A statistical test for the optimality of deliberative time allocation

Rahul Bhui¹

Published online: 25 January 2019
© The Psychonomic Society, Inc. 2019

Abstract

Whenever we make a choice, we must also decide how much time to spend making it. Many theories of decision-making crucially assume that this deliberation perfectly balances the costs of time expenditure and the benefits of better decisions. However, might we “overthink” or “underthink” decisions? Here, I propose and implement a method to precisely determine whether people are optimally spending their time on deliberation, accounting for individual preferences. This test evaluates the consistency of underlying preferences for time when incentives change, which is a necessary condition for optimality. This enables a more comprehensive analysis of rationality in a variety of contexts. I demonstrate how the test can reveal departures from optimality using two motion discrimination experiments in which I vary task difficulty and monetary incentives.

Keywords Decision making · Response time · Optimality · Speed-accuracy tradeoff · Sequential sampling models

When we decide whether to leap left or right to return a tennis serve, we expend time evaluating the options. The purpose of this deliberation is to reduce uncertainty as to which option is better, and thereby improve decision quality. Accordingly, choice often involves an inverse relationship between speed and accuracy; we can make judgments that are fast but error-prone, or slow but high quality. Theories based on utility maximization predict that people optimally balance the costs of spending time with the benefits of better decisions. I propose a new test of this prediction that more closely targets the theoretical core of optimality compared to past attempts, and I conduct two experiments to show how suboptimality can be accordingly investigated in deliberative time allocation.

A long-standing body of work in psychology involves measuring response time in order to improve inferences about cognitive processing and predictions about behavior (Donders, 1869/1969; Hick, 1952; Sternberg, 1966; Shepard & Metzler, 1971; Jensen, 2006), and research in

economics has followed suit (Wilcox, 1993; Rubinstein, 2007; Chabris, Morris, Taubinsky, Laibson, & Schuldt, 2009; Rand, Greene, & Nowak, 2012; Spiliopoulos & Ortmann, 2018; Clithero, 2016, 2018). To rigorously apply these insights, researchers have developed computational models that jointly incorporate stochastic choice and time expenditure (e.g., Ratcliff, 1978; Busemeyer & Townsend, 1993; Woodford, 2014; Fudenberg, Strack, & Strzalecki, 2018; Webb, 2018). Most of these models are motivated by principles of optimality. They lay out mathematical descriptions of noisy information processing over time and posit decision rules that are optimal subject to this processing. Such models have been shown to capture important features of individual behavior relating to the joint distribution of time and accuracy as well as patterns of neural activity (Gold & Shadlen, 2007; Ratcliff & McKoon, 2008). Although these theories have enjoyed extensive study, their foundations in utility maximization and portability across environments have remained less well studied empirically.

Do people optimally balance the costs and benefits of time spent and reward attained? When conditions change, does behavior change commensurately? The answers to these questions are important because they inform us about how we can best generalize, predict, and influence behavior across various contexts. If people are choosing optimally, we can predict their actions precisely using optimization models even when task demands move outside of their original confines. If, on the other hand, people are not choosing optimally, then alternative models must be sought to achieve

Electronic supplementary material The online version of this article (<https://doi.org/10.3758/s13423-018-1555-1>) contains supplementary material, which is available to authorized users.

✉ Rahul Bhui
rbhui@g.harvard.edu

¹ Departments of Psychology and Economics & Center for Brain Science, Harvard University, 52 Oxford St, Cambridge MA 02138, USA

better predictions, and there may be room for interventions to improve the efficiency of decision-making. For instance, time limits could prevent people from spending excess time on chronic deliberation between similar courses of action where deliberating yields little return (Krajbich, Oud, & Fehr, 2014; Oud et al., 2016).

The answers are also challenging to determine because costs and benefits are subjective. How one feels about spending time towards some end varies from person to person. To fully explore individual decision-making, we must allow for individual preferences. Lacking ways to overcome this challenge, past attempts to test optimal time allocation were based on alternative criteria that overlooked subjectivity or only coincided with utility maximization in special cases. Research on the portability of deliberative choice models has neglected this crucial element, resulting in tests that have not directly addressed the theory.

In this paper, I propose a flexible way to test expected utility maximization in stochastic time allocation. The approach consists of a hypothesis test for individual preference consistency across different environments: when the objective costs and benefits of deliberation change, time expenditure should change in such a way as to reflect consistent subjective preferences for time versus reward. This test can be applied in a variety of scenarios spanning perceptual and value-based decision-making. In tandem, I conduct online experiments in the perceptual domain to illustrate how the test can be used to investigate optimality according to the drift diffusion model, one of the most popular theories combining stochastic choice and time allocation. The experiments include manipulations of task difficulty and of monetary incentives. I find that nearly half of the participants do not appear to respond optimally to changes in task difficulty, and the majority of participants do not appear to be sensitive to changes in monetary incentives.

Optimality and sequential sampling models

A strain of stochastic choice models originating in the mid 20th century formally connect response time and choice by detailing the process of deliberation (Stone, 1960; Laming, 1968; Ratcliff, 1978). These models, known as sequential sampling or bounded accumulation models, describe decision-making as one or more stochastic processes representing the evolution of confidence in one's answer as noisy information is processed. Once confidence reaches a particular threshold level, the agent stops processing information and commits to a choice. The oldest and most famous model in this class, the drift diffusion model, assumes the confidence trajectory follows a Wiener process with drift. It has been found to closely match configurations of choice and response times in a variety of tasks ranging from perceptual

discrimination (Ratcliff & Rouder, 2000, 2002; Smith, Ratcliff, & Wolfgang, 2004) to recognition memory (Ratcliff, 1978; Starns & Ratcliff, 2014), and more recently, value-based judgment (Krajbich, Armel, & Rangel, 2010; Milosavljevic, Malmaud, Huth, Koch, & Rangel, 2010; Krajbich & Rangel, 2011; Krajbich et al., 2012, 2015). Moreover, direct measurements of neural activity reveal the implementation of evidence accumulation processes that fit the model's structure (Hanes & Schall, 1996; Shadlen & Newsome, 2001; Gold & Shadlen, 2002; Ratcliff, Cherian, & Segraves, 2003; Smith & Ratcliff, 2004; Gluth, Rieskamp, & Büchel, 2012; Green, Biele, & Heekeren, 2012).

Part of the drift diffusion model's original motivation was its formal analogy with efficient statistical algorithms. The model was built as the continuous-time limit of the sequential probability ratio test (Wald, 1947) and is often theoretically characterized as inheriting its optimal stopping properties. For instance, it attains the speed–accuracy frontier (Wald & Wolfowitz, 1948; Arrow, Blackwell, & Girshick, 1949), meaning that it achieves the highest accuracy for any given response time and the quickest response time for any given accuracy level (Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006). The decision threshold governs the exact tradeoff between speed and accuracy; when a higher standard of confidence is demanded, choices become more accurate but take more time. Theoretical derivations of the drift diffusion model as well as some of its extensions assume the threshold is selected to optimally manage the tradeoff between the cost of time and the benefit of reward (e.g., Laming, 1968; Fudenberg et al., 2018). This means the threshold and induced behavior should vary across environments in a precise manner tied to costs and benefits (Bogacz et al., 2006). Estimated thresholds and corresponding brain activity do sometimes respond to information and incentives (e.g., Domenech & Dreher, 2010; Green et al., 2012; Gluth et al., 2013) but whether these responses are precisely optimal remains to be seen.

To be more specific, in a forced choice between two alternatives, the key outcomes which decision makers are thought to consider are time spent and performance attained. The former is penalized due to objective or subjective costs of time, and the latter is rewarded by association with some payoff. In the special but commonly encountered case that occurs when the rewards from the correct and incorrect options are fixed, performance reduces to accuracy, the probability of choosing the correct option. Two criteria for optimality have been used in the neuroeconomic literature based on these central elements.

The first criterion is to minimize a weighted sum of error rate (ER) and decision time (DT), which is known as the Bayes risk (Wald & Wolfowitz, 1948):

$$\min_{x \in X} BR(x; \theta) = rER(x; \theta) + \psi DT(x; \theta), \quad (1)$$

where x and θ are choice variables and fixed parameters, respectively, that determine the outcomes, and r is the reward for a correct answer. In the simple two-parameter drift diffusion model that I focus on, for example, x consists of the decision threshold and θ consists of the drift rate parameter (relative to noise) that summarizes the rate of information processing. The optimal choice of x depends not only on the reward r for a correct answer but also on the free preference parameter ψ that determines how much importance is placed on a unit of time relative to a unit of reward, and which may include a sizable subjective component that is not directly observable.¹ Thus ψ is interpretable as the subjective flow cost of time.² The Bayes risk expression is a special case of expected utility when the reward depends only on whether the response is correct. More generally, the decision-maker is assumed to solve

$$\max_{x \in X} \mathbb{E}[\text{reward}(x; \theta) - \text{cost} \times \text{time}(x; \theta)]. \quad (2)$$

This is widely considered to be the most universally appropriate optimization criterion.

The second criterion is to maximize the ratio between accuracy and decision time, which is the reward rate (Gold & Shadlen, 2002):

$$\max_{x \in X} RR(x; \theta) = \frac{r(1 - ER(x; \theta))}{DT(x; \theta) + T_0 + D + ER(x; \theta) \times D_p}, \quad (3)$$

where T_0 is the time required for sensory and motor processing, D is the time interval between a correct response and the following stimulus, and D_p is the additional time delay which penalizes an incorrect response on top of D . This expression is relatively more common in psychology and ecology due to its origins in reinforcement rate analysis, but remains little used outside of those fields. The optimal choice here does not require the specification of any free parameters.

If the environment is homogeneous in difficulty, then the sequential probability ratio test (and drift diffusion model) maximizes any reward criterion based on error rate and decision time that is monotonically nonincreasing in decision time (Bogacz et al., 2006). Thus the Bayes risk and reward rate optimal solutions happen to coincide in certain aspects (Moran, 2015). However, even in the homogeneous setting, predicted changes in behavior across different environments will vary depending on the structure of the model and which of its components are assumed to be invariant. Because reward rate maximization is

nominally parameter-free, empirical tests of optimal choice have focused almost exclusively on this criterion (Simen et al., 2009; Bogacz, Hu, Holmes, & Cohen, 2010; Starns & Ratcliff, 2010; Starns & Ratcliff, 2012; Zacksenhouse, Bogacz, & Holmes, 2010; Balci et al., 2011; Karşılar, Simen, Papadakis, & Balci, 2014; Drugowitsch et al., 2012, 2014, 2015). Optimality in the utility maximization sense has thus been empirically neglected despite its theoretical importance. Krajovich et al. (2014) and Oud et al. (2016) are among the few to have conducted more direct tests of utility maximization. They found that participants under time pressure spent longer on low-stakes compared to high-stakes trials, and earned more when given per-trial deadlines. This evidence of irrationality, while convincing, is based only on coarse benchmarks of optimality. Participants could spend more time when stakes are high and still be violating expected utility theory.

Testing optimality via preference consistency

I seek to redress the imbalance by introducing a method to test expected utility maximization in stochastic decision settings. The test comprises a check for consistency of preferences across different conditions. The idea is that when changes occur in environmental parameters such as problem difficulty or external incentives, optimal behavior should change commensurately to reflect consistent underlying subjective preferences for a unit of time relative to a unit of reward. Although past work studying sequential sampling models has identified preferences with objective costs and benefits, this is an assumption maintained for convenience and not a core element of optimality. To attain a given reward, some individuals might be willing to spend quite a while even when others remain reluctant to part with their time. The presence of individual differences does not necessarily imply irrationality, any more than would the simultaneous existence of one person who intrinsically prefers apples to oranges and another person who has the opposite predilection.

The method proposed in the present paper finds its basis in economic tests of revealed preference (Samuelson, 1938; Houthakker, 1950; Afriat, 1967; Crawford & De Rock, 2014). If an analyst wishes to avoid making extraneous claims about the inherent quantitative sensibility of possible preferences (such as to what extent a decision maker values an apple versus an orange, or a unit of time versus a unit of reward), the most foundational constraint that can be imposed is internal consistency. Thus the question asked is: do there exist consistent preferences that can account for the observed behavior? Alternatively stated, can the behavior be *rationalized*? It must be stressed that such consistency is not a tangential property of rational choice but its cornerstone, to the point that the two are virtually

¹The reward amount r should not be conflated with the preference parameter ψ . Reward r is denominated in dollars while the value ψ of (a unit of) time is denominated in dollars per hour.

²The value of time is theoretically its opportunity cost, thus I refer to the “value” and “cost” of time interchangeably.

equated in economic analysis. For example, Binmore (2009) states that “the words *rationality* and *consistency* are treated almost as synonyms in much modern work,” and Gintis (2009) even declares that “a *rational actor* is an individual with *consistent preferences*.” The approach proposed here thus strikes directly at the heart of optimality.

To explain how the test works, suppose data is available from two sets of problems, one with uniformly low difficulty and one with uniformly high difficulty, similar to Experiment 1. A person’s ability is likely lower among the high difficulty problems. The optimal decision rule (captured in the drift diffusion model by the confidence threshold) in each condition will be based on their ability and their preference for speed versus accuracy. If we have reason to believe their underlying preference should not vary across these conditions (or at least that it changes in line with some specified relationship), then the estimated decision rule in conjunction with the estimated ability parameter in each condition should imply the same preference parameter across conditions. This forms the core of the test. While estimates of ability and decision rule will vary due to task difficulty, they should jointly indicate the same inferred preference if people are indeed optimally balancing the costs and benefits according to expected utility maximization.

Or suppose, as in Experiment 2, that we are interested in studying whether people can optimally adjust their deliberative behavior according to possible earnings. Although problems of varying rewards may lead to different decision rules, these differences should nonetheless be generated by a similar underlying (subjective) tradeoff. This consistency can be checked by estimating preference parameters across different sets of problems characterized by payoffs, and testing whether fundamental preferences are the same even when payoffs change. It is important to not conflate these deeper preferences characterized by ψ with the superficial behavioral “preference” for spending time on the task characterized by the decision rule x (e.g., the decision threshold). Even though the choice of x may change when the reward amount r is altered, the newly selected x should continue to reflect the same underlying preference for a given unit of time relative to a given unit of reward as encapsulated in ψ . To use an analogy, a rational agent may buy relatively more apples than oranges when the price of apples decreases relative to oranges, but this behavior does not mean that how much they intrinsically like apples versus oranges has changed; what has changed is merely how the same underlying preference translates into behavior because of the different constraints faced by the agent.

More formally, the analysis can be implemented by a general hypothesis testing approach that relies on a combination of highly flexible tools, and may thus be applied to a wide range of models. The key question

embodied by the test is whether the preference parameter ψ remains constant across environmental conditions. Model parameters—including the preference parameter—can be estimated using maximum likelihood estimation under different assumptions. To determine whether there is sufficient evidence of variation to reject consistency, the model fit assuming a single value of ψ across conditions can be compared to the fit allowing ψ to be free in each condition. The difference can be assessed with a likelihood ratio test between the restricted (single ψ) likelihood and the unrestricted (condition-dependent ψ) likelihood.

This approach is feasible as long as the model makes specific predictions about the joint distribution of choice and time expenditure, even if that distribution is analytically intractable (since it can be simulated). As long as preferences can be identified using some form of maximum likelihood estimation, consistency can be assessed using a likelihood ratio test where the restriction comes from the specified relation between preferences. In particular, if preferences are believed to be the same across multiple conditions, the restricted likelihood is based on equality between the preference parameters. Alternatively, weaker restrictions based on inequalities can be made if one only wants to assume ordinal relationships between conditions.

The approach has natural advantages and limitations. Because it is a consistency check, data is needed from multiple distinct but comparable conditions such that the analyst believes in some specific relationship between preferences across environments. Since preferences are allowed to be idiosyncratic, this imposes a meaningful constraint on when the test can be applied in practice. Moreover, model fitting may be computationally intensive due to the nonstandard transformations and restrictions involved. However, the test is flexible and can be applied in a general form with any model that makes precise predictions about decision time and performance. It can be implemented using numerical methods, which is useful because the class of sequential sampling models does not typically permit analytical characterizations of its key properties. For clarity, in this paper I will restrict my attention to the pure drift diffusion model due to its analytical tractability.

Optimality in the pure drift diffusion model

The drift diffusion model (DDM) formally describes the process of choice as the stochastic accumulation of net evidence hitting an absorbing boundary. In its basic form, the model consists of four elements:³ the difference in

³There is often assumed to be a period of time after stimulus onset but before evidence accumulation, captured by a non-decision time parameter. This parameter does not affect the optimization problem if it remains fixed across conditions. Non-decision time will be included in the model fitting later on.

evidence between the two alternatives at a point in time ($x(t)$), the drift rate that captures the agent's speed in integrating evidence (A), the noise in the accumulation process (c), and the threshold that stops the accumulation process ($\pm z$). Evidence for an option is integrated noisily in continuous time as represented by a Wiener process,

$$dx = A dt + c dW, \quad (4)$$

that begins at $x(0) = 0$ in an unbiased decision, and stops as soon as it hits the decision threshold $\pm z$, at which time a choice is made according to whether $+z$ or $-z$ was hit. This threshold governs the speed–accuracy tradeoff. If it is set high, then the standard of evidence required to make a decision is stringent, and the resulting decision will be slow but accurate. If it is set low, then only weak evidence is needed, and the decision will be fast but inaccurate. This confidence threshold is chosen depending on one's preferences for time and accuracy; if optimal, it solves the Bayes risk minimization problem.

The properties of this model are well understood. Mathematically, its predicted outcomes arise as solutions to a first passage problem in stochastic processes—when does a Wiener process with drift first hit an absorbing boundary, and which boundary does it hit? Analytical characterizations of accuracy and time contingent on its parameters are available due to its tractability (Bogacz et al., 2006).⁴ The error rate is given as

$$ER = \frac{1}{1 + e^{2Az/c^2}}, \quad (5)$$

and the mean decision time is given as

$$DT = \frac{z}{A} \tanh\left(\frac{Az}{c^2}\right). \quad (6)$$

The optimal decision threshold z^* is thus

$$z^* = \arg \min_{z \in \mathbb{R}_{++}} BR(z; A, c, \psi) = \arg \min_{z \in \mathbb{R}_{++}} rER(z) + \psi DT(z), \quad (7)$$

and is the solution to

$$\frac{r}{\psi} \frac{2A^2}{c^2} - \frac{4Az^*}{c^2} + e^{-(2Az^*/c^2)} - e^{2Az^*/c^2} = 0. \quad (8)$$

Although this does not produce a closed form solution, it is known to uniquely determine z since the equation can

Fig. 1 The optimal DDM threshold (z^*) for different drift rates (A) and preferences for time (ψ), according to the Bayes risk criterion. The optimal threshold is nonmonotonically related to drift rate, because the benefit of deliberating is low when the drift rate is low, and the need to deliberate is low when the drift rate is high. When time has low value, the decision-maker is more willing to spend time deliberating, and has a raised threshold

be written as an equality between increasing and decreasing functions. The properties of the optimal threshold are, however, non-trivial; for instance, it depends nonmonotonically on the drift rate as illustrated in Fig. 1. When the drift rate is low, noise dominates and the value of deliberation is low. At intermediate levels, enough information becomes available to make deliberation valuable. When the drift rate is high, even quick judgments contain plenty of information so there is no need to continue accumulating evidence, and the optimal threshold converges to zero. The preference parameter modulates the location and shape of the resulting curve. Figure 1 also illustrates how a given threshold–drift-rate pair uniquely identifies the preference parameter (for some given reward amount). Thus we see that the interpretation of parameters in terms of optimality is complex, underscoring the need for formal analysis.

The optimality test centers on checking for consistency of individual preference parameters across environmental conditions (such as difficulty level). To do so, under the assumption of optimality the DDM can be equivalently recast in terms of A and ψ (instead of z) using Eq. 8. These parameters can be estimated using maximum likelihood estimation either forcing ψ to be the same across all conditions or allowing it to vary.⁵ These constitute the restricted and unrestricted models, respectively, that are entered into a likelihood ratio test with the null hypothesis $\psi_i = \psi_j$ for all conditions i and j . Intuitively, this

⁴The properties of the simple DDM depend only on the ratios z/c and A/c , so as is typically done, c will be fixed to 0.1 in what follows, and z and A will refer to the ratios when c is not mentioned.

⁵The ability parameter A may be free across conditions in both cases, and indeed this must be permitted when conditions are characterized by changes in task difficulty.

Fig. 2 The distribution of p values resulting from the optimality test in Experiment 1, in which task difficulty varied across conditions. Optimality was rejected at the 5% significance level for eight out of 20 participants

is checking that the subjective cost of time implied by behavior is the same across conditions.

Experimental tests of optimality

I conduct two online experiments that vary different task parameters to illustrate use of the test. First, the decisions we make often vary in their difficulty level. Some have obvious answers and can be made quickly while others are unclear and may become drawn-out. Do people optimally adjust their deliberative behavior according to the difficulty of decision problems? Experiment 1 varies task difficulty which alters the ability parameter A . Second, the possible rewards for answering correctly may be small or large. High-value trials rationally demand more time than low-value trials. Do people optimally adjust their deliberative behavior according to the value of decision problems? Experiment 2 varies the reward r for a correct answer.⁶

I study these questions in a common perceptual judgment paradigm, the random dot motion task (e.g., Newsome, Britten, & Movshon, 1989; Britten, Shadlen, Newsome, & Movshon, 1992) in which behavior and brain activity have been shown to fit the structure of the drift diffusion model. In each trial of the random dot motion task, 100 white dots are displayed on a screen. A large fraction of them are “noise dots” moving in random directions, while the remaining few are “signal dots” moving in a consistent direction—either all to the left or all to the right. The agent must determine in which direction the signal dots are moving.

⁶As discussed earlier, this should not be confused with changing the preference parameter ψ itself.

Fig. 3 The estimated value of time (ψ) for each participant across the two conditions in Experiment 1. Optimal behavior is characterized by stability of the preference parameter when task difficulty changes, depicted as the identity line. The optimality test detects those who are statistically far enough away from this benchmark

Experiment 1

Participants

Twenty-four participants (12 female, two unreported; age range, 23–62 years) were recruited from Amazon Mechanical Turk. They were paid a \$1 participation fee in addition to earnings for performance as described below. The experiment was approved by the Harvard Committee on the Use of Human Subjects and all participants gave informed consent before the start of the task.

Materials

Participants faced 200 trials of random dot kinematograms each with 100 white dots moving on a black background. Task difficulty was determined as usual by the number of dots moving in a consistent direction (i.e., the coherence). Trials were split evenly between two difficulty levels defined by coherences set at 10% or 20%.⁷ These were grouped into 20 blocks each comprising 10 trials of the same difficulty level, with blocks ordered randomly. This structure preserved a counterbalanced design while relieving participants of the need to constantly adjust their strategy across individual trials. The difficulty level was displayed in large bold font before each block for 2s (as “Easy” for 20% coherence and “Hard” for 10% coherence)

⁷These coherence values were selected so that most participants were willing and able to attain intermediate accuracy levels in both conditions, to avoid ceiling and floor effects which might complicate the analysis. This was determined by pilot testing with other participants from the same population.

Fig. 4 The estimated drift rate (A) for each participant across the two conditions in Experiment 1. Drift rates decrease for everyone when difficulty is high

and was continuously displayed during each trial in case participants missed or forgot the information.⁸ Feedback on the correct direction was presented immediately after each trial for 1s in order to facilitate optimal behavior. Inter-trial intervals were fixed at 500ms. Participants were paid a fixed amount of \$0.03 for each correct answer and nothing for wrong answers, on top of the \$1 participation fee. The direction of coherent motion was determined with equal probability randomly across blocks, and participants were informed of this. The task was programmed in jsPsych (De Leeuw, 2015) using a random dot kinematogram plugin (Rajananda, Lau, & Odegaard, 2018).⁹ Six practice trials at various coherence levels (two at coherence 80%, two at 20%, and two at 10%) were shown before the regular trials to help explain the task. At the end of the experiment, participants completed a questionnaire in which they were asked about their task strategies and motivations.

Results

Models were fitted using quantile maximum probability estimation, with response time bins defined by the 10th, 30th, 50th, 70th, and 90th percentiles (Heathcote et al., 2002, 2004; Brown & Heathcote, 2003). Distributions of choices and response times were simulated using a random walk approximation with time step 50ms and 100,000

⁸Since participants were informed of the difficulty level before trials began, this design does not violate the homogeneity required for the DDM's constant threshold to be optimal.

⁹The kinematogram had an elliptical aperture of 400×600 pixels, and each dot had a radius of 3 pixels, a velocity of 1 pixel per frame, and a lifetime of 60ms. Participants hit the “a” key to signify “left” and the “l” key to signify “right.”

Fig. 5 The distribution of p values resulting from the optimality test in Experiment 2, in which the payoff for a correct answer varied across conditions. Optimality was rejected at the 5% significance level for 18 out of 20 participants

replicates (Tuerlinckx, Maris, Ratcliff, & De Boeck, 2001), and optimization was carried out via differential evolution (Mullen, Ardia, Gil, Windover, & Cline, 2011).¹⁰ As depicted in the [Supplementary Materials](#) Figs. A1, A2 and A3, the estimated models provide a very good fit to the empirical response time distributions and accuracy data. In the following analyses, four individuals are excluded due to performance statistically indistinguishable from chance. Among the remaining individuals, trials that took longer than 20s are excluded, comprising 0.2% of trials. Mean accuracy was 80.4% (range, 64–92%) and mean response time was 2.48s (range, 0.95–5.43s).

The optimality test was implemented for each individual in the sample. The histogram of resulting p values is displayed in Fig. 2. Optimality was rejected at the 5% significance level for eight out of 20 people (four at the 1% level, two at the 0.1% level).¹¹ Figures 3 and 4 display the parameters for each individual estimated according to the unrestricted model (i.e., both parameters allowed to be free in each condition). Each point represents the parameters of a single participant in the easy condition (x dimension) and the hard condition (y dimension). Points near the identity line in Fig. 3 reflect preference parameters that are stable (i.e., consistent with optimality). Figure 4 reveals that drift

¹⁰Due to the customization required, this procedure was computationally intensive, taking on the order of 1000 core-hours to fit the models for 20 individuals. A time step of 10ms was used for two individuals with especially short response times.

¹¹It may be that participants gain consistency during the course of the experiment, particularly since they must learn the task parameters through experience. Attempting to assess this possibility, I conduct the test separately on the first and second halves of the experiment. Due to the reduction in power, optimality is rejected at the 5% level for four people in the first 100 trials and two people in the last 100 trials.

Fig. 6 The estimated value of time (ψ) for each participant across the two conditions in Experiment 2. Optimal behavior is characterized by stability of the preference parameter when payoffs change, as depicted by the identity line

rates are lower for nearly everyone when the task is hard, including those individuals who are not picked out by the test. Thus, the test does not appear to be merely detecting any kind of change in information processing due to task difficulty (though more extreme variation in task difficulty could help produce evidence of suboptimality). Rather, it is selectively pinpointing violations of preference consistency.

As mentioned earlier, the test requires the relationship between preference parameters across conditions to be specified. Consistency in the present case was taken to mean that the value of time should be identical across difficulty levels; that is, $\psi_{\text{Hard}} = \psi_{\text{Easy}}$. However, it may be plausible that time spent on hard problems feels

Fig. 7 The estimated drift rate (A) for each participant across the two conditions in Experiment 2

subjectively worse than the same amount of time spent on easy problems. Superficially in line with this notion, every individual spent less time on hard trials (mean time 2.02 s vs. 2.94 s). More directly, in response to a post-experiment questionnaire, six individuals indicated that they found easy blocks to be “more fun” (two individuals said they found hard blocks more fun, while the remaining 12 felt no difference). In this case, a wider range of behavior would be rationalizable, as certain preference changes could still reflect an underlying consistency. While this aspect of subjectivity is worth further study, it does not explain the violations of optimality observed in the present experiment. If hard trials were subjectively worse than easy trials, this would raise the estimated cost of time in hard trials—but all individuals who were rejected by the test exhibit the opposite pattern.

I note furthermore that the test can be adapted to such alternative criteria while maintaining the same general approach, by relaxing the equality restriction. Rather than assuming $\psi_{\text{Hard}} = \psi_{\text{Easy}}$, one could enforce an order-restricted hypothesis, $\psi_{\text{Hard}} \geq \psi_{\text{Easy}}$. Theoretical results imply that the likelihood ratio test statistic is then usually distributed as $\bar{\chi}^2$, which is a mixture of χ^2 distributions with varying degrees of freedom (e.g., Robertson, 1978; Robertson, Wright, & Dykstra, 1988).¹² In practice, the mixture weights for each component of the $\bar{\chi}^2$ distribution are difficult to calculate, but its critical value will be bounded by the χ^2 critical value with equality-restricted degrees of freedom, which can be used as a conservative criterion.

Although participants might have been less willing to spend time on hard trials due to subjective effort costs, this cannot by itself account for observed behavior. Perhaps individuals placed a fixed personal premium on accuracy, which is psychologically plausible. Indeed, they were asked about their strategy on a post-experiment multiple choice question which allowed them to select one of “I focused more on getting the direction correct and taking my time,” “I tried to go as quickly as possible,” and “neither, I tried to balance these things.” The first of these seems consistent with an accuracy premium, and ten participants chose this option (the remaining ten chose the last option). However, the data also cannot be explained by this idea within the maintained framework. Such a premium would already factor into the estimated preference parameter equivalently across conditions, because the parameter summarizes the value of time expenditure relative to the value of a

¹²Intuitively, the $\bar{\chi}^2$ distribution formalizes what can be thought of as a χ^2 distribution with non-integer degrees of freedom. Consider that restricting two parameters to be equal to each other removes a single degree of freedom; it follows that restricting one parameter to be weakly greater than the other is less stringent and should in some sense remove only a fractional degree of freedom.

correct decision. The estimated parameters thus reveal some excess willingness to work on hard trials (or insufficient willingness to work on easy trials) beyond that induced by the value of a correct choice, intrinsic or extrinsic.¹³ If such subjective preferences did play a role, they must have entered into the mental calculation differently than prescribed by simple expected utility theory.

Experiment 2

Participants

Twenty-four participants (who did not take part in Experiment 1; six female, six unreported; age range, 19–57) were recruited through Amazon Mechanical Turk. They received a \$1 show-up fee in addition to a bonus payment for performance as described below. The experiment was approved by the Harvard Committee on the Use of Human Subjects and all participants gave informed consent before the start of the task.

Materials

Participants engaged in 200 trials of the random dot motion task constructed from the same basic elements as in Experiment 1. However, here the coherence was maintained at a constant 10% (the “Hard” difficulty level in Experiment 1), which participants were informed of. In half of the trials (in the same block structure as Experiment 1) the payoff for each correct answer was \$0.01, while in the other half the payoff was \$0.05. In all trials, nothing was earned for incorrect answers. The potential reward was displayed in large bold gold-colored font before each block for 2s, was continuously displayed during each trial, and was shown as part of the feedback as well. As in Experiment 1, feedback was provided after each trial for 1s, inter-trial intervals were fixed at 500ms, and six practice trials were shown.

Results

Models were fitted in the same way as Experiment 1.¹⁴ As shown in the [Supplementary Materials](#) Figs. A4, A5 and A6, the model predictions fit the empirical response

¹³Notwithstanding the potential impact of alternative preferences, monetary rewards did seem to be a prime motivator here. On a post-experiment survey question asking how motivated participants were to win (performance-based) bonus money, 19 out of 20 said they were “very motivated” (the highest rating on the five-point Likert scale) and the remaining individual was “somewhat motivated” (the second highest rating).

¹⁴A time step of 10ms was used for all but two individuals (for whom a 50ms time step was used) who had relatively long response times in both conditions. Again, the procedure was computationally intensive, taking over 6000 core-hours to fit the models for 20 individuals.

time distributions and accuracy data quite well. Four individuals were excluded from the following analyses due to performance at chance level or for having shorter response times in the high-payoff condition. Of the remaining 20 participants, 16 exhibited a statistically significant difference (increase) in response times across conditions according to a nonparametric Mann–Whitney *U* test at the 5% level; thus, the payoff manipulation did seem to successfully influence behavior.¹⁵ Trials longer than 20s were excluded from analysis (comprising 0.2% of trials). Mean accuracy was 69.8% (range, 60–79%) and mean response time was 2.20 s (range, 0.57–4.28 s).

Assuming the value of the reward *r* to be its objective monetary amount (either \$0.01 or \$0.05), optimality was rejected at the 5% significance level for 18 out of 20 individuals (16 at the 1% level, 13 at the 0.1% level), displayed in Fig. 5.¹⁶ When the reward *r* changes by a factor of 5, time allocation should rationally change as well to a substantial degree; for instance, the individual with the most consistent preferences increased their mean response time from 1.05 s to 4.56 s. However, nearly everyone else underreacted to the altered payoffs; participants on average increased their response times from 1.66 s to only 2.74 s. If people are presumed to maximize utility (i.e., they are presumed to react perfectly to their subjective preferences), then “underreaction” (or “overreaction”) is accounted for by an artificial change in the subjective cost of time (because this is the free utility parameter). In this case, the inferred ψ parameter is forced to contort itself to explain this insensitivity within the utility maximization framework, and this contortion is what the test is designed to detect. It is depicted in Fig. 6 as instability in the estimated preference parameters across payoff regimes. Figure 7 reveals little obvious pattern in the drift rate parameter across conditions.

Stickiness here could conceivably reflect some kind of inflexible bias in decision-making, in contrast to Experiment 1. For example, people might place intrinsic value on making the correct choice independent of any monetary payoff. Roughly consistent with this idea, on a post-experiment multiple choice question, six participants said that they “focused more on getting the direction correct and taking [their] time” (13 indicated that they “tried to balance [speed and accuracy],” and the remaining one said they “tried to go as quickly as possible”). Although the above analysis assumes that benefits derive only from

¹⁵Several participants also explicitly reported in comments that they strategically spent more time on the 5-cent blocks.

¹⁶As before, attempting to see whether experience helps participants behave closer to optimal, I split the data into the first and second halves and conduct the test on each half. Optimality is rejected at the 5% level for 13 participants in the first half and 15 participants in the second half. Experience thus does not appear to have much effect.

monetary rewards,¹⁷ it is important to recognize that this is largely for practical convenience and can be relaxed in some cases. The approach laid out in this paper is capable of accommodating alternative preference structures if the number of conditions is increased to compensate for the added degrees of freedom. Though beyond the present scope, understanding what type of intrinsic preferences help account for apparent deviations from optimality is an interesting topic for future study (e.g., Bhui 2018). While it is known that the speed–accuracy tradeoff can be influenced by instructions emphasizing speed or accuracy (e.g., Ratcliff & McKoon, 2008) more remains to be understood about subjective preferences in this context.

General discussion

Time expenditure has proven to be a useful predictor of choice behavior, and sequential sampling models have become standard tools to analyze deliberative time allocation. Part of their appeal lies in the optimal properties they inherit from efficient statistical rules (Wald & Wolfowitz, 1948; Arrow et al., 1949; Bogacz et al., 2006). The extent to which observed deliberation can be grounded in optimization is an open question, but most attention to date has been focused on criteria that are not universally appropriate definitions of optimality. I lay out a method to test expected utility maximization in these settings that accommodates flexible characterizations of subjective individual preferences. This method combines versatile statistical tools to check for consistency of underlying preferences across different environmental conditions, in the spirit of economic tests of revealed preference. The prime virtue of the approach is its flexibility, though analysts must specify how preferences should manifest in varying environments, and the estimation procedure may be computationally intensive.

In conjunction, I run perceptual decision-making experiments to empirically assess optimal deliberation by varying problem difficulty and incentive scheme. In the data, the hypothesis of optimality is rejected for nearly half of participants when difficulty changes, and for nearly all participants when monetary stakes change. It must be recognized that such assessments are conditional on the specified model. If the test rejects the null hypothesis, then either the model (including the preference structure) is correctly specified

and the agent is behaving inconsistently, or the model itself is misspecified. Notwithstanding this limitation, both response time and accuracy data were fit quite well by the standard drift diffusion model used. Moreover, the intent of this paper is not to provide a global indictment of optimality, but to propose a method that can help determine under what circumstances people might behave closer to or farther from optimal benchmarks.

One possible empirical caveat is that the experiments were conducted online with limited control over factors like display settings and ambient lighting. Nonetheless, research increasingly supports the reliability of web-based experiments in terms of presentation and response timing (Brand & Bradley, 2012; Reimers & Stewart, 2015; Chetverikov & Upravitelev, 2016; Hilbig, 2016). Response timing variability in the specific JavaScript library I use (jsPsych) has been found comparable to the commonly used Psychophysics Toolbox in MATLAB (De Leeuw & Motz, 2016). Classic response time tasks yield similar results whether experiments are run in lab or online (Crump, McDonnell, & Gureckis, 2013; Simcox & Fiez, 2014; Slote & Strand, 2016), and I note that similar results in the present task were obtained in pilot experiments run using PsychToolbox in the lab (Caltech's Social Science Experimental Laboratory). Here, noise in stimulus presentation should affect a person's information processing ability like any other source of noise captured by the drift rate, while fixed timing offsets would be captured in non-decision time. Such variability should be unproblematic since comparisons are made on a within-subject basis. Furthermore, the relatively long response times observed in the experiments reduce the impact of small-scale variability caused by the medium. Importantly, as shown in the [Supplementary Materials](#), the estimated model predictions provide an excellent fit to individual response time and accuracy data. Thus, the use of these online experiments appears to be appropriate.

Although the experiments were restricted to two conditions for simplicity, the test can accommodate any number of these. In fact, raising the number of conditions would be beneficial by providing more opportunities to detect inconsistency. This would be especially interesting when varying task difficulty due to the non-monotonicity in the optimal decision threshold as a function of drift rate depicted in Fig. 1. Task difficulty and incentives could even be varied simultaneously to produce particularly strong tests. I leave these as important steps for future work.

To the extent that past research can be compared to the current experiments, there seems to be agreement that suboptimality is not uncommon. Zacksenhouse et al. (2010) analyzed data from an experiment with homogeneous task difficulty, and claimed that approximately 70% of

¹⁷This is not unreasonable here; in response to a post-experiment survey question, 13 out of 20 participants replied they were “very motivated” to win bonus money (the highest rating) and the other seven were “somewhat motivated” (the second highest rating).

participants might be considered suboptimal. Krajbich et al. (2014) and Oud et al. (2016) conducted experiments with mixtures of high stakes and low stakes trials faced in a fixed amount of total time. More time should be spent when stakes are high, but participants at large violated this prediction. It must be kept in mind, though, that these studies are not perfectly comparable to mine since Zacksenhouse et al. (2010) focused on reward rate and Krajbich et al. (2014) and Oud et al. (2016) used a value-based task which implicitly has heterogeneous difficulty.

There are two other differences between my study and related experiments that are worth remarking upon. First, decisions in my task took longer than in most similar perceptual studies. Mean response times in dot motion paradigms are often on the order of half a second, and even the 90th percentile tends to be under 1 s (e.g., Simen et al., 2009; Bogacz et al., 2010; Balci et al. 2011). My experiments had mean response times of around 2 s and 90th percentiles of around 5 s.¹⁸ What this entails for optimality in general is unclear. On longer timescales, control over timing does not need to be as precise to approximate optimality. For example, noise on the order of 1 s is enormous if decisions take about 1 s but negligible if decisions take 20 s. However, other cognitive elements could contaminate the decision-making process as time passes, inhibiting optimal behavior. These elements could also change the choice process itself, leading to different standards of optimality. Given that many meaningful economic decisions take place on longer timescales, it is imperative that we push the frontier of sequential sampling models along this dimension, and determine when they continue to work well and when they need to be modified. The current data indicates that the basic drift diffusion model fits behavior well even over somewhat longer time periods.

The second difference is that participants had limited experience in my task. In some studies participants engage in many thousands of dot motion trials over ten or more cumulative hours (Simen et al., 2009; Balci et al., 2011). My participants completed 200 trials and spent between 15 and 40 min on the task. Nor did they have prior experience, as only three participants in Experiment 1 and four participants in Experiment 2 stated that they had done a very similar task before. It is possible that people could learn towards optimal behavior with training. Balci et al. (2011) find that subjects do indeed approach optimality (by their definition) after experience. In particular, they observe that an accuracy

bias dominates early performance but appears to decline with practice. Studying the process by which this occurs, as with research on threshold-setting algorithms (Myung and Busemeyer, 1989; Simen, Cohen, & Holmes, 2006; Bogacz et al., 2006) is a valuable endeavor.

Note finally that questions surrounding optimality are not restricted to humans. The ability to deliberate appropriately holds substantial evolutionary consequences for many species, such as for prey deciding how to dodge a predator. Versions of the random dot motion task have already been administered to rats and mice (Douglas, Neve, Quittenbaum, Alam, & Prusky, 2006), pigeons (Nguyen et al., 2004), and rhesus macaques (Kim & Shadlen, 1999), and the testing approach I propose can be applied in this task (using learned cues to signal the environmental context) and far beyond. Analyses of economic rationality in various settings have sometimes found striking support across a variety of species (e.g., Kagel, Battalio, & Green, 1995). The degree to which nonhuman animals adhere to optimality in deliberative time allocation may shed light on human behavior as well.

Conclusions

In this paper, I propose a method of testing whether patterns of stochastic choice and time expenditure are consistent with expected utility maximization, allowing for individual subjective preferences. This test determines whether the preference for time remains consistent when environmental conditions change, which is a core feature of economic rationality. I conduct motion discrimination experiments in which task difficulty and monetary incentives are varied, demonstrating how the test can be used in conjunction with the drift diffusion model, a popular model of deliberative time allocation. Utility maximization plays a central role in many theories of deliberation including the drift diffusion model, but has been neglected empirically in this context due to the challenge of accounting for subjective preferences. The test proposed here addresses this issue, enabling evaluation of behavior against this important benchmark across a wide range of paradigms.

Acknowledgements Thanks to Colin Camerer, Jaron Colas, Hayley Dorfman, Sam Gershman, Taisuke Imai, Ian Krajbich, and Tomasz Strzalecki for helpful comments and discussions. Funding from the Social Sciences and Humanities Research Council of Canada and the Harvard Mind Brain Behavior Interfaculty Initiative is gratefully acknowledged.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

¹⁸Similar relatively long response times were also observed in pilot experiments conducted in the lab (Caltech's Social Science Experimental Laboratory) using Psychophysics Toolbox with participants experienced in a variety of tasks, indicating this is not an online artifact.

References

- Afriat, S. N. (1967). The construction of utility functions from expenditure data. *International Economic Review*, 8(1), 67–77.
- Arrow, K. J., Blackwell, D., & Girshick, M. A. (1949). Bayes and minimax solutions of sequential decision problems. *Econometrica*, 17(3/4), 213–244.
- Balcı, F., Simen, P., Niyogi, R., Saxe, A., Hughes, J. A., Holmes, P., & Cohen, J. D. (2011). Acquisition of decision-making criteria: Reward rate ultimately beats accuracy. *Attention Perception, & Psychophysics*, 73(2), 640–657.
- Bhui, R. (2018). Falling behind: Time and expectations based reference dependence. *Decision*. Advance online publication, <https://doi.org/10.1037/dec0000102>.
- Binmore, K. (2009). *Rational decisions*. Princeton: Princeton University Press.
- Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision-making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, 113(4), 700.
- Bogacz, R., Hu, P. T., Holmes, P. J., & Cohen, J. D. (2010). Do humans produce the speed–accuracy trade-off that maximizes reward rate? *Quarterly Journal of Experimental Psychology*, 63(5), 863–891.
- Brand, A., & Bradley, M. T. (2012). Assessing the effects of technical variance on the statistical outcomes of web experiments measuring response times. *Social Science Computer Review*, 30(3), 350–357.
- Britten, K. H., Shadlen, M. N., Newsome, W. T., & Movshon, A. J. (1992). The analysis of visual motion: A comparison of neuronal and psychophysical performance. *Journal of Neuroscience*, 12(12), 4745–4765.
- Brown, S., & Heathcote, A. (2003). QMLE: Fast, robust, and efficient estimation of distribution functions based on quantiles. *Behavior Research Methods Instruments, & Computers*, 35(4), 485–492.
- Bussemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision-making in an uncertain environment. *Psychological Review*, 100(3), 432.
- Chabris, C. F., Morris, C. L., Taubinsky, D., Laibson, D., & Schuldt, J. P. (2009). The allocation of time in decision-making. *Journal of the European Economic Association*, 7(2-3), 628–637.
- Chetverikov, A., & Upravitelev, P. (2016). Online versus offline: The Web as a medium for response time data collection. *Behavior Research Methods*, 48(3), 1086–1099.
- Clithero, J. A. (2016). Response times in economics: Looking through the lens of sequential sampling models. SSRN Working Paper 2795871.
- Clithero, J. A. (2018). Improving out-of-sample predictions using response times and a model of the decision process. *Journal of Economic Behavior & Organization*, 148, 344–375.
- Crawford, I., & De Rock, B. (2014). Empirical revealed preference. *Annual Review of Economics*, 6(1), 503–524.
- Crump, M. J. C., McDonnell, J. V., & Gureckis, T. M. (2013). Evaluating Amazon's Mechanical Turk as a tool for experimental behavioral research. *PLOS One*, 8(3), e57410.
- De Leeuw, J. R. (2015). jsPsych: A JavaScript library for creating behavioral experiments in a Web browser. *Behavior Research Methods*, 47(1), 1–12.
- De Leeuw, J. R., & Motz, B. A. (2016). Psychophysics in a Web browser? Comparing response times collected with JavaScript and Psychophysics Toolbox in a visual search task. *Behavior Research Methods*, 48(1), 1–12.
- Domenech, P., & Dreher, J.-C. (2010). Decision threshold modulation in the human brain. *Journal of Neuroscience*, 30(43), 14305–14317.
- Donders, F. C. (1869/1969). On the speed of mental processes. *Acta Psychologica*, 30, 412–431.
- Douglas, R. M., Neve, A., Quittenbaum, J. P., Alam, N. M., & Prusky, G. T. (2006). Perception of visual motion coherence by rats and mice. *Vision Research*, 46(18), 2842–2847.
- Drugowitsch, J., DeAngelis, G. C., Angelaki, D. E., & Pouget, A. (2015). Tuning the speed–accuracy trade-off to maximize reward rate in multisensory decision-making. *eLife*, 4, e06678.
- Drugowitsch, J., DeAngelis, G. C., Klier, E. M., Angelaki, D. E., & Pouget, A. (2014). Optimal multisensory decision-making in a reaction-time task. *eLife*, 3, e03005.
- Drugowitsch, J., Moreno-Bote, R., Churchland, A. K., Shadlen, M. N., & Pouget, A. (2012). The cost of accumulating evidence in perceptual decision-making. *Journal of Neuroscience*, 32(11), 3612–3628.
- Fudenberg, D., Strack, P., & Strzalecki, T. (2018). Speed, accuracy, and the optimal timing of choices. *American Economic Review*, 108(12), 3651–3684.
- Gintis, H. (2009). The bounds of reason: Game theory and the unification of the behavioral sciences.
- Gluth, S., Rieskamp, J., & Büchel, C. (2012). Deciding when to decide: time-variant sequential sampling models explain the emergence of value-based decisions in the human brain. *Journal of Neuroscience*, 32(31), 10686–10698.
- Gluth, S., Rieskamp, J., & Büchel, C. (2013). Classic EEG motor potentials track the emergence of value-based decisions. *NeuroImage*, 79, 394–403.
- Gold, J. I., & Shadlen, M. N. (2002). Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. *Neuron*, 36(2), 299–308.
- Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision-making. *Annual Review of Neuroscience*, 30, 535–574.
- Green, N., Biele, G. P., & Heekeren, H. R. (2012). Changes in neural connectivity underlie decision threshold modulation for reward maximization. *Journal of Neuroscience*, 32(43), 14942–14950.
- Hanes, D. P., & Schall, J. D. (1996). Neural control of voluntary movement initiation. *Science*, 274(5286), 427–430.
- Heathcote, A., Brown, S., & Cousineau, D. (2004). QMPE: Estimating Lognormal, Wald, And Weibull RT distributions with a parameter-dependent lower bound. *Behavior Research Methods Instruments, & Computers*, 36(2), 277–290.
- Heathcote, A., Brown, S., & Mewhort, D. J. K. (2002). Quantile maximum likelihood estimation of response time distributions. *Psychonomic Bulletin & Review*, 9(2), 394–401.
- Hick, W. E. (1952). On the rate of gain of information. *Quarterly Journal of Experimental Psychology*, 4(1), 11–26.
- Hilbig, B. E. (2016). Reaction time effects in lab-versus Web-based research: Experimental evidence. *Behavior Research Methods*, 48(4), 1718–1724.
- Houthakker, H. S. (1950). Revealed preference and the utility function. *Economica*, 17(66), 159–174.
- Jensen, A. R. (2006). *Clocking the mind: Mental chronometry and individual differences*. New York: Elsevier.
- Kagel, J. H., Battalio, R. C., & Green, L. (1995). *Economic choice theory: An experimental analysis of animal behavior*. Cambridge: Cambridge University Press.
- Karşılar, H., Simen, P., Papadakis, S., & Balcı, F. (2014). Speed accuracy trade-off under response deadlines. *Frontiers in Neuroscience*, 8, 248.
- Kim, J.-N., & Shadlen, M. N. (1999). Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nature Neuroscience*, 2(2), 176–185.
- Krajich, I., & Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences*, 108(33), 13852–13857.

- Krajibich, I., Armel, C., & Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nature Neuroscience*, 13(10), 1292–1298.
- Krajibich, I., Hare, T., Bartling, B., Morishima, Y., & Fehr, E. (2015). A common mechanism underlying food choice and social decisions. *PLOS Computational Biology*, 11(10), e1004371.
- Krajibich, I., Lu, D., Camerer, C., & Rangel, A. (2012). The attentional drift-diffusion model extends to simple purchasing decisions. *Frontiers in Psychology*, 3, 193.
- Krajibich, I., Oud, B., & Fehr, E. (2014). Benefits of neuroeconomic modeling: New policy interventions and predictors of preference. *American Economic Review*, 104(5), 501–506.
- Laming, D. R. J. (1968). *Information theory of choice-reaction times*. New York: Academic Press.
- Milosavljevic, M., Malmaud, J., Huth, A., Koch, C., & Rangel, A. (2010). The drift diffusion model can account for the accuracy and reaction time of value-based choices under high and low time pressure. *Judgment and Decision Making*, 5(6), 437–449.
- Moran, R. (2015). Optimal decision making in heterogeneous and biased environments. *Psychonomic Bulletin & Review*, 22(1), 38–53.
- Mullen, K. M., Ardia, D., Gil, D. L., Windover, D., & Cline, J. (2011). DEOptim: An R package for global optimization by differential evolution. *Journal of Statistical Software*, 40(6), 1–26.
- Myung, I. J., & Busemeyer, J. R. (1989). Criterion learning in a deferred decision-making task. *American Journal of Psychology*, 102(1), 1–16.
- Newsome, W. T., Britten, K. H., & Movshon, A. J. (1989). Neuronal correlates of a perceptual decision. *Nature*, 341(6237), 52–54.
- Nguyen, A. P., Spetch, M. L., Crowder, N. A., Winship, I. R., Hurd, P. L., & Wylie, D. R. W. (2004). A dissociation of motion and spatial-pattern vision in the avian telencephalon: Implications for the evolution of visual streams. *Journal of Neuroscience*, 24(21), 4962–4970.
- Oud, B., Krajibich, I., Miller, K., Cheong, J. H., Botvinick, M., & Fehr, E. (2016). Irrational time allocation in decision-making. *Proceedings of The Royal Society B: Biological Sciences*, 283(1822), 20151439.
- Rajananda, S., Lau, H., & Odegaard, B. (2018). A random-dot kinematogram for web-based vision research. *Journal of Open Research Software*, 6(1), 6.
- Rand, D. G., Greene, J. D., & Nowak, M. A. (2012). Spontaneous giving and calculated greed. *Nature*, 489(7416), 427–430.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2), 59.
- Ratcliff, R. (2002). A diffusion model account of response time and accuracy in a brightness discrimination task: Fitting real data and failing to fit fake but plausible data. *Psychonomic Bulletin & Review*, 9(2), 278–291.
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, 20(4), 873–922.
- Ratcliff, R., & Rouder, J. N. (2000). A diffusion model account of masking in two-choice letter identification. *Journal of Experimental Psychology: Human Perception and Performance*, 26(1), 127.
- Ratcliff, R., Cherian, A., & Segraves, M. (2003). A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of two-choice decisions. *Journal of Neurophysiology*, 90(3), 1392–1407.
- Reimers, S., & Stewart, N. (2015). Presentation and response timing accuracy in Adobe Flash and HTML5/javascript Web experiments. *Behavior Research Methods*, 47(2), 309–327.
- Robertson, T. (1978). Testing for and against an order restriction on multinomial parameters. *Journal of the American Statistical Association*, 73(361), 197–202.
- Robertson, T., Wright, F. T., & Dykstra, R. (1988). *Order restricted statistical inference*. New York: Wiley.
- Rubinstein, A. (2007). Instinctive and cognitive reasoning: A study of response times. *Economic Journal*, 117(523), 1243–1259.
- Samuelson, P. A. (1938). A note on the pure theory of consumer's behaviour. *Economica*, 5(17), 61–71.
- Shadlen, M. N., & Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *Journal of Neurophysiology*, 86(4), 1916–1936.
- Shepard, R. N., & Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science*, 171(3972), 701.
- Simcox, T., & Fiez, J. A. (2014). Collecting response times using Amazon Mechanical Turk and Adobe Flash. *Behavior Research Methods*, 46(1), 95–111.
- Simen, P., Cohen, J. D., & Holmes, P. (2006). Rapid decision threshold modulation by reward rate in a neural network. *Neural Networks*, 19(8), 1013–1026.
- Simen, P., Contreras, D., Buck, C., Hu, P., Holmes, P., & Cohen, J. D. (2009). Reward rate optimization in two-alternative decision making: Empirical tests of theoretical predictions. *Journal of Experimental Psychology: Human Perception and Performance*, 35(6), 1865.
- Slote, J., & Strand, J. F. (2016). Conducting spoken word recognition research online: Validation and a new timing method. *Behavior Research Methods*, 48(2), 553–566.
- Smith, P. L., & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends in Neurosciences*, 27(3), 161–168.
- Smith, P. L., Ratcliff, R., & Wolfgang, B. J. (2004). Attention orienting and the time course of perceptual decisions: Response time distributions with masked and unmasked displays. *Vision Research*, 44(12), 1297–1320.
- Spiliopoulos, L., & Ortmann, A. (2018). The BCD of response time analysis in experimental economics. *Experimental Economics*, 21(2), 383–433.
- Starns, J. J., & Ratcliff, R. (2010). The effects of aging on the speed–accuracy compromise: Boundary optimality in the diffusion model. *Psychology and Aging*, 25(2), 377.
- Starns, J. J., & Ratcliff, R. (2012). Age-related differences in diffusion model boundary optimality with both trial-limited and time-limited tasks. *Psychonomic Bulletin & Review*, 19(1), 139–145.
- Starns, J. J., & Ratcliff, R. (2014). Validating the unequal-variance assumption in recognition memory using response time distributions instead of ROC functions: A diffusion model analysis. *Journal of Memory and Language*, 70, 36–52.
- Sternberg, S. (1966). High-speed scanning in human memory. *Science*, 153(3736), 652–654.
- Stone, M. (1960). Models for choice-reaction time. *Psychometrika*, 25(3), 251–260.
- Tuerlinckx, F., Maris, E., Ratcliff, R., & De Boeck, P. (2001). A comparison of four methods for simulating the diffusion process. *Behavior Research Methods Instruments, & Computers*, 33(4), 443–456.
- Wald, A. (1947). *Sequential analysis*. New York: Wiley.
- Wald, A., & Wolfowitz, J. (1948). Optimum character of the sequential probability ratio test. *Annals of Mathematical Statistics*, 19(3), 326–339.
- Webb, R. (2018). The dynamics of stochastic choice. *Management Science*. Advance online publication, <https://doi.org/10.1287/mnsc.2017.2931>.
- Wilcox, N. T. (1993). Lottery choice: Incentives, complexity and decision time. *Economic Journal*, 103(421), 1397–1417.
- Woodford, M. (2014). Stochastic choice: An optimizing neuroeconomic model. *American Economic Review*, 104(5), 495–500.
- Zacksenhouse, M., Bogacz, R., & Holmes, P. (2010). Robust versus optimal strategies for two-alternative forced choice tasks. *Journal of Mathematical Psychology*, 54(2), 230–246.