & FitzGibbon, L. (2016). Cognitive flexibility in young children: Beyond perseveration. *Child Development Perspectives*, *10*(4), 211–215.
- Clithero, J. (2018). Response times in economics: Looking through the lens of sequential sampling models. *Journal of Economic Psychology*, *69*, 61–86.
- Commodari, E., & Guarnera, M. (2008). Attention and aging. *Aging clinical and experimental research*, *20*(6), 578–584.
- Croson, R., & Gneezy, U. (2009). Gender differences in preferences. *Journal of Economic literature*, *47*(2), 448–474.
- Cross, C. P., Copping, L. T., & Campbell, A. (2011). Sex differences in impulsivity: a meta-analysis. *Psychological bulletin*, *137*(1), 97.
- Curtis, C., & Lee, D. (2010). Beyond working memory: the role of persistent activity in decision making. *Trends in cognitive sciences*, *14*(5), 216–222.
- Dai, J., Pachur, T., Pleskac, T. J., & Hertwig, R. (2019). What the future holds and when: A description–experience gap in intertemporal choice. *Psychological Science*, *30*(8), 1218–1233.
- Dasgupta, I., Schulz, E., Goodman, N., & Gershman, S. (2018). Remembrance of inferences past: Amortization in human hypothesis generation. *Cognition*, *178*, 67–81.
- Daw, N., Courville, A., & Dayan, P. (2008). Semi-rational models of conditioning: The case of trial order. *The probabilistic mind*, 431–452.
- Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *The quarterly journal of economics*, *75*(4), 643–669.
- ES, B. (1983). Impulsivity: Cognitive, behavioral, and psychophysiological correlates. *Biological bases of sensation seeking, impulsivity, and anxiety*.
- Estle, S. J., Green, L., Myerson, J., & Holt, D. D. (2006). Differential effects of amount on temporal and probability discounting of gains and losses. *Memory & cognition*, *34*, 914–928.
- Evenden, J. L. (1999). Varieties of impulsivity. *Psychopharmacology*, *146*(4), 348–361.

- Fischer, J., & Whitney, D. (2014). Serial dependence in visual perception. *Nature neuroscience*, *17*(5), 738–743.
- Fischer-Baum, S., Miozzo, M., Laiacona, M., & Capitani, E. (2016). Perseveration during verbal fluency in traumatic brain injury reflects impairments in working memory. *Neuropsychology*, *30*(7), 791.
- Foldi, N. S., Helm-Estabrooks, N., Redfield, J., & Nickel, D. G. (2003). Perseveration in normal aging: A comparison of perseveration rates on design fluency and verbal generative tasks. *Aging, Neuropsychology, and Cognition*, *10*(4), 268–280.
- Forbes, N., Carrick, L., McIntosh, A., & Lawrie, S. (2009). Working memory in schizophrenia: a meta-analysis. *Psychological medicine*, *39*(6), 889–905.
- Forstmann, B. U., Ratcliff, R., & Wagenmakers, E.-J. (2016). Sequential sampling models in cognitive neuroscience: Advantages, applications, and extensions. *Annual review of psychology*, *67*, 641–666.
- Frey, R., Pedroni, A., Mata, R., Rieskamp, J., & Hertwig, R. (2017). Risk preference shares the psychometric structure of major psychological traits. *Science advances*, *3*(10), e1701381.
- Frey, R., Richter, D., Schupp, J., Hertwig, R., & Mata, R. (2021). Identifying robust correlates of risk preference: A systematic approach using specification curve analysis. *Journal of Personality and Social Psychology*, *120*(2), 538.
- Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annu. Rev. Neurosci.*, *30*, 535–574.
- Goldberg, E. (1986). Varieties of perseveration: A comparison of two taxonomies. *Journal of Clinical and Experimental Neuropsychology*, *8*(6), 710–726.
- Green, L., Fry, A. F., & Myerson, J. (1994). Discounting of delayed rewards: A life-span comparison. *Psychological science*, *5*(1), 33–36.
- Grubb, M. A., Tymula, A., Gilaie-Dotan, S., Glimcher, P. W., & Levy, I. (2016). Neuroanatomy accounts for age-related changes in risk preferences. *Nature communications*, *7*(1), 13822.
- Guan, M. (2019). *A cognitive modeling analysis of risk in sequential choice tasks*. University of California, Irvine.
- Gureckis, T., Martin, J., McDonnell, J., Rich, A., Markant, D., Coenen, A., . . . Chan, P. (2016). psiturk: An open-source framework for conducting replicable behavioral experiments online. *Behavior research methods*, *48*(3), 829–842.
- Hasher, L., & Zacks, R. T. (1988). Working memory, comprehension, and aging: A review and a new view. *Psychology of learning and motivation*, *22*, 193–225.

- Hayden, B. Y., & Platt, M. L. (2007). Temporal discounting predicts risk sensitivity in rhesus macaques. *Current Biology*, *17*(1), 49–53.
- Hertwig, R., & Erev, I. (2009). The description–experience gap in risky choice. *Trends in cognitive sciences*, *13*(12), 517–523.
- Hunter, D. R., et al. (2002). *Risk perception and risk tolerance in aircraft pilots* (Tech. Rep.). United States. Department of Transportation. Federal Aviation Administration
- Hunter, L., Bornstein, A., & Hartley, C. (2018). A common deliberative process underlies model-based planning and patient intertemporal choice. *bioRxiv*, 499707.
- Kable, J. W., & Glimcher, P. W. (2007). The neural correlates of subjective value during intertemporal choice. *Nature neuroscience*, *10*(12), 1625–1633.
- Kahneman, D., & Tversky, A. (2013). Prospect theory: An analysis of decision under risk. In *Handbook of the fundamentals of financial decision making: Part i* (pp. 99–127). World Scientific.
- Khaw, M., Li, Z., & Woodford, M. (2018). Temporal discounting and search habits: evidence for a task-dependent relationship. *Frontiers in Psychology*, *9*, 2102.
- Khaw, M., Li, Z., & Woodford, M. (2021). Cognitive imprecision and small-stakes risk aversion. *The review of economic studies*, *88*(4), 1979–2013.
- Kirby, K. N. (2009). One-year temporal stability of delay-discount rates. *Psychonomic bulletin & review*, *16*(3), 457–462.
- Kiyonaga, A., Scimeca, J., Bliss, D., & Whitney, D. (2017). Serial dependence across perception, attention, and memory. *Trends in Cognitive Sciences*, *21*(7), 493–497.
- Konova, A., Lopez-Guzman, S., Urmanche, A., Ross, S., Louie, K., Rotrosen, J., & Glimcher, P. (2020). Computational markers of risky decision-making for identification of temporal windows of vulnerability to opioid use in a real-world clinical setting. *JAMA psychiatry*, *77*(4), 368–377.
- Konovalov, A., & Krajbich, I. (2019). Revealed strength of preference: Inference from response times. *Judgment and Decision making*, *14*(4).
- Koop, G. J., & Johnson, J. G. (2012). The use of multiple reference points in risky decision making. *Journal of Behavioral Decision Making*, *25*(1), 49–62.
- Koop, G. J., & Johnson, J. G. (2013). The response dynamics of preferential choice. *Cognitive psychology*, *67*(4), 151–185.
- Krajbich, I. (2019). Accounting for attention in sequential sampling models of decision making. *Current opinion in psychology*, *29*, 6–11.
- Krajbich, I., Lu, D., Camerer, C., & Rangel, A. (2012). The attentional drift-diffusion model extends to simple purchasing decisions. *Frontiers in psychology*, *3*, 193.

- Kugler, T., Connolly, T., & Ordóñez, L. D. (2012). Emotion, decision, and risk: Betting on gambles versus betting on people. *Journal of Behavioral Decision Making*, *25*(2), 123–134.
- Kumar, M., & Persaud, A. (2002). Pure contagion and investors' shifting risk appetite: analytical issues and empirical evidence. *International Finance*, *5*(3), 401–436.
- Kwako, L., Momenan, R., Litten, R., Koob, G., & Goldman, D. (2016). Addictions neuroclinical assessment: a neuroscience-based framework for addictive disorders. *Biological psychiatry*, *80*(3), 179–189.
- Lazzaro, S., Rutledge, R., Burghart, D., & Glimcher, P. (2016). The impact of menstrual cycle phase on economic choice and rationality. *PLoS One*, *11*(1), e0144080.
- Ledgerwood, D. M., Alessi, S. M., Phoenix, N., & Petry, N. M. (2009). Behavioral assessment of impulsivity in pathological gamblers with and without substance use disorder histories versus healthy controls. *Drug and alcohol dependence*, *105*(1-2), 89–96.
- Lee, M. (2018). Bayesian methods in cognitive modeling. *The Stevens' handbook of experimental psychology and cognitive neuroscience*, *5*, 37–84.
- Lee, M., & Wagenmakers, E.-J. (2014). *Bayesian cognitive modeling: A practical course*. Cambridge university press.
- Leiser, D., Azar, O. H., & Hadar, L. (2008). Psychological construal of economic behavior. *Journal of Economic Psychology*, *29*(5), 762–776.
- Lempert, K., & Phelps, E. (2016). The malleability of intertemporal choice. *Trends in cognitive sciences*, *20*(1), 64–74.
- Lempert, K. M., Glimcher, P. W., & Phelps, E. A. (2015). Emotional arousal and discount rate in intertemporal choice are reference dependent. *Journal of Experimental Psychology: General*, *144*(2), 366.
- Levy, I., Snell, J., Nelson, A., Rustichini, A., & Glimcher, P. (2010). Neural representation of subjective value under risk and ambiguity. *Journal of neurophysiology*, *103*(2), 1036–1047.
- Liberman, N., & Trope, Y. (2003). Construal level theory of intertemporal judgment and decision. In *Time and decision: Economic and psychological perspectives on intertemporal choice*. Russell Sage Foundation.
- Lockhead, G. R., & King, M. C. (1983). A memory model of sequential effects in scaling tasks. *Journal of Experimental Psychology: Human Perception and Performance*, *9*(3), 461.
- Loewenstein, G., O'Donoghue, T., & Bhatia, S. (2015). Modeling the interplay between affect and deliberation. *Decision*, *2*(2), 55.
- Long, N. M., Lee, H., & Kuhl, B. A. (2016). Hippocampal mismatch signals are modulated by the strength of neural predictions and their similarity to outcomes. *Journal of Neuroscience*, *36*(50), 12677–12687.

- Lopez-Guzman, S., Konova, A. B., Louie, K., & Glimcher, P. W. (2018). Risk preferences impose a hidden distortion on measures of choice impulsivity. *PLoS One*, *13*(1), e0191357.
- Luck, S., & Gold, J. (2008). The construct of attention in schizophrenia. *Biological psychiatry*, *64*(1), 34–39.
- MacKillop, J., Amlung, M. T., Few, L. R., Ray, L. A., Sweet, L. H., & Munafò, M. R. (2011). Delayed reward discounting and addictive behavior: a meta-analysis. *Psychopharmacology*, *216*, 305–321.
- Manassi, M., Ghirardo, C., Canas-Bajo, T., Ren, Z., Prinzmetal, W., & Whitney, D. (2021). Serial dependence in the perceptual judgments of radiologists. *Cognitive research: principles and implications*, *6*, 1–13.
- Marr, D. (1971). Simple memory: a theory for archicortex. *Philosophical Transactions of the Royal Society of London. B, Biological Sciences*, *262*(841), 23–81.
- May, C. P., Zacks, R. T., Hasher, L., & Multhaup, K. S. (1999). Inhibition in the processing of garden-path sentences. *Psychology and Aging*, *14*(2), 304.
- Orquin, J. L., & Loose, S. M. (2013). Attention and choice: A review on eye movements in decision making. *Acta psychologica*, *144*(1), 190–206.
- Pedroni, A., Frey, R., Bruhin, A., Dutilh, G., Hertwig, R., & Rieskamp, J. (2017). The risk elicitation puzzle. *Nature Human Behaviour*, *1*(11), 803–809.
- Peters, J., & Büchel, C. (2011). The neural mechanisms of inter-temporal decision-making: understanding variability. *Trends in cognitive sciences*, *15*(5), 227–239.
- Peters, J., & D’Esposito, M. (2020). The drift diffusion model as the choice rule in inter-temporal and risky choice: a case study in medial orbitofrontal cortex lesion patients and controls. *PLoS computational biology*, *16*(4), e1007615.
- 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* (Vol. 124, pp. 1–10).
- Pulcu, E., Trotter, P., Thomas, E., McFarquhar, M., Juhász, G., Sahakian, B., ... Elliott, R. (2014). Temporal discounting in major depressive disorder. *Psychological medicine*, *44*(9), 1825–1834.
- Radulescu, A., Holmes, K., & Niv, Y. (2020, May). *On the convergent validity of risk sensitivity measures*. PsyArXiv. Retrieved from psyarxiv.com/qdhx4 doi: 10.31234/osf.io/qdhx4
- Raio, C. M., Lu, B. B., Grubb, M., Shields, G. S., Slavich, G. M., & Glimcher, P. (2022). Cumulative lifetime stressor exposure assessed by the strain predicts economic ambiguity aversion. *Nature communications*, *13*(1), 1686.

- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological review*, 85(2), 59.
- Reeck, C., Wall, D., & Johnson, E. J. (2017). Search predicts and changes patience in intertemporal choice. *Proceedings of the National Academy of Sciences*, 114(45), 11890–11895.
- Rogers, T., & Bazerman, M. H. (2008). Future lock-in: Future implementation increases selection of ‘should’ choices. *Organizational behavior and human decision processes*, 106(1), 1–20.
- Salthouse, T. (1994). The aging of working memory. *Neuropsychology*, 8(4), 535.
- Salthouse, T. A. (2009). When does age-related cognitive decline begin? *Neurobiology of aging*, 30(4), 507–514.
- Samuelson, P. A. (1937). A note on measurement of utility. *The review of economic studies*, 4(2), 155–161.
- Sandson, J., & Albert, M. L. (1984). Varieties of perseveration. *Neuropsychologia*, 22(6), 715–732.
- Savage, L. J. (1972). *The foundations of statistics*. Courier Corporation.
- Scherbaum, S., Dshemuchadse, M., & Goschke, T. (2012). Building a bridge into the future: dynamic connectionist modeling as an integrative tool for research on intertemporal choice. *Frontiers in psychology*, 3, 514.
- Scherder, E., Eggermont, L., Visscher, C., Scheltens, P., & Swaab, D. (2011). Understanding higher level gait disturbances in mild dementia in order to improve rehabilitation: ‘last in–first out’. *Neuroscience & Biobehavioral Reviews*, 35(3), 699–714.
- Schuch, S. (2016). Task inhibition and response inhibition in older vs. younger adults: A diffusion model analysis. *Frontiers in psychology*, 7, 1722.
- Schulte-Mecklenbeck, M., Johnson, J. G., Böckenholt, U., Goldstein, D. G., Russo, J. E., Sullivan, N. J., & Willemsen, M. C. (2017). Process-tracing methods in decision making: On growing up in the 70s. *Current Directions in Psychological Science*, 26(5), 442–450.
- Schwartz, O., Hsu, A., & Dayan, P. (2007). Space and time in visual context. *Nature Reviews Neuroscience*, 8(7), 522–535.
- Serpell, L., Waller, G., Fearon, P., & Meyer, C. (2009). The roles of persistence and perseveration in psychopathology. *Behavior therapy*, 40(3), 260–271.
- Simon, H. A. (1990). Bounded rationality. *Utility and probability*, 15–18.
- Simoncelli, E. P., & Olshausen, B. A. (2001). Natural image statistics and neural representation. *Annual review of neuroscience*, 24(1), 1193–1216.

- Slovic, P. (2020). The construction of preference. In *Shaping entrepreneurship research* (pp. 104–119). Routledge.
- Sokol-Hessner, P., Raio, C. M., Gottesman, S. P., Lackovic, S. F., & Phelps, E. A. (2016). Acute stress does not affect risky monetary decision-making. *Neurobiology of stress*, *5*, 19–25.
- Solway, A., Lohrenz, T., & Montague, P. R. (2017). Simulating future value in intertemporal choice. *Scientific Reports*, *7*(1), 43119.
- Stark, C. E., Noche, J. A., Ebersberger, J. R., Mayer, L., & Stark, S. M. (2023). Optimizing the mnemonic similarity task for efficient, widespread use. *Frontiers in Behavioral Neuroscience*, *17*, 1080366.
- Stark, S. M., Kirwan, C. B., & Stark, C. E. (2019). Mnemonic similarity task: A tool for assessing hippocampal integrity. *Trends in cognitive sciences*, *23*(11), 938–951.
- Steinglass, J. E., Lempert, K. M., Choo, T.-H., Kimeldorf, M. B., Wall, M., Walsh, B. T., ... Simpson, H. B. (2017). Temporal discounting across three psychiatric disorders: anorexia nervosa, obsessive compulsive disorder, and social anxiety disorder. *Depression and Anxiety*, *34*(5), 463–470.
- Stewart, N., Brown, G. D., & Chater, N. (2005). Absolute identification by relative judgment. *Psychological review*, *112*(4), 881.
- Stillman, P. E., Krajbich, I., & Ferguson, M. J. (2020). Using dynamic monitoring of choices to predict and understand risk preferences. *Proceedings of the National Academy of Sciences*, *117*(50), 31738–31747.
- Tal, E. (2020). Measurement in Science. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Fall 2020 ed.). Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/fall2020/entries/measurement-science/>.
- Team, S. D. (2023). *Rstan: the r interface to stan*. Retrieved from <https://mc-stan.org/> (R package version 2.21.8)
- Thaler, R. (1981). Some empirical evidence on dynamic inconsistency. *Economics letters*, *8*(3), 201–207.
- Thurstone, L. L. (1994). A law of comparative judgment. *Psychological review*, *101*(2), 266.
- Turner, B. M., Palestro, J. J., Miletić, S., & Forstmann, B. U. (2019). Advances in techniques for imposing reciprocity in brain-behavior relations. *Neuroscience & Biobehavioral Reviews*, *102*, 327–336.
- Tversky, A., & Kahneman, D. (1989). Rational choice and the framing of decisions. In *Multiple criteria decision making and risk analysis using microcomputers* (pp. 81–126). Springer.

- Tymula, A., Glimcher, P. W., Levy, I., & Belmaker, L. A. R. (2012). Separating risk and ambiguity preferences across the life span: Novel findings and implications for policy. *Unpublished manuscript*.
- Ubfal, D. (2016). How general are time preferences? eliciting good-specific discount rates. *Journal of Development Economics*, *118*, 150–170.
- Van Patten, R., Kaufman, D., Mitchell, S., Sachs, B., & Loring, D. (2015). Perseverative error subtypes in patients with alzheimer’s disease and mild cognitive impairment. *J Neurol Psychol*.
- Verplanck, W. S., & Blough, D. S. (1958). Randomized stimuli and the non-independence of successive responses at the visual threshold. *The Journal of general psychology*, *59*(2), 263–272.
- Vlaev, I. (2018). Local choices: Rationality and the contextuality of decision-making. *Brain sciences*, *8*(1), 8.
- Von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior* princeton. *Princeton University Press*, 1947, 1953.
- Wabersich, D., & Vandekerckhove, J. (2014). Extending jags: A tutorial on adding custom distributions to jags (with a diffusion model example). *Behavior research methods*, *46*(1), 15–28.
- Wahlheim, C. N., Christensen, A. P., Reagh, Z. M., & Cassidy, B. S. (2022). Intrinsic functional connectivity in the default mode network predicts mnemonic discrimination: A connectome-based modeling approach. *Hippocampus*, *32*(1), 21–37.
- Weafer, J., & de Wit, H. (2014). Sex differences in impulsive action and impulsive choice. *Addictive behaviors*, *39*(11), 1573–1579.
- Webb, R. (2019). The (neural) dynamics of stochastic choice. *Management Science*, *65*(1), 230–255.
- Weber, E. U. (2004). Perception matters: Psychophysics for economists. *The psychology of economic decisions*, *2*(163-176), 14–41.
- Weber, E. U. (2010). Risk attitude and preference. *Wiley Interdisciplinary Reviews: Cognitive Science*, *1*(1), 79–88.
- Weber, E. U., Blais, A.-R., & Betz, N. E. (2002). A domain-specific risk-attitude scale: Measuring risk perceptions and risk behaviors. *Journal of behavioral decision making*, *15*(4), 263–290.
- Westfall, H. A., & Lee, M. D. (2021). A model-based analysis of the impairment of semantic memory. *Psychonomic Bulletin & Review*, *28*, 1484–1494.
- Yassa, M. A., & Stark, C. E. (2011). Pattern separation in the hippocampus. *Trends in neurosciences*, *34*(10), 515–525.

Zimmermann, J., Glimcher, P., & Louie, K. (2018). Multiple timescales of normalized value coding underlie adaptive choice behavior. *Nature communications*, *9*(1), 1–11.