Title
Sequential effects: A Bayesian analysis of prior bias on reaction time and behavioral choice

Permalink
https://escholarship.org/uc/item/9jp3j5c7

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 36(36)

ISSN
1069-7977

Authors
Zhang, Shunan
Huang, He Crane
Yu, Angela J.

Publication Date
2014

Peer reviewed

Sequential effects: A Bayesian analysis of prior bias on reaction time and behavioral choice

Shunan Zhang   **He Crane Huang**   **Angela J. Yu**

(s6zhang, heh001, ajyu@ucsd.edu)
Department of Cognitive Science, University of California, San Diego
9500 Gilman Drive, La Jolla, CA 92093-0515

Abstract

Human subjects exhibit "sequential effects" in many psychological experiments, in which they respond more rapidly and accurately to a stimulus when it reinforces a local pattern in stimulus history, compared to when it violates such a pattern. This is often the case even if the local pattern arises by chance, such that stimulus history has no real predictive power, and therefore any behavioral adjustment based on these erroneous predictions essentially amounts to superstition. Earlier, we proposed a normative Bayesian learning model, the Dynamic Belief Model (DBM), to demonstrate that such behavior reflects the engagement of mechanisms that identify and adapt to changing patterns in the environment (Yu & Cohen, 2009). In that earlier work, we assumed a monotonic relationship between prior bias and response time (bias *toward* a stimulus was assumed to result in faster reaction time when that was the actual stimulus; conversely, when the other stimulus was present, it was assumed to result in a slower response). Here, we present a more detailed and quantitative analysis of the relationship between prior bias and behavioral outcome, in terms of response time and choice accuracy. We also present novel behavioral data, along with a framework for jointly identifying subject-specific parameters of the trial-by-trial learning (Dynamic Belief Model, DBM) and within-trial sensory processing and decision-making (Drift-Diffusion Model, DDM) based on the behavioral data. Our results provide strong evidence for DBM, and reveal potential individual differences, in their differential beliefs about the timescale at which local patterns persist in sequential data.

Keywords: Perceptual Decision Making; Sequential Effects; Bayesian Model; Drift-Diffusion Model

Introduction

In a variety of behavioral experiments, human subjects display "sequential effects", a modulation of response time and/or accuracy by recent trial history (e.g. Bertelson, 1961; Laming, 1968; Kornblum, 1973; Soetens, Boer, & Hueting, 1985; Cho et al., 2002; Jones, Curran, Mozer, & Wilder, 2013). For example, in two-alternative forced choice experiments, in which subjects discriminate between two types of stimuli (A or B), subjects respond more accurately and rapidly if a trial is consistent with the recent pattern (e.g. $AAAAA$ followed by A, $ABABA$ followed by B), than if it is inconsistent (e.g. $AAAAA$ followed by B, $ABABA$ followed by A). This sequential effect depends on the length of the run (Cho et al., 2002). For instance, an alternation following four repetitions affects responding more than one following only two repetitions. Figure 1 illustrates a robust finding of the dependence of RT and error rate on recent trial history, both being largest when a relative long run of repetitions or alternations are broken by the current observation (middle two trial types), and smallest when such runs are extended (left and right end).

Previously, we proposed a Bayesian learning model, the Dynamic Belief Model (DBM), to account for sequential effects, via a human learning mechanism that assumes the potential for discrete, un-signaled changes in the environment. Consequently, DBM repeatedly modifies internal estimates of the relative probability of one stimulus type versus another occurring, based on recent stimulus history (Yu & Cohen, 2009). By assuming reaction time and error rate to be monotonically and inversely correlated with the estimated prior probability of observing the actual stimulus prior to stimulus onset, DBM can qualitatively reproduce the empirically observed sequential effects shown by Cho et al. (2002).

In this work, we give a more precise and quantitative treatment of the influence of prior expectations on sensory processing and decision-making *within a trial*, by assuming an evidence-integration-to-bound process (Gold, 2002), which is formally similar to the Drift-Diffusion Model (DDM) (e.g. Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006) and appears to explain activities of parietal cortical neurons during primate perceptual decision-making (Roitman & Shadlen, 2014). We present a Bayesian method for simultaneously identifying subject-specific parameters of DBM and DDM based on an individual's choice accuracy and reaction times, and apply it to behavioral data collected in a simple 2-alternative choice perceptual discrimination task. Using this quantitative method, we will compute the relative support, measured in Bayes factors, the data lend to DBM versus a competing model, the Fixed Belief Model (FBM) (Yu & Cohen, 2009), which assumes that human subjects do not believe the task statistics to be changeable over time. We will also characterize the population distributions of subject-specific Bayesian model parameters, which correspond to semantically readily interpretable variables, such as subjects' beliefs about the rate of change in the environmental statistics, the overall relative frequency of repetition and alternation trials, and the subjective diffi-

culty (or sensory/perceptual noise) associated with processing the sensory stimulus.

The paper is organized as follows. First, we will describe the experiment and the data collected. Second, we will review DBM and FBM, showing their qualitative differences in trial-wise behavior. Next, we will introduce the quantitative model of prior bias on choice RT and accuracy. Subsequently, we will analyze the empirical data using the models. And finally, we will conclude with a discussion of the results, their implications, and directions for future inquiries.

Experiment

42 college students in UCSD participated in this experiment. All with normal or corrected vision. In this 2AFC discrimination task, subjects were instructed to decide if the coherent motion of a patch of moving dots on the computer screen was toward left or right by pressing the corresponding arrow keys. Subjects were seated in a chair approximately 60 cm from the computer screen. When a trial started, a patch (diameter: 5 deg visual angel) with coherent moving dots would appear on the screen, moving at a speed of 5 deg/second. The density of the dots was 16.7 dots per sq deg/second, with 3 pixels per dot.

All subjects needed to complete two practice sessions and reach an accuracy of at least 80% to proceed to the main experiment. In the main experiment, there were 7 coherence levels, ranging from 0 to 100%. There were 14 blocks in total, with 120 trials/block, and 2 blocks/coherence. As our first modeling attempt, we only considered data collected from the blocks with 100% coherence. There was no time limit on each trial. Stimuli repetition or alternation were d650 ms. Subjects received feedback at the end of each trial with a beep to indicate if their response was correct. A crisp beep indicated a correct response, while a low-frequency beep indicated an incorrect response. There was a 4-second black screen penalty for premature response (respond within 100 ms after stimuli appears).

Figure 1 C;D demonstrate that very similar sequential effects were observed in this experiment as in (Cho et al., 2002).

Learning Models

We give a brief summary of the two Bayes-optimal, ideal observer models in (Yu & Cohen, 2009), which have different assumptions about the temporal persistence of statistical contingencies in the world.

Fixed Belief Model (FBM): Learning about a Fixed World

Suppose the subject has an internal model that on each trial t, there is a fixed probability γ of seeing a repetition ($x_t = 1$), and therefore a probability $1 - \gamma$ of seeing an
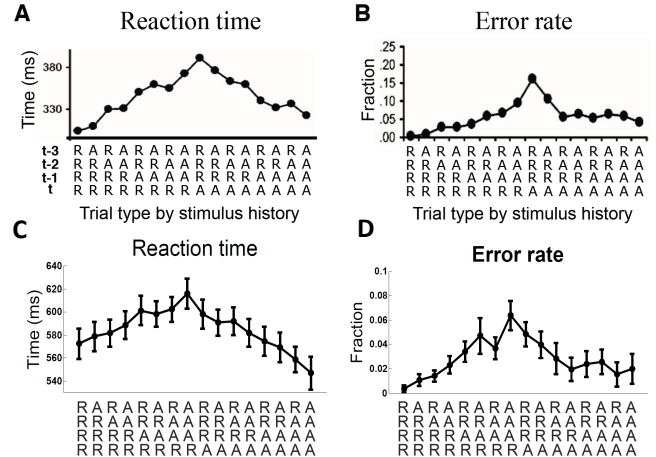

Figure 1: Sequential effects in 2AFC tasks manifested in previous and current experiments. (A) Median reaction time (RT) and (B) error rate (ER) averaged across six subjects, adapted from Figure 1 of (Cho et al., 2002). Along the abscissa are all of the possible five-stimulus-long sequences, where R stands for repetition, and A stands for alternation. Each sequence, read from top to bottom, proceeds from the earliest stimulus progressively toward the present stimulus. (C) Median RT and (D) error rate from the current experiment show similar patterns; error bars: s.e.m. (each error bar shows the standard error of subjects' median RT's for the corresponding sequence)

alternation ($x_t = 0$). $p_0(\gamma)$ is the generic prior to capture the subject's belief about γ before any observations. The prior is modeled as a Beta distribution, Beta (a, b). Over time, the subject can use the number of observed repetitions versus alternations to gain an increasingly precise estimate of the underlying γ. After t observations, the posterior belief of γ is

$$p(\gamma|\mathbf{x}_t) \sim p(\mathbf{x}_t|\gamma)p_0(\gamma) \tag{1}$$

which is simply Beta $(a + r_t,\, b + t - r_t)$, where r_t is the number of repetitions observed up to trial t; $\mathbf{x}_t$ is shorthand for the vector of observed sequence $(x_1, \ldots, x_t)$. The probability of observing a repetition on trial $t + 1$ is the mean of the posterior distribution over γ:

$$p(x_{t+1} = 1|\mathbf{x}_t) = \frac{r_t + a}{t + a + b} \tag{2}$$

Dynamic Belief Model (DBM): Learning about a Changing World

Suppose the subjects believe that the relative frequency of repetition (versus alternation) can undergo discrete

Figure 2: Generative models of the fixed and dynamic belief models. (A) Fixed Belief Model (FBM): a hidden bias parameter γ specifies the frequency of repetitions (and alternations) in the experiment. (B) Dynamic Belief Model (DBM): the hidden bias parameter can change from trial t to trial $t + 1$.

changes at unsignaled times during the experimental session (see Figure 2B for graphical representation of the model), the subject's implicit task, then, is to track the *evolving* frequency of repetition versus alternation over the course of the experiment. The crucial assumption is that γ_t has a Markovian dependence on γ_{t-1}, such that there is a large probability α of them being the same, and a small probability $1 - \alpha$ of γ_t being redrawn from the generic prior distribution, p_0. As with the FBM, the observer would then need to combine the sequentially developed prior belief about stimulus identity, with the incoming stream of sensory inputs, $x_1, x_2, \ldots, x_t, \ldots$, to infer the identity of the stimulus in each trial in an iterative fashion:

$$p(\gamma_t = \gamma | \mathbf{x}_{t-1}) = \alpha p(\gamma_{t-1} = \gamma | \mathbf{x}_{t-1}) + (1 - \alpha)p_0(\gamma) \quad (3)$$

and the posterior distribution is:

$$p(\gamma_t = \gamma | \mathbf{x}_t) \propto p(x_t | \gamma_t = \gamma)p(\gamma_t = \gamma | \mathbf{x}_{t-1}) \quad (4)$$

The probability of seeing a new repetition is thus:

$$
\begin{aligned}
p(x_t = 1 | \mathbf{x}_{t-1}) &= \int p(x_t = 1 | \gamma_t)p(\gamma_t | \mathbf{x}_{t-1})d\gamma_t \\
&= (1 - \alpha)\langle \gamma \rangle_{p_0(\gamma)} + \alpha \langle \gamma_t | \mathbf{x}_{t-1} \rangle
\end{aligned}
\quad (5)
$$

One important consequence of the diminishing uncertainty in the FBM versus the persisting uncertainty in the DBM is that the influence of individual observations persist indefinitely in FBM, but decreases over time for DBM, with the parameter α determining the effective time window over which individual events exert predictive influence on future events. Figure 3 graphically illustrates analyses of the consequences of the different assumptions made by the two models. These simulation results support our hypothesis that the trial-to-trial adjustments seen in subjects' behavior in 2AFC tasks reflect a
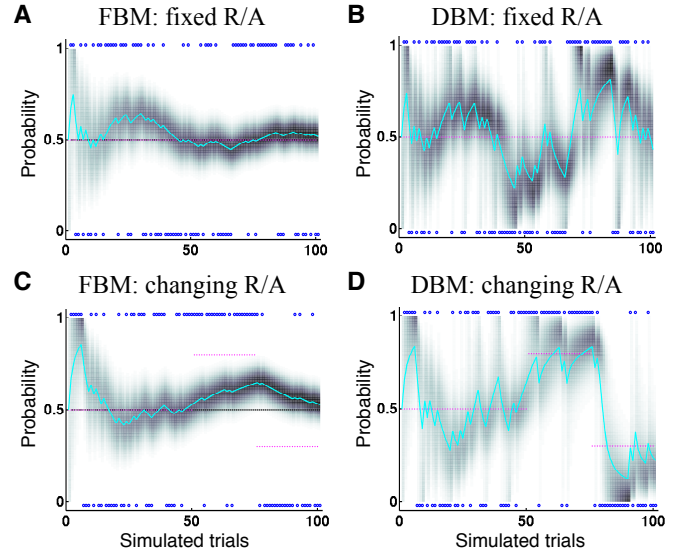

Figure 3: Sequential effects transient in FBM, persistent in DBM. (A) illustrates the trial-wise predictive probability of repetition by FBM (cyan), to a sequence of 100 data points, random with $\gamma = .5$ for a repetition. An ideal observer implementing FBM is superior: it becomes immune to the noisy fluctuations in the sequence of observations (the darkening and narrowing band) and converges to the true value. (B) illustrates the trial-wise predictive probability by DBM under the same process. As DBM is (erroneously) applied to learn about a stationary (and random) process, it is strongly and persistently influenced by spurious local patterns. (C) When the underlying environment is truly volatile, FBM cannot easily adapt to new values of γ, whereas as (D) DBM negotiates these changes adroitly.

(perhaps implicit and unconscious) assumption that statistical regularities, such as runs of repetitions or alternations, exist and persist on a characteristic timescale. Such a strategy is useful for a truly volatile environment but inappropriate for the experimental environment, in which stimulus statistics are held fixed.

Bayesian Model of Prior Belief on Reaction Time and Choice

Previously (Yu & Cohen, 2009), we suggested within-trial perceptual inference and decision-making to be analogous to the sequential hypothesis ratio test (SPRT), and made a loose argument for the prior bias (inferred by DBM) to have an approximately linear effect on mean RT. Separately, we have found that the approximately linear relationship to RT hold for a wide range of α values (data not shown). Here, we take a similar approach but explicitly model the relationship between prior bias and RT. The
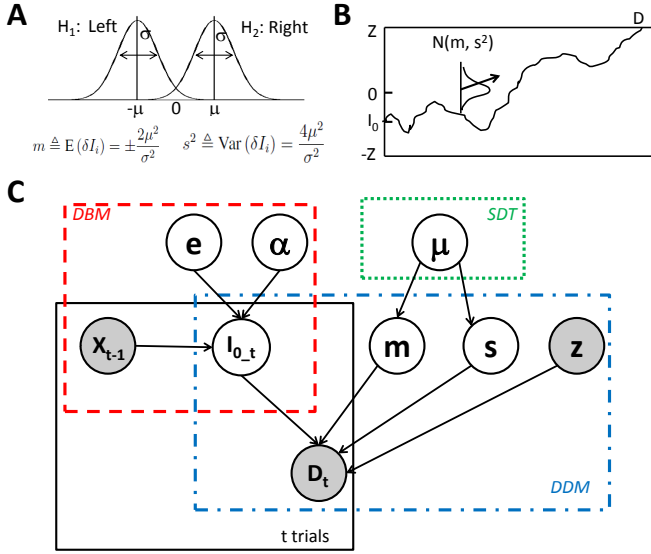
Figure 4: Graphical representation of the joint inference of the cross-trial learning (DBM) and within-trial decision making (DDM). Discriminability of the hypothetical distributions in SDT (captured by μ, since we can fix σ at 1 without loss of generality) determines the rate (m) and variability (s^2) of sensory evidence accumulation in DDM within a trial. The bias of DDM (I_0) is determined by the prior probability of repetition as inferred by DBM from cross-trial learning, conditioned on the parameter associated with a stationary environment (α) and the generic prior belief of the probability of a repetition trial (e). Shaded nodes represent variables that are observable to the experimenter, or can be calculated in a model-free fashion, such as the stimulus sequence (x_t), the observed choices and RT's (D_t), and the log odds of overall accuracy of decision making (Z). The black solid frame indicates t repeated plates of the cross-trial variables. Colored broken frames illustrate different components in the composite model.

general idea is to specify a within-trial mapping from the belief state (prior probability impending stimulus type), and the stimulus strength/sensory uncertainty, to choice and RT of a given trial.

Within-Trial Processing

We introduce our method in the context of our experiment, whereby the decision maker needs to decide whether the coherent motion of the dots is toward left or right. For simplicity, and similar to the signal-detection-theory (SDT) formulation (Green & Swets, 1966), we assume each of the two possible stimuli generates normally distributed noisy neural responses at some intermediate stage of the visual pathway, based on which the subsequent brain region(s) must decide which stimulus was

present (and thus which response is required) and when to respond. We assume that the perceived strength and uncertainty of the motion does not depend on its direction, thus the two distributions under hypothesis H_1 (motion toward the left) and H_2 (motion toward the right) have means $-\mu$ and μ, and equal variance, σ^2. The distribution under the true hypothesis is termed the *target* distribution. Figure 4A illustrates SDT.

SPRT solves the problem of deciding between H_1 and H_2, based on an ongoing stream of independent series of sensory signals from the stimulus (the target distribution), $y_1, \ldots, y_t, \ldots$, perceived at discrete steps. The total length (sample size) of sensory signals is also under the observer's control. SPRT says that the observer should keep tracking the relative likelihoods of the two hypotheses being true, and choose the more likely one as soon as the likelihood ratio crosses some decision threshold Z_1 (in which case, stop and decide H_1) or Z_2 (stop and decide H_2). Suppose the prior probability of H_1 being true is p; the probability of the sensory signals up to time t, $\mathbf{y}_t := \{y_1, \ldots, y_t\}$, conditioned on hypothesis H_1 being true, is $f_1(\mathbf{y}_t)$, and the probability of the same sequence of sensory signals being generated by hypothesis H_2 is $f_2(\mathbf{y}_t)$, then SPRT says to stop as soon as $S_n := \frac{p f_1(\mathbf{y}_t)}{(1-p) f_2(\mathbf{y}_t)} \geq Z_1$ or if $S_n \leq Z_2$, and continue otherwise (i.e. if $Z_2 < S_n < Z_1$). Suppose ϵ is the type I error to be controlled for deciding on either hypothesis, then $Z_1 = \frac{1-\epsilon}{\epsilon}$, and $Z_2 = \frac{\epsilon}{1-\epsilon}$.

It has long been noted that SPRT is formally equivalent to a bounded random-walk model (Laming, 1968; Bogacz et al., 2006). When the observations have statistically independent noise, we have $f_j(y_1, \ldots, y_t) = \prod_{i=1}^{t} f_j(y_i)$ for $j = 1, 2$, and thus

$$I_n := \log S_n = \log \frac{p}{1-p} + \sum_{i=1}^{t} \log \frac{f_1(y_i)}{f_2(y_i)}$$

Notice that f_1 and f_2 are density functions of Gaussian distributions $N(-\mu, \sigma^2)$ and $N(\mu, \sigma^2)$. The increment of information gained from a sensory signal y_i, and its mean and variance are

$$\delta I_i = \log \frac{f_1(y_i)}{f_2(y_i)} = -\frac{2\mu}{\sigma^2} y_i$$

$$m := \mathrm{E}\,(\delta I_i) = \pm \frac{2\mu^2}{\sigma^2} \qquad (6)$$

$$s^2 := \mathrm{Var}\,(\delta I_i) = \frac{4\mu^2}{\sigma^2}$$

Since $\log S_n$ is strictly monotonically related to S_n, the decision policy is equivalent to stopping as soon as $I_n \geq Z$ (and choose H_1) or $I_n \leq -Z$ (and choose H_2), for $Z := \log(1-\epsilon) - \log \epsilon$. In other words, the sensory signal accumulation in SPRT is equivalent to a bounded random

Figure 5: Distributions of the MAP estimates of α, e and μ

Figure 6: Left: distribution of Bayes factors of DBM against FBM, with each brick showing one subject; Right: model predicted RT's compared to data.

walk with noisy increments that have a mean drift rate of m per time step (positive if H_1 is true, negative if H_2 is true).

We can then rewrite the total sensory evidence accrued with n steps as

$$I_n = I_0 + \sum_{i=1}^{n} \delta I_i \qquad (7)$$

where $I_0 := \log p - \log(1-p)$ is the starting bias toward H_1. Assuming sensory signals are obtained at small time intervals, we consider the continuous-time limiting process, $I(t)$, which satisfies the stochastic differential equation

$$dI = m dt + s dW, \quad I(0) = I_0 \qquad (8)$$

where m and s^2 are defined by Equation 6. Equation 8 is a drift diffusion process. The thresholds for I_n, with respect to the thresholds in the original SPRT form, are just $\pm Z$. Figure 4B illustrates DDM.

Cross-Trial Processing

Both DBM and FBM infer the identity of the stimulus (repetition vs alternation) for each trial, based on the previous observed sequence (Equation 2 and Equation 5). At trial t, the prior probability of seeing a repetition $p(x_t = 1 | \mathbf{x}_{t-1})$ can be readily translated to the bias in SPRT: $I_0 = \pm (\log p(x_t = 1 | \mathbf{x}_{t-1}) - \log(1 - p(x_t = 1 | \mathbf{x}_{t-1})))$ which takes positive (negative) sign if H_1 (H_2) was true on the previous trial.

Graphical Model Implementation

Figure 4C shows a graphical representation of our join inference of DBM, SDT, and DDM. Unshaded nodes represent model parameters to be inferred from data. Because only the ratio of μ and σ (relative discriminability of the two hypotheses) matters in determining the drift of DDM, we can fix σ at 1. We model the generic prior of the probability of repetition in DBM using $\mathrm{Beta}(a, b)$, and denote its mean by $e := a/(a+b)$. To reduce computational complexity, we fix $a + b$ at 2, assuming subjects have (equally) low certainty of the environment before observing any stimulus. We also use the simplest treatment for non-decision time by subtracting the smallest RT for each subject. In fitting the joint model, we first

generated the sequential predictions by DBM, given the true sequence of stimuli observed by each subject, using a grid (.02 increment) of α and e values ranging between 0 and 1. We then used MCMC sampling for the graphical model inference using a uniform prior between 0 and 10 for μ, and "discretized uniform" priors for α and e. We fit the model to each individual subject, and conducted a formal model comparison between DBM and FBM by examining the Bayes factors (e.g. Kass & Raftery, 1995).

Results

The top row of Figure 5 shows distributions (over all subjects) of the maximum a posterior (MAP) estimates of the DBM parameter, α, the prior belief of the mean probability of repetition, e, and the psychological discriminability of the target direction, μ. The distribution of α has large variation indicating individual differences. The mean of the distribution of e is smaller than .5, implying a bias toward alternation in general. However, we can clearly see that some subjects hold a greater bias toward either repetition or alternation, as indicated by the more extreme estimates of e. There is only one subject who had poor choice accuracy (.74), captured by a low μ value in SDT representation. We did not find any significant correlations between any pair of parameters.

We examine DBM and FBM in their abilities of capturing empirical data, using the Bayes factor as a model comparison measure. We calculate the Bayes factor of DBM against FBM, both combined with DDM, for each individual subject. The left panel of Figure 6 shows the distribution of Bayes factors over all subjects. Evidence for DBM against FBM is positive for 35 (out of 42) subjects, with a Bayes factor greater than 3 (according to the interpretation scale proposed by Kass and Raftery (1995)). We also compared the RT conditioned on 5 previous trials predicted by DBM and FBM at their best parameterization. The right panel of Figure 6 shows that FBM predicts a much smaller sequential effect as compared to real-world behavior, whereas DBM can capture

a stronger sequential effect as seen in the data.

Discussion

Previous computational approach to sequential effects focused on Bayes-optimal learning mechanisms, while giving a simplified treatment to the decision process, such as assuming approximate linear effect of prior on choice RT (Yu & Cohen, 2009; Jones et al., 2013). In this study, we assume an explicit model for within-trial processing, and develop a method for the joint inference of cross-trial learning and within-trial decision making, by augmenting the computational learning model with a principled, sequential hypothesis testing paradigm that is proven to be optimal in both the frequentist and the Bayesian sense (Wald & Wolfowitz, 1948).

Our joint inference and model comparison results support DBM as a better account of human sequential learning than FBM. On the other hand, our results also provide strong evidence for individual differences in their belief of the rate of changing of the environment. The distribution of the inferred α values across all subjects has a large variation that implies potential individual differences.

Our model builds prior knowledge in the starting point of SPRT (and its continuous-time limit, DDM). One of the main theoretical points of proponents of alternative race/accumulator type of models is that the starting point confounds prior knowledge with decision utilities (values). We do not manipulate utilities in the current study, yet future work involves an analysis of how the model would handle decision utilities.

Another future direction is to consider a joint inference of computational and neural models, by extending a newly developed, statistical approach of combining neural and behavioral measures to study cognition (Turner, in press). By jointly fit the computational and neural models, it would become feasible to make simultaneous inference about the correlation between parameters at these different levels.

Acknowledgments

References

Bertelson, P. (1961). Sequential redundancy and speed in a serial two-choice responding task. *Quarterly Journal of Experimental Psychology*, *13*, 90–102.

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*, 700–765.

Cho, R., Nystrom, L., Brown, E., Jones, A., Braver, T., Holmes, P., & Cohen, J. D. (2002). Mechanisms underlying dependencies of performance on stimulus history in a two-alternative forced-choice task. *Cognitive, Affective and Behavioral Neuroscience*, *2*, 283–299.

Gold, J. I. (2002). Good vibrations. *Neuron*, *33*, 842–4.

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: Wiley.

Jones, M., Curran, T., Mozer, M. C., & Wilder, M. H. (2013). Sequential effects in response time reveal learning mechanisms and event representations. *Psychological Review*, *120*, 628–666.

Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, *90*, 377–395.

Kornblum, S. (1973). *Sequential effects in choice reaction time. a tutorial review.* (S. Kornblum, Ed.). New York: Academic Press.

Laming, D. R. J. (1968). *Information theory of choice reaction time*. London: Academic Press.

Roitman, J. D., & Shadlen, M. N. (2014). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *The Journal of Neuroscience*, *22*, 9475–9489.

Soetens, E., Boer, L. C., & Hueting, J. E. (1985). Expectancy or automatic facilittion? separating sequential effects in two-choice reaction time. *Journal of Experimental Psychology: Human Perception and Performance*, *11*, 598–616.

Turner, B. M. (in press). Constraining cognitive abstractions through bayesian modeling. In B. U. Forstmann & E.-J. Wagenmakers (Eds.), *An introduction to model-based cognitive neuroscience*. New York: Springer.

Wald, A., & Wolfowitz, J. (1948). Optimal character of the sequential probability ratio test. *Annals of Mathematical Statistics*, *19*, 326–339.

Yu, A. J., & Cohen, J. D. (2009). Sequential effects: Superstition or rational behavior? In *Advances in neural information processing systems* (Vol. 21, pp. 1873–1880). Cambridge, MA: MIT Press.