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

Nidhi V. Banavar
Department of Cognitive Sciences
UC - Irvine
nbanavar@uci.edu

Aaron M. Bornstein
Department of Cognitive Sciences
UC - Irvine
aaron.bornstein@uci.edu

Abstract

Decision making under uncertainty necessitates complex computations which are traded-off with the need for efficiency. This is particularly relevant in the context of experiments where individuals make a sequence of choices and previous computations may be leveraged to support efficiency. However, it is an open question as to whether humans do indeed reference the recent past, especially in complex environments where it is task-incongruent to do so (e.g. non-sequential experiments). In behavioral economic experiments with randomized or unstructured choice sets, trial-level sequential dependencies are generally assumed to be present only in motor or perceptual operations. Here, we explicitly model trial-property-driven sequential effects in response time data in two data sets: intertemporal choice and risky and ambiguous choice. We find evidence for widespread sequential effects. These effects are modulated by decision difficulty and trial-level uncertainty in both tasks. Our results add to the growing literature demonstrating trial-level sequential dependencies in higher-order cognition.

Keywords: sequential decision making; economics; Bayesian cognitive models; response time analysis

Acknowledgements

We thank Catherine A. Hartley for providing the intertemporal choice data. We also thank Michael D. Lee and Joachim Vandekerckhove for helpful discussions. This work was supported by NIA R21AG072673 and a NARSAD Young Investigator Award by the Brain and Behavior Research Foundation to AMB.

1 Introduction and Background

When making decisions under uncertainty and with limited time, humans and animals must balance efficiency with completeness. We have previously shown that humans exhibit 'spillover' between adjacent trials in non-sequential decision tasks[1]. Here, we examine whether this effect is modulated by the relative properties of the decision computations required on successive trials. Such effects could be the result of rational computational approaches – for instance if people are showing sequential effects in nonsequential environments, it could be because they are trying to infer the underlying Markov Decision Process (MDP) [4, 14]. While this idea has been studied extensively in lower-order decision making, it is a major open question whether individuals reference recent history in higher-order non-sequential decision tasks [2]. This is conceptually distinct from notions of motor or perceptual perseveration, which may themselves also be present during complex choice [1].

We therefore consider economic decision making in two non-sequential tasks: intertemporal choice (ITC) – decision making under temporal uncertainty – and risky and ambiguous choice (RAC) – choice under immediate known (risk) or unknown (ambiguous) uncertainty. A common method of measuring these constructs in behavioral economics is through experiments: typically, individuals make a sequence of choices in a randomized choice set and are explicitly instructed to treat each decision independently [8, 11]. The data is then analyzed as if each choice was indeed made independently and not in sequence. Our previous work suggests that this may be a flawed assumption [1].

First we consider the cognitive processes involved in making decisions about immediate or temporal uncertainty. Popular theories of intertemporal choice suggest that decisions are in part informed by the deliberative simulation of potential future outcomes [12]. These sorts of computations and simulations are expensive and time-consuming, so when successive decisions are made under time pressure and involve similar future dates and values it may be inefficient to construct preferences anew on each trial [3]. We may further expect such re-using of computations or, more broadly, reference to the past, as recent work suggests that neurons that code for value do so relatively (i.e. what has *changed* now compared to what was seen before) [16]. Thus, taken together, individuals may *act as though* there are structured temporal dependencies during decisions made in a sequence – even if those decisions are not explicitly related to each other.

The mechanisms that generate, or propagate these sequential effects are an open question even in perceptual and visual working memory, where such task in-congruent sequential behavior has been studied extensively [5, 9]. Candidate mechanisms include serial biases in attention [5] and efficient coding influences on working memory representations [9].

However, it remains unknown whether such 'spillover' effects apply to choice behavior more broadly and, if so, whether they are mediated by a general inclination to optimize speed/accuracy tradeoff in behavior, or by local uncertainty about task demands. To address this question, we leverage our previously introduced methodology to test and account for trial properties-driven sequential effects in such environments [1]. In our previous analysis, we observed that some sequential effects were driven by decision difficulty: they were present when the current trial was more difficult than the preceding. In particular, we found that a subset of subjects were sensitive to relative changes in Expected Value (i.e. when the previous trial had a large difference in Expected Value - "easy" - and the current trial had a small difference in Expected Value - "hard"). However, decision difficulty in ITC tasks, for example, is potentially conflated with the inherent uncertainty in receipt of future reward outcomes [7]. Here, we further extend our previous framework by re-parameterizing the stimulus space (all potential trial properties) in a simpler fashion and present a re-analysis of the intertemporal choice data in [1]. Then, we apply this approach to a newly collected risky and ambiguous choice task to test directly whether relative changes in decision difficulty or uncertainty more generally drive sequential effects. Across these datasets, we find evidence in favor of sequential effects, and in the RAC task, specifically that these effects are being driven by a tradeoff against choice difficulty, rather than by a need to resolve contextual uncertainty.

2 Methods

2.1 Experiments and Data

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). All data were collected previously (for details, see [8]). 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.

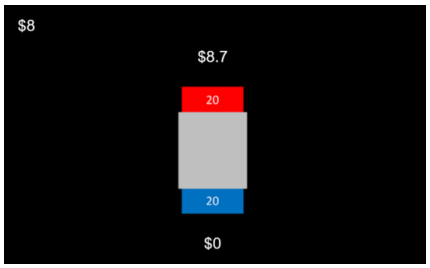

Figure 1: An example ambiguous trial from the RAC task, where the subject has 3 seconds from stimulus onset to make a choice between a certain reward of $8 and a chance to win $8.7 by playing the lottery.

Risky and Ambiguous Choice (RAC). We model $n = 98$ 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 1). The amount a subject could win by choosing the lottery was almost

always larger than the certain reward, except during 16 catch trials. All data were collected on Amazon Mechanical Turk. Lotteries were either *risky* (1/7) 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%). While we did not design any trial-level dependencies, we ensured that 50% of successive trials increased in ambiguity, and 50% decreased. A risk trial followed by an ambiguous trial is 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 3s after stimulus onset to make a choice.

No outcomes were realized over the course of either experiment (i.e no feedback after each choice other than a confirmation of the subject's selection). However, the experiments were incentive compatible and a single trial was selected at random, realized, and paid out at the end of the experiment as a bonus. For the RAC experiment, participants on Amazon MTurk were paid 10% of their winnings to be consistent with pay rates on the platform.

We exclude any responses that were made in less than 300ms. We also exclude any missed trials and the trial immediately after them from the following analyses.

2.2 Model

2.2.1 Baseline Models

We implement a hierarchical Bayesian drift diffusion model (DDM) to model response times using the Wiener module [15] in JAGS [13]. 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 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 RAC), 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 [10]. For α_i, τ_i, and β_i, we use the same prior and hyperprior specifications for both tasks, referencing [15] for mean hyperpriors and using 'noninformative' priors for the standard deviation:

$$\mu_\alpha \sim \textit{Uniform}(0.001, 3) \quad \mu_\tau \sim \textit{Uniform}(0, 0.6) \quad \mu_\beta \sim \textit{Uniform}(0.01, 0.99) \quad \sigma_\alpha, \sigma_\tau, \sigma_\beta \sim \textit{Uniform}(0.001, 4)$$

We take a cognitive psychometrics approach to modeling the drift rate by allowing it to be driven by (combinations of) trial properties. 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:

$$\delta_{ITC,ij} = \beta_{0,i} + \beta_{1,i} \cdot (\textit{value}_{LL,ij} - \textit{value}_{SS,ij}) + \beta_{2,i} \cdot \textit{delay}_{ij}^{-1} \tag{1}$$

$$\delta_{RAC,ij} = \beta_{0,i} + \beta_{1,i} \cdot (\textit{Expected Value}_{\textit{Lottery},ij} - \textit{Expected Value}_{\textit{Fixed},ij}) + \beta_{2,i} \cdot A_{ij} \tag{2}$$

In Equation 2, the Expected Value of a choice option is given by $EV = p \cdot v$, where p is the probability of reward and v is the monetary value. For a risky trial, p is the percentage chance of winning reward. For an ambiguous trial, $p = 0.5$ as in [11]. A represents the degree of ambiguity of a given lottery. On ambiguous trials, this is the fraction of the lottery that is occluded by the grey bar as seen in Figure 1. On risky trials, $A = 0$. Finally, we allow all drift rate decomposition parameters ($\beta_0, \beta_1, \beta_2$) to be hierarchical Standard Normals.

2.2.2 Sequential Effects

We build stimulus-driven sequential effects into each of the β terms on the drift rate decomposition and on the bias. Specifically, as we only consider influence of recent history that is transient (1-trial-back), we incorporate relative changes in stimulus values between the previous and current trial. For the ITC task, we consider successive trials that *increase* or *decrease* in (objective) value difference, delay difference, and value x delay difference. For the RAC task, we similarly consider value differences, ambiguity difference, and value x ambiguity difference. We use Indicator Variables to subset sequences of trials that follow any of the above specified patterns, resulting in a 8-fold tiling of stimulus space for both experiments. We then augment our baseline models by allowing these stimulus properties to exert linear additive influences on the parameters of interest. For example, if we consider trials that increase in value from trial $j-1$ to trial j:

$$\beta'_{0,ij} = \beta_{0,i} + \gamma_i \cdot 1([V_{a,j} - V_{b,j}] > [V_{a,j-1} - V_{b,j-1}]) \tag{3}$$

We allow all sequential effect parameters to be hierarchical Standard Normals. We simultaneously assess the influence of relative stimulus properties on all drift rate decomposition and bias parameters. Thus, in Equation 3, $\beta_{0,i}$ 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. To answer our core question of interest, we test whether the sequential effect parameters (e.g. γ_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 Results

3.1 Drift Rate Decompositions

We fit a variety of different drift rate decompositions and present DIC differences in two model comparisons between the decompositions presented in the *Methods* section (Equations 1 and 2) that may be of particular interest to the reader. For ITC, we report that our drift rate decomposition that allows an inverse relationship between delay and value, as generally given in economic models of temporal discounting, outperforms the decomposition that where both value and delay have directly proportional relationships with drift rate ($DIC_{ITC,\text{Inverse}} = 116886$, $DIC_{ITC,\text{Direct}} = 11875$). In RAC task, we report that the drift rate that only includes the Expected Value difference weight (β_1) performs more poorly than the drift rate that incorporates all trial properties $DIC_{RAC,\text{All}} = 21824$, $DIC_{RAC,\text{EV Only}} = 22134$).

ITC	Mean(CI)	RAC	Mean(CI)
β_0	$-0.66(-1.26, -0.06)$	β_0	$-0.72(-2.99, 0.37)$
β_1	$3.45(1.19, 5.70)$	β_1	$0.84(-0.87, 6.14)$
β_2	$2.34(-2.15, 6.84)$	β_2	$-0.02(-0.72, 0.67)$
α	$2.47(1.70, 3.24)$	α	$1.87(0.97, 2.47)$
τ	$0.76(0.36, 1.16)$	τ	$0.40(0.19, 1.14)$
$bias$	$0.51(0.39, 0.63)$	$bias$	$0.50(0.38, 0.59)$

Figure 2: Aggregate posterior means (95% Credible Intervals) for Drift Rate decompositions and other DDM parameters. Note that these values represent the spread across individuals.

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), all held constant. Both make sense, 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.

As in the ITC task, subjects appear to be sensitive to all trial properties in the RAC task. All else held constant, the average subject's drift rate increases as the Expected Value difference between choice options increase (β_1). Subjects are also sensitive to the degree of ambiguity, with drift rate decreasing with increasing ambiguity (β_2) (all else held constant). However, we highlight the credible interval ranges and note that there are considerable individual differences.

3.2 Sequential Effects

We observed reliable trial-property-driven sequential effects on almost all sequences, regardless of task. Indeed, for almost every possible combination of stimulus sequences, 100% of subjects showed evidence of sensitivity to previous stimuli in the ITC task. Consistent with our hypothesis, effects in the RAC task are more specific to the combination of trial properties considered – in particular, most individuals show sequential effects when there are relative increases and decreases in value.

ITC	Proportion	RAC	Proportion
$value \uparrow$	1	$value \uparrow$	0.80
$value \downarrow$	0.998	$value \downarrow$	0.93
$delay \uparrow$	1	$amb \uparrow$	0.05
$delay \downarrow$	1	$amb \downarrow$	0.01
$v. \uparrow d. \uparrow$	1	$v. \uparrow a. \uparrow$	0
$v. \uparrow d. \downarrow$	1	$v \uparrow a. \downarrow$	0.18
$v. \downarrow d. \uparrow$	1	$v. \downarrow a. \uparrow$	0
$v. \downarrow d. \downarrow$	0.84	$v. \downarrow a. \downarrow$	0.57

Figure 3: 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. $value \uparrow$ subsets successive trials that increased in value difference as noted in Equation 3). The top four rows can be thought of as "main effects" of specific trial properties and the bottom four "interactions."

In Figure 5, we list the posterior group-level means and standard deviations of the sequential effect parameters themselves for the RAC task only, omitting ITC results due to space constraints. We find broad individual variability in terms of which parameters seem to be most sensitive to sequential trial properties: while one subject may manifest this sensitivity on the drift rate intercept parameter, another may see the same on multiple parameters.

We show this more concretely in the ribbon plots Figure 4. Indeed, we find variability in the magnitude and directionality of sequential effects even when considering related sequences (value difference $\downarrow$ vs. value difference $\downarrow$ *and* amb. $\downarrow$). For subjects that showed non zero sequential effects on successive trials that increased or decreased in value difference regardless of ambiguity, we recovered these effects only on β_1, the Expected Value difference parameter on the drift rate. This serves as a sanity check of sorts, as reward value is explicitly weighted by β_1 through Expected Value. Similarly, for the handful of subjects that were only sensitive to relative differences in cross-trial ambiguity, we only recovered

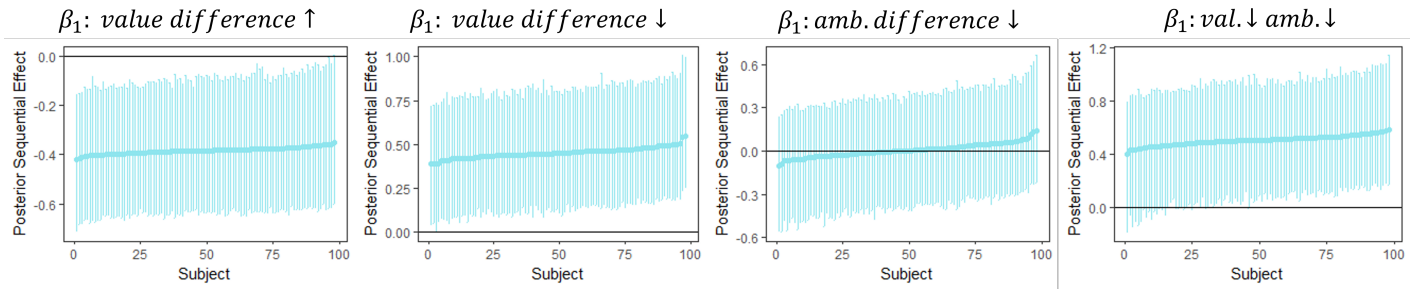

Figure 4: Sorted posterior 95% Credible Intervals of sequential effects on the drift rate Expected Value term (β_1) when successive trials (a) increase in value difference, (b) decrease in value difference, (c) decrease in ambiguity, and (d) decrease in both value difference and ambiguity difference. (Equation 2, RAC task).

effects on β_2, the weight of ambiguity on the drift rate. For decrease in value difference and decreases in both value and ambiguity difference, inferred Bayes Factors were near threshold ($BF_v = 2.97, BF_{va} = 2.9$).

4 Conclusion

Using a generative model and a simple parameterization of trial properties, we have demonstrated that individuals are influenced by recent history during economic decision making, even when it is task-incongruent to do so. We do not argue that this influence is a conscious one, but more specifically that these sorts of sequential influences are fundamental to human cognition and information processing, and thus ought to be explicitly accounted for. We find near-ubiquitous evidence of sequential effects in the ITC task. Our results from the RAC task showed more specific sequential effects and suggest that these effects are driven by relative differences in difficulty, and not necessarily uncertainty (ambiguity). We further show that in this task, explicitly incorporating trial properties into response time modeling via drift rates is fundamental to recovering these sequential effects (Figure 5).

Properties	β_0	β_1	β_2	*bias*
$value\uparrow$	H_0	$-0.38(0.12)$	H_0	H_0
$value\downarrow$	NEE	$0.45(0.11)$	NEE	H_0
$amb\uparrow$	H_0	NEE	$-0.29(0.07)$	H_0
$amb\downarrow$	H_0	H_0	$0.32(0.09)$	H_0
$v.\uparrow a.\uparrow$	H_0	NEE	H_0	H_0
$v\uparrow a.\downarrow$	H_0	NEE	$0.59(0.13)$	H_0
$v.\downarrow a.\uparrow$	H_0	NEE	NEE	H_0
$v.\downarrow a.\downarrow$	NEE	$0.5(0.16)$	H_0	H_0

Figure 5: Posterior sequential effect group means (γ_i) for Risk and Ambiguity Choice Task. All numerical means and standard deviations presented in a cell (*mean (standard deviation)*) have a BF > 3 that 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 "NEE" then the data does not contain enough evidence to favor either the null or alternative hypothesis (NEE: Not enough evidence).

We also note the value of simple parametrizations of trial properties: by explicitly modeling only relative increases or decreases, we were able to recover substantially more individuals who were sensitive to sequential effects. Thus, we have found strong evidence of effects that are likely widespread and therefore should be explicitly accounted for in all measures (Figure 3). Indeed, this could be one reason why task measures do not seem to correlate with each other or have as good test-retest reliability as other measures like surveys [6].

References

[1] N.V. Banavar, M.D. Lee, and A.M. Bornstein. Sequential effects in non-sequential tasks. *Proceedings of the 19th International Conference on Cognitive Modeling*.

[2] H.R. Brooks and P. Sokol-Hessner. Quantifying the immediate computational effects of preceding outcomes on subsequent risky choices. *Scientific reports*, 10(1):1–10, 2020.

[3] I. Dasgupta, E. Schulz, N.D. Goodman, and S.J. Gershman. Remembrance of inferences past: Amortization in human hypothesis generation. *Cognition*, 178:67–81, 2018.

[4] N.D. Daw, A.C. Courville, and P. Dayan. Semi-rational models of conditioning: The case of trial order. *The probabilistic mind*, pages 431–452, 2008.

[5] J. Fischer and D. Whitney. Serial dependence in visual perception. *Nature neuroscience*, 17(5):738–743, 2014.

[6] R. Frey, An. Pedroni, R. Mata, J. Rieskamp, and R. Hertwig. Risk preference shares the psychometric structure of major psychological traits. *Science advances*, 3(10):e1701381, 2017.

[7] X. Gabaix and D. Laibson. Myopia and discounting. Technical report, National bureau of economic research, 2017.

[8] L.E. Hunter, A.M. Bornstein, and C.A. Hartley. A common deliberative process underlies model-based planning and patient intertemporal choice. *bioRxiv*, page 499707, 2018.

[9] A. Kiyonaga, J.M. Scimeca, D.P. Bliss, and D. Whitney. Serial dependence across perception, attention, and memory. *Trends in Cognitive Sciences*, 21(7):493–497, 2017.

[10] M.D. Lee. Bayesian methods in cognitive modeling. *The Stevens' handbook of experimental psychology and cognitive neuroscience*, 5:37–84, 2018.

[11] I. Levy, J. Snell, A.J. Nelson, A. Rustichini, and P.W. Glimcher. Neural representation of subjective value under risk and ambiguity. *Journal of neurophysiology*, 103(2):1036–1047, 2010.

[12] J. Peters and C. Büchel. The neural mechanisms of inter-temporal decision-making: understanding variability. *Trends in cognitive sciences*, 15(5):227–239, 2011.

[13] M. Plummer et al. Jags: A program for analysis of bayesian graphical models using gibbs sampling. In *Proceedings of the 3rd international workshop on distributed statistical computing*, volume 124, pages 1–10. Vienna, Austria., 2003.

[14] F. Schuur, B.P. Tam, and L.T. Maloney. Learning patterns in noise: Environmental statistics explain the sequential effect. In *CogSci*, 2013.

[15] D. Wabersich and J. Vandekerckhove. Extending jags: A tutorial on adding custom distributions to jags (with a diffusion model example). *Behavior research methods*, 46(1):15–28, 2014.

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