

Sequential Effects in Response Time Reveal Learning Mechanisms and Event Representations

Matt Jones, Tim Curran, Michael C. Mozer, and Matthew H. Wilder
University of Colorado Boulder

Binary choice tasks, such as 2-alternative forced choice, show a complex yet consistent pattern of sequential effects, whereby responses and response times depend on the detailed pattern of prior stimuli going back at least 5 trials. We show this pattern is well explained by simultaneous incremental learning of 2 simple statistics of the trial sequence: the base rate and the repetition rate. Both statistics are learned by the same basic associative mechanism, but they contribute different patterns of sequential effects because they entail different representations of the trial sequence. Subtler aspects of the data that are not explained by these 2 learning processes alone are explained by their interaction, via learning from joint error correction. Specifically, the cue-competition mechanism that has explained classic findings in animal learning (e.g., blocking) appears to operate on learning of sequence statistics. We also find that learning of the base rate and repetition rate are dissociated into response and stimulus processing, respectively, as indicated by event-related potentials, manipulations of stimulus discriminability, and reanalysis of past experiments that eliminated stimuli or prior responses. Thus, sequential effects in these tasks appear to be driven by learning the response base rate and the stimulus repetition rate. Connections are discussed between these findings and previous research attempting to separate stimulus- and response-based sequential effects, and research using sequential effects to determine mental representations. We conclude that sequential effects offer a powerful means for uncovering representations and learning mechanisms.

Keywords: sequential effects, two-alternative forced choice, incremental learning, event-related potential, representation

Supplemental materials: <http://dx.doi.org/10.1037/a0033180.supp>

Sequential effects are ubiquitous in human and animal behavior. Across a wide range of domains, spanning stimulus detection (Posner & Cohen, 1984), perceptual identification (Garner, 1953; Treisman & Williams, 1984), probability learning (Myers, 1976), categorization (M. Jones, Love, & Maddox, 2006; Stewart, Brown,

& Chater, 2002), and decision-making (Hogarth & Einhorn, 1992), choice probabilities and response time (RT) are sensitive not just to the current stimulus or experimental condition, but also to the local history of preceding events. Sequential effects are often marginalized (literally) in statistical analyses, an approach that equates them with random noise. However, careful consideration of sequential effects can reveal a great deal about knowledge representation and learning mechanisms. This article argues that many sequential effects reflect trial-by-trial learning of statistics of the task environment, and that their detailed pattern can reveal which statistics are being tracked and how they are learned.

The present study investigates sequential effects in speeded perceptual tasks, in particular two-alternative forced choice (2AFC). This paradigm is extremely simple, involving, for example, pressing one of two keys in response to a dot that appears either above or below visual fixation. However, even this simple setting gives rise to a complex pattern of sequential effects in RT extending over many trials, which has been replicated in detail across several previous studies (e.g., Cho et al., 2002; Jentzsch & Sommer, 2002; Soetens, Boer, & Hueting, 1985). We propose these sequential effects reflect learning of statistics of the stimulus and response sequences, such as the binary sequence of dot positions (above vs. below) and of response hands (left vs. right). Our results show that the observed pattern of sequential effects can be explained in impressive detail by the combination of two simple learning mechanisms. In particular, subjects appear to be continually engaged in learning the *base rate* (i.e., the proportion of trials

Matt Jones and Tim Curran, Department of Psychology and Neuroscience, University of Colorado Boulder; Michael C. Mozer and Matthew H. Wilder, Department of Computer Science, University of Colorado Boulder.

This research was supported by a National Science Foundation (NSF) Science of Learning Center catalyst grant from the Institute of Cognitive Science at the University of Colorado to Matt Jones, Tim Curran, and Michael C. Mozer; Air Force Office of Scientific Research Grant FA9550-10-1-0177 to Matt Jones; NSF Grants BCS-0339103 and BCS-720375 to Michael C. Mozer; and NSF Grant SMA-1041755 to the Temporal Dynamics of Learning Center, an NSF Science of Learning Center, to Tim Curran and Michael C. Mozer. We thank Alaa Ahmed and Ines Jentzsch for providing their data, Brent Young for programming Experiment 1, Kyler Eastman for sharing stimulus code for Experiment 2, and the following people for research assistance and subject testing: Chris Bird, Casey Debusse, Colin Argys, Maxine Brumer, Michelle Cho, Alex Eichenbaum, Elizabeth Eustis, Megan Freeman, Robert Garcia, William Hall, Lindsey Johnson, Kathryn Keller, Emily Kleinfelder, Kyle McKelligott, Cindy Mendez, Kelly O'Neill, and Darien Taie.

Correspondence concerning this article should be addressed to Matt Jones, Department of Psychology and Neuroscience, University of Colorado Boulder, 345 UCB, Boulder, CO 80309-0345. E-mail: mcj@colorado.edu

on which each stimulus or response occurs) and the *repetition rate* (i.e., the proportion of trials that repeat vs. alternate from the previous trial). We refer to the base rate and repetition rate, respectively, as *first- and second-degree statistics* of the trial sequence, to highlight their close relationship and the fact that the repetition rate is the more complex of the two.

A common explanation for sequential effects is incremental learning, a mechanism assumed by classical theories of association learning (Estes, 1950; Rescorla & Wagner, 1972), more modern connectionist models (Rumelhart, Hinton, & Williams, 1986), and the computational framework of reinforcement learning (Sutton & Barto, 1998). Incremental learning naturally leads to *recency effects* (REs), whereby response probabilities are biased toward recent feedback, and RT is speeded whenever the current trial matches recent trials (Estes, 1957). We show here that in 2AFC, incremental learning of the base rate and repetition rate can both lead to REs, which separately are simple but which in combination produce a complex pattern of sequential effects.

We define a *first-degree RE* as a reduction in RT due to a physical match between the current stimulus and stimuli on recent trials, or between the current response and the responses on recent trials. (The question of whether these effects are driven by stimuli or responses is a further focus of this article, discussed shortly.) In the models considered here, first-degree REs arise because the estimate of the base rate is adjusted in the direction of the stimulus or response on each trial. This adjustment increases the expectancy that the same stimulus or response will occur on subsequent trials, and if this expectancy is met then processing is faster.

Likewise, we define a *second-degree RE* as a reduction of RT due to a match between the current trial and recent trials in terms of repetition versus alternation. A repetition trial is one on which the stimulus and response match those of the previous trial, whereas an alternation trial is one that mismatches the previous trial. It is well established in 2AFC that responding on repetition trials tends to be faster when recent trials were repetitions, whereas responding on alternation trials tends to be faster when recent trials were alternations (e.g., Soetens et al., 1985). The models considered here produce second-degree REs through incremental learning of the repetition rate, adjusting their estimate upward following repetition trials and downward following alternation trials.¹

There is broad evidence for both first- and second-degree REs (e.g., Myers, 1976; Soetens et al., 1985), but they have largely been studied in separate literatures and their relationship has not been closely examined. As explained below, open questions regarding each of these sequential effects can be resolved by consideration of the other. Furthermore, we show that models incorporating both first- and second-degree learning (i.e., learning of the base rate and repetition rate) provide excellent fits to the complete pattern of sequential effects observed in 2AFC tasks. We argue this finding has implications for the psychological representation of stochastic binary sequences, specifically that they are encoded in terms of base rate and repetition rate, and that both of these statistics are learned incrementally. This conclusion illustrates an important theoretical contribution of sequential effects, that they can reveal the nature of cognitive representations, because different representations lead to different patterns of sequential effects.

In addition to their implications for representation, the present data have implications for learning mechanisms. Specifically, we compare a *parallel-learning model* that learns the base rate and

repetition rate independently with a *joint-learning model* in which both statistics are learned from their combined prediction error (Rescorla & Wagner, 1972). The parallel-learning model predicts sequential effects to be a pure combination of first- and second-degree REs, whereas the joint-learning model predicts additional sequential effects arising from what amounts to competitive learning between the two statistics. The latter predictions are confirmed in the present experiments, supporting the assumptions of the joint-learning model and more generally supporting our learning-based approach to sequential effects. Learning from joint prediction error has been fundamental to explanations of many animal conditioning phenomena (e.g., blocking; Kamin, 1968), but it has not previously been applied to learning sequence statistics. Previous models of sequential effects built on first- and second-degree REs that did not explicitly connect them to learning treated them as arising from independent mechanisms and cannot explain the results found here (Cho et al., 2002; Squires, Wickens, Squires, & Donchin, 1976).

A final focus of this article is the question of whether sequential effects are driven by stimuli or responses. This is a longstanding issue in the study of sequential effects in speeded choice (e.g., Pashler & Baylis, 1991). The framework proposed here offers a novel approach to this question, because of the possibility that first- and second-degree REs originate from different processing stages. We refer to this possibility as the *separate-stages hypothesis* (although it is not meant to imply a strictly serial model). In particular, several lines of evidence from the present experiments and from reanalysis of previous data indicate that, at least in the simple binary tasks considered here, first-degree REs arise in response processing whereas second-degree REs arise in stimulus processing. That is, sequential effects in 2AFC reflect learning of *response base rates* and *stimulus repetition rates*. This conclusion is supported most strongly by the present Experiment 1, which examines sequential effects in components of event-related potential (ERP) believed to separate stimulus and response processing, specifically the amplitude of the P100 signal in visual cortex and the timing of the lateralized readiness potential (LRP) in motor cortex. These analyses indicate that sequential effects in stimulus processing are characterized by second-degree sequential effects alone, whereas sequential effects in response processing are characterized by first-degree sequential effects alone. Furthermore, model fits to the RT data give good parameter-free predictions of the ERP measures, under the assumption that the second-degree component of the model resides in stimulus processing and the first-degree component resides in response processing. These results provide a neurophysiological foundation for the psychological separation of the two learning mechanisms proposed here and, more generally, illustrate how latent variables of learning models fit to behavioral data can be used to explain sequential effects in neurophysiological data. Although the separate-stages hypothesis

¹ It is important not to confuse the present distinction between first- and second-degree REs with the common distinction in the literature between first-order and higher-order sequential effects (e.g., Soetens et al., 1985). First-order sequential effects are defined as effects of the immediately preceding trial, and higher-order effects are from earlier trials. As the results reported here demonstrate, the two distinctions are independent. That is, both first- and second-degree REs contribute to both first- and higher-order sequential effects.

might seem surprising a priori, we explain in the Separating Stimulus- and Response-Based Sequential Effects section how it is consistent with previous theories of stimulus- versus response-based sequential effects.

The remainder of this article is organized as follows. We begin by reviewing evidence for first- and second-degree REs. Then we show how a model that produces both, through learning of the base rate and repetition rate, can explain the detailed pattern of sequential effects found in previous 2AFC experiments. Next, data from three previous studies (Jentzsch & Sommer, 2002; Maloney, Dal Martello, Sahn, & Spillmann, 2005; Wilder, Jones, Ahmed, Curran, & Mozer, 2013) are reanalyzed to motivate the hypothesis that first- and second-degree REs are dissociated into response and stimulus processing, respectively. Two new experiments are then reported that test the proposed theoretical framework. Experiment 1 varied the repetition rate between sessions, to evaluate model predictions concerning the relationship between learning and sequential effects. The RT data from this experiment are used to evaluate different possible mechanisms by which first- and second-degree statistics are learned, in support of the joint-learning model. This model is then found to provide good predictions of the ERP data, supporting the separation of first- and second-degree learning into response and stimulus processing. Experiment 2 provides a further test of this separation, by varying the difficulty of the stimulus discrimination and assessing the impact on first- and second-degree sequential effects.

Overall, the results provide a remarkably clear and coherent picture of the origins of sequential effects in 2AFC tasks. According to this view, sequential effects are a signature of learning certain statistics of the trial sequence. By considering the details of this signature, one can uncover which statistics are learned, and how. We conclude that sequential effects offer a powerful tool for revealing learning mechanisms, the representations on which they operate, and their neurophysiological underpinnings.

First-Degree Recency Effects

Perhaps the most ubiquitous sequential effect observed in psychology is a facilitation of performance when the current stimulus or response matches those of recent events. We refer to such an effect as a first-degree RE because it depends on the identities of individual events and not on sequential properties (such as whether events are repetitions or alternations). First-degree recency effects are observed in essentially any behavioral paradigm involving repeated trials and a fixed set of stimuli or responses (see references in the Introduction section).

An illustrative example comes from the paradigm of probability learning, in which a subject repeatedly chooses one of a few response options (usually two), without any cues, and then is told which response was correct for that trial. Research in probability learning was originally concerned with people's ability to discern differential reward rates for different responses, but it was discovered that behavior in this paradigm is dominated by the recent trial history (see Myers, 1976, for a review). For example, Engler (1958) ran a two-choice task in which responses 1 and 2 were each correct on 50% of trials. He found that subjects chose 1 on 60% of trials following reinforcement of 1 on the previous trial, whereas they chose 1 on only 40% of trials following reinforcement of 2.

This effect is not limited to the previous trial, but can instead extend back many trials (e.g., Gambino & Myers, 1967).

First-degree REs are well explained by incremental learning models that update expectancies or associations following each trial. Incremental learning models have a long history in learning theory (Estes, 1950; Hebb, 1949) and are fundamental to modern neurophysiological and computational theories of reinforcement learning (Schultz, Dayan, & Montague, 1995; Sutton & Barto, 1998). As explained in the A Simple Two-Component Model section, incremental learning naturally leads to an exponentially decaying influence of past experience on current behavior, with more recent events having the greatest impact (Estes, 1957).

Although this mechanistic explanation of REs suggests they are merely byproducts of the learning process, there is also reason to suspect they have a functional or rational purpose. When an environment is nonstationary, meaning that contingencies or outcome probabilities change over time, it is generally optimal to weight decisions on more recent events (Anderson & Schooler, 1991; Cuthill, Kacelnik, Krebs, Haccou, & Iwasa, 1990; Flood, 1954; M. Jones & Sieck, 2003; Real, 1991). This normative view of REs suggests they should be adaptive, such that they become stronger when autocorrelation in the outcome sequence is increased (M. Jones & Sieck, 2003; Yu & Cohen, 2009). In the context of incremental-learning models, adaptation of REs could naturally be explained by adjusting the learning rate. A faster learning rate is adaptive for highly nonstationary environments because it allows more rapid tracking of changes in outcome probabilities, and it also produces stronger REs.

M. Jones and Sieck (2003) found evidence for an adaptive first-degree RE in a category-learning task with two categories. They manipulated the autocorrelation in the sequence of correct categories, so that the repetition probability was 30% (negative autocorrelation), 50% (independent trials), or 70% (positive autocorrelation), and found that subjects in the positive condition exhibited the strongest REs while subjects in the negative condition exhibited the weakest. They interpreted this finding under the normative explanation given above: REs improve performance in the positive autocorrelation condition, whereas they hurt performance in the negative condition. However, M. Jones and Sieck also presented a computational model embodying a different explanation. This model learned to use the outcome (i.e., correct category) on the previous trial as a cue, in addition to the current stimulus, for predicting the current outcome. Thus, rather than adapting the strength of the first-degree RE (e.g., by changing the learning rate), the model directly learned the repetition rate of the outcome sequence. Hence, we have two competing hypotheses for why first-degree REs adapt in autocorrelated environments: the *adaptation hypothesis*, whereby first-degree REs become stronger when nonstationarity is greater, and the *direct hypothesis*, whereby subjects directly learn the repetition rate in the trial sequence. The contrast between these explanations bears on the fundamental question of functionality of sequential effects in learning. Testing between them requires understanding the contribution of learning of second-degree sequential statistics such as repetition rates.

Second-Degree Recency Effects

Evidence for learning of repetition rates comes from the pattern of sequential effects in RT commonly observed in 2AFC tasks. In

this paradigm, RT appears to depend primarily on whether the current trial is a repetition or an alternation—that is, whether the stimulus and response are the same as on the previous trial. On repetition trials, responding is faster if recent trials were primarily also repetitions, whereas the reverse is true on alternation trials. We refer to this phenomenon as a second-degree RE, because it depends on the second-degree match between current and recent trials (i.e., whether each is a repetition or an alternation).

Figure 1 shows two typical examples of this pattern (Cho et al., 2002; Jentsch & Sommer, 2002). In each of these experiments, subjects made speeded key presses to one of two stimuli (lower or uppercase “O,” or a dot above or below visual fixation). RT for correct responses is plotted as a function of trials $n - 3$ through n (where n is the current trial), yielding 16 possible histories, each shown as a sequence of repetitions (R) and alternations (A). For example, AAAR indicates cases in which the current trial is a repetition and the previous three trials were alternations. In both experiments, there is a clear trend for RT on repetition trials to be faster when repetitions are more frequent and more recent in the trial history, indicated by the rising pattern in the left half of each graph. Likewise, RT for alternation trials is faster when alterna-

tions are more frequent and more recent, as indicated by the falling pattern in the right half of each graph.

This pattern of second-degree REs is consistent with the same explanations given above for first-degree REs, but applied to learning of the repetition rate. That is, subjects estimate the probability that the current trial will be a repetition versus an alternation, based on the past history of repetitions and alternations. As with first-degree REs, this explanation can be framed mechanistically, in terms of incremental learning, or computationally, in terms of optimal Bayesian inference assuming a nonstationary environment (Wilder, Jones, & Mozer, 2010; Yu & Cohen, 2009).

However, second-degree REs alone cannot explain the patterns in Figure 1. For example, consider the histories ARAR and RAAR. In both cases, the current trial is a repetition, and there is exactly one repetition in the previous three trials. (Because trials were independently sampled in these experiments, both data points represent approximately half repetitions and half alternations on each of trials $n - 4$ and earlier.) This previous repetition is more recent in the ARAR case, and hence, ARAR should show shorter RTs than RAAR. In contrast, both experiments in Figure 1 show the opposite pattern, which we also replicate in the experiments reported below.

This discrepancy from the predictions of second-degree REs can be explained by a contribution of first-degree REs. This can be seen by translating the second-degree history representations to equivalent first-degree representations. The sequence ARAR for trials $n - 3$ through n corresponds to a sequence XYX for trials $n - 4$ through n , where X and Y are the two first-degree trial types (i.e., the two possible stimuli or responses), with the final X representing the value of trial n . Likewise, the RAAR history corresponds to XXY . In terms of this first-degree representation, the two sequences differ only in trial $n - 3$, which matches trial n for XXY but not for XYX . Thus, first-degree REs should contribute to faster responding for the XXY /RAAR history.

As a second example, consider RARR versus AARR. Second-degree REs would produce faster responding to RARR, because of the additional second-degree match (the R on trial $n - 3$), but the data contradict this prediction. Translating the second-degree history representations to first-degree representations yields YYX for RARR and XYX for AARR. Thus, first-degree REs could explain the faster responding to XYX /AARR, due to the extra first-degree match (the X on trial $n - 4$).

The picture of sequential effects in 2AFC therefore appears somewhat complicated. The bulk of the variation in RT by trial history is well explained by second-degree REs, but there are reliable discrepancies that are consistent with a contribution of first-degree REs.

To summarize the last two sections, previous research in a variety of domains offers strong support for both first- and second-degree REs. These effects are consistent with explanations based on learning of base rates and repetition rates, respectively. However, neither offers a complete theory of sequential effects, even in simple tasks. Understanding the relationship between first- and second-degree REs might explain open questions regarding each, including whether first-degree REs adapt to autocorrelated environments and whether deviations from the predictions of second-degree REs are due to first-degree effects. Therefore, we turn now to a modeling framework that describes how the learning mecha-

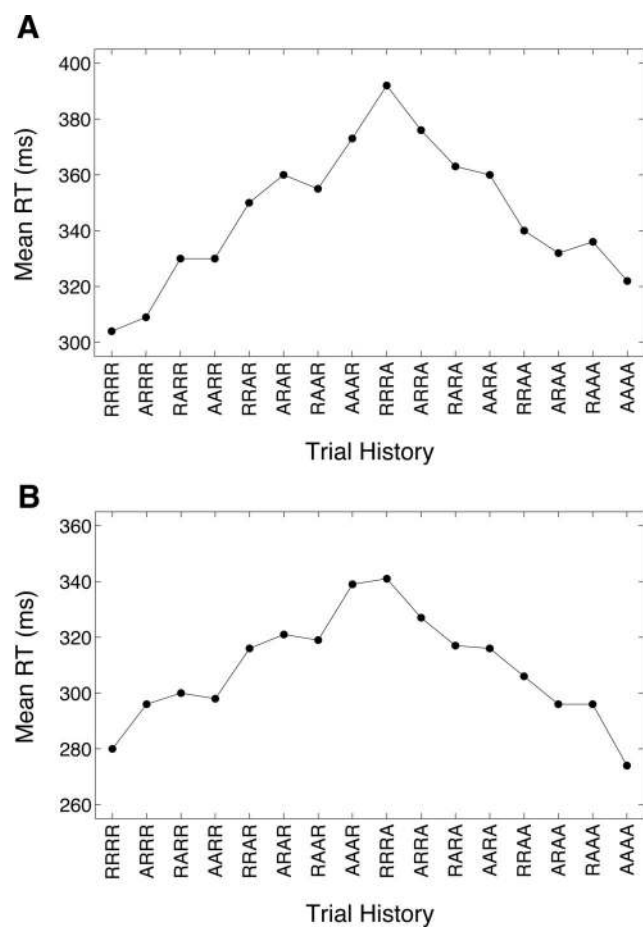


Figure 1. Sequential effects in response time (RT) in two-alternative forced choice. Trial histories indicate trials $n - 3$ through n (current trial) as repetitions (R) or alternations (A). A: Cho et al. (2002). B: Jentsch and Sommer (2002), Experiment 1.

This document is copyrighted by the American Psychological Association or one of its allied publishers. This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

nisms underlying these two sequential effects might combine to influence behavior.

A Simple Two-Component Model

Consider a 2AFC task with two stimuli mapped to two responses. Let E_n represent the stimulus-response pair on trial n (e.g., a left button press for a stimulus above fixation, or a right press for a stimulus below fixation). We set aside the distinction between stimuli and responses for now, implicitly limiting the model to trials on which the current and recent responses are all correct. To simplify notation in a way that will be clear shortly, we encode E_n as ± 1 .

To motivate the full model, assume first that the subject learns only an estimate of the base rate—that is, the expected value of E_n . A simple and standard incremental learning rule updates this estimate after each trial, in proportion to the prediction error:

$$\Delta w_{\text{base}} = \epsilon_{\text{base}}(E_n - w_{\text{base}}). \quad (1)$$

Here, w_{base} is the estimate of the base rate, and ϵ_{base} is a learning-rate parameter ($0 < \epsilon_{\text{base}} \leq 1$). It is easy to show that iteration of this rule leads the value of w_{base} at the beginning of trial n to be a weighted average of all past events, with weights an exponentially decreasing function of lag (Estes, 1957):

$$w_{\text{base}} = \sum_{k=1}^{n-1} \epsilon_{\text{base}}(1 - \epsilon_{\text{base}})^{k-1} E_{n-k}. \quad (2)$$

Here k indexes the lag from the current trial to each past trial. For simplicity, we assume that w_{base} is initialized to 0 at the start of the experiment (i.e., neutral expectation between the two outcomes).

In advance of each trial, the current value of w_{base} acts as an expectancy for E_n . A positive value of w_{base} corresponds to an expectancy for $E_n = 1$, and a negative value corresponds to an expectancy for $E_n = -1$. A natural assumption when modeling RT is that responding will be faster to the extent the expectancy matches the actual outcome (Bertelson, 1961; Kirby, 1976; Lam-ing, 1968). Assuming a simple linear relationship with no bias between the two responses yields the following prediction for mean RT (where β_0 and β_{base} are scaling parameters):

$$RT = \beta_0 - \beta_{\text{base}} w_{\text{base}} E_n. \quad (3)$$

Thus, when w_{base} is positive, RT will be faster (i.e., smaller) if $E_n = 1$ than if $E_n = -1$. Likewise, if w_{base} is negative, RT will be faster if $E_n = -1$. Combining Equations 2 and 3 leads to a direct expression for predicted mean RT in terms of the trial history:

$$RT = \beta_0 - \beta_{\text{base}} \sum_{k=1}^{n-1} \epsilon_{\text{base}}(1 - \epsilon_{\text{base}})^{k-1} M_1(n - k, n). \quad (4)$$

Here, $M_1(n - k, n) = E_{n-k} \cdot E_n$ denotes the *first-degree match* between trials $n - k$ and n . When these trials have the same stimulus and response, $M_1(n - k, n)$ will equal 1, and when they mismatch, $M_1(n - k, n)$ will equal -1 . Writing the model's prediction in this way highlights how RT depends on these matches between the current and past trials. RT is reduced when there are more matches, and matches to more-recent trials have a greater effect. In short, this model shows how incremental learning of the base rate naturally produces exponentially decaying first-degree REs.

The same analysis applies to second-degree REs arising from incremental learning of the repetition rate. First, note that $E_{n-1} \cdot E_n$ encodes the second-degree outcome of trial n : It equals 1 when trial n is a repetition and -1 when it is an alternation (this is just a notational convenience, due to the ± 1 encoding). Therefore, incremental learning of the repetition rate follows an update rule analogous to Equation 1, with the first-degree outcome, E_n , replaced by the second-degree outcome, $E_{n-1} \cdot E_n$:

$$\Delta w_{\text{rep}} = \epsilon_{\text{rep}}(E_{n-1}E_n - w_{\text{rep}}). \quad (5)$$

Iteration of this rule leads the estimate of the repetition rate, w_{rep} , to be an exponentially decaying weighted average of past trials, this time in terms of whether they were repetitions or alternations:

$$w_{\text{rep}} = \sum_{k=1}^{n-2} \epsilon_{\text{rep}}(1 - \epsilon_{\text{rep}})^{k-1} E_{n-k-1}E_{n-k}. \quad (6)$$

(The sum terminates at $k = n - 2$ because learning about repetitions can only begin after trial 2.)

The current value of w_{rep} generates an expectancy for each upcoming trial, just as does w_{base} , but with a direction depending on the identity of the previous trial. When w_{rep} is positive, the expectancy is in the direction of E_{n-1} (expectancy of repetition), and when w_{rep} is negative, it is in the opposite direction (expectancy of alternation). Assuming that both expectancies combine to determine mean RT, we have


$$RT = \beta_0 - \beta_{\text{base}} w_{\text{base}} E_n - \beta_{\text{rep}} w_{\text{rep}} E_{n-1} E_n. \quad (7)$$

Combining Equations 2, 6, and 7 yields an explicit expression for the predicted RT as a function of trial history:

$$RT = \beta_0 - \beta_{\text{base}} \sum_{k=1}^{n-1} \epsilon_{\text{base}}(1 - \epsilon_{\text{base}})^{k-1} M_1(n - k, n) - \beta_{\text{rep}} \sum_{k=1}^{n-2} \epsilon_{\text{rep}}(1 - \epsilon_{\text{rep}})^{k-1} M_2(n - k, n). \quad (8)$$

Here, $M_2(n - k, n) = E_{n-k-1}E_{n-k} \cdot E_{n-1}E_n$ denotes the *second-degree match* between trials $n - k$ and n , equal to 1 when both are repetitions or both are alternations, and -1 otherwise. Thus, the model's predictions are a sum of first- and second-degree REs, both of which decay exponentially as a function of lag. The rates of decay are determined by the learning-rate parameters, ϵ_{base} and ϵ_{rep} . The β_{base} and β_{rep} parameters (both constrained to be non-negative) determine the magnitudes of the two REs, via the influences of first- and second-degree expectancies on RT. (We later interpret these parameters as speeds of response and stimulus processing, respectively, in the context of the separate-stages hypothesis.) We refer to Equation 4 as the parallel-learning model, because it produces sequential effects from simultaneous and independent learning of the base rate and repetition rate.

Figure 2 shows how the parallel-learning model can produce excellent fits to the full pattern of sequential effects in a 2AFC task. Figure 2A presents a fit of the model to the data from Figure 1A (Cho et al., 2002). The model was fit by least squares to the mean RTs of the 16 trial histories. Predictions were derived directly from Equation 8, with the sums truncated after $k = 4$ for first-degree effects and $k = 3$ for second-degree effects (i.e., at the longest lags determined by these trial histories). Because Cho et al. (2002) used independently sampled trials, the contributions from all



earlier lags should approximately cancel out. Model parameters were constrained to satisfy $\beta_0 \geq 0$, $\beta_{base} \geq 0$, $\beta_{rep} \geq 0$, $0 < \epsilon_{base} \leq 1$, and $0 < \epsilon_{rep} \leq 1$ (these constraints were used in all later model fits unless noted otherwise). The model explains 98.98% of the variance in mean RT across the 16 trial histories, on 5 free parameters. The estimated values of these parameters are given in Table 1.

To facilitate interpretation of Figure 2, the trial histories are labeled using both first- and second-degree representations. The first-degree representations indicate the identity of the stimulus and response (which are perfectly correlated on correct trials) on trials $n - 4$ through n . We collapse over the symmetry between the two possible trial types, using the convention that X always indicates the identity of the current trial (e.g., sequences 12221 and 21112 are combined as YXXYX). Therefore, the information in the first-degree history is whether each of trials $n - 4$ through $n - 1$ is a first-degree match with trial n . The second-degree history indicates whether each of trials $n - 3$ through n is a repetition (R) or an alternation (A). It should be apparent that both representations contain exactly the same information. We refer to the trial histories under both of these representations as *4-deep histories*. We use both first- and second-degree history representations because the relationship between them is nontrivial, which is also why the two simple REs assumed by the model can give rise to complex sequential effects.

To illustrate the separate contributions of first- and second-degree REs to the model's predictions, the model was broken down into its first- and second-degree components, corresponding to the second and third terms on the right side of Equation 7 (equivalently, the second and third main terms on the right side of Equation 8). Figures 2B and 2C plot the predictions of these separate components (the β_0 term is included in both of these figures, so that the partial predictions are aligned on the grand mean RT). These separate predictions are the patterns produced by first- and second-degree REs, respectively. Second-degree REs generate the pyramid pattern that characterizes the overall shape of the data, and first-degree REs explain the deviations from that pattern (second-degree effects dominate because $\beta_{rep} > \beta_{base}$; see Table 1). Adding together the predictions of the two components (and subtracting the duplicated β_0) yields the predictions of the full model.

In conclusion, the complex pattern of sequential effects observed empirically appears to be well explained by two simple REs. The parallel-learning model formalizes this idea, based on the following three assumptions: (1) estimates of the base rate and repetition rate are both learned incrementally, updated from one trial to the next; (2) the current estimates each generate an expect-

Figure 2. Fit of parallel learning model to data of Cho et al. (2002). Mean response time (RT) is shown for each 4-deep trial history. First-degree history representations (e.g., XXXXX) indicate identities of trials $n - 4$ through n , with X always defined as the value of trial n . Second-degree history representations (e.g., RRRR) indicate status of trials $n - 3$ through n as repetitions (R) or alternations (A). A: Fit of full model (Equation 8). B: Predictions of first-degree (base-rate learning) component of model only, plus intercept term (β_0). C: Predictions of second-degree (repetition-rate learning) component of model only, plus intercept term (β_0).

Table 1
Fits of Parallel-Learning Model to Sequential Effects in Previous Data Sets

Data set	Parameter						Var
	β_0	β_{base}	ϵ_{base}	β_{rep}	ϵ_{rep}	δ	
Cho et al. (2002)	345.75	23.20	.301	32.62	.472	—	.9898
Jentsch & Sommer (2002), Experiment 1, LRPr							
Full model	71.88	23.20	.592	4.21	.527	—	.8826
First-degree model	71.88	27.04	.517	—	—	—	.8703
Second-degree model	71.88	—	—	9.65	.793	—	.2098
Jentsch & Sommer (2002), Experiment 1, sLRP							
Full model	234.06	0.00	—	28.12	.297	19.06	.9392
First-degree model	234.06	128,602.70	.999931	—	—	128,612.89	.8593
Second-degree model	234.06	—	—	28.12	.297	19.06	.9392
Maloney et al. (2005), Experiment 1							
Full model	88.34	12.99	.100	18.42	.139	-0.31	.9834
First-degree model	88.34	14.51	.581	—	—	6.82	.6453
Second-degree model	88.34	—	—	11.71	.307	-1.61	.9039
Maloney et al. (2005), Experiment 2							
Full model	87.94	0.00	—	2.99	.607	1.91	.9733
First-degree model	87.94	306,770.55	.999994	—	—	306,770.64	.9070
Second-degree model	87.94	—	—	2.99	.607	1.91	.9733
Wilder et al. (2013), Experiment 2							
Full model	2.57	1.15	.340	0.00	—	—	.9797
First-degree model	2.57	1.15	.340	—	—	—	.9797
Second-degree model	2.57	—	—	0.28	.837	—	.2124

Note. Var column indicates proportion of explained variance across trial histories. Units for β and δ parameters are milliseconds for Cho et al. (2002) and Jentsch and Sommer (2002), degrees for Maloney et al. (2005), and centimeters for Wilder et al. (2013); ϵ parameters are dimensionless. Dashes indicate parameters held to zero (β) or irrelevant (ϵ , when corresponding β parameter equals 0). Bold rows indicate the model taken as the best explanation for each data set. LRPr = latency from lateralized readiness potential (LRP) onset to response; sLRP = latency from stimulus onset to LRP onset.

tancy for each upcoming trial; and (3) responding is faster when these expectancies more strongly match the actual outcome. Interestingly, each RE taken alone produces additive effects of all past trials, and the model combines the two REs additively. Nevertheless, the model's predictions (and the data) show strongly nonadditive effects of past trials. The explanation is that the two learning mechanisms operate on different types of trial representation (physical stimulus-response identities vs. repetition/alternation) which themselves are nonlinearly related. This result thus supports an important link between sequential effects, learning, and psychological representation of binary event sequences: The complex sequential effects observed in 2AFC are the product of simple learning mechanisms operating simultaneously on different representations.

Previous investigators have offered models that combine first- and second-degree sequential effects, but in a more ad hoc manner. For example, Squires et al. (1976) assumed exponentially decaying first-degree REs, combined with a special expectancy for alternation if at least two of the past three trials were alternations. Cho et al. (2002) explored a complex set of detectors that develop expectancies for upcoming trials based on various patterns of past trials. Their most successful model combined a detector that develops an expectancy for each trial type based on past occurrences of that trial type, and a second detector that develops expectancy for alternations based on past occurrences of alternations. These detectors, respectively, generate exponentially decaying first- and second-degree REs, and together they behave almost identically to the parallel-learning model of Equation 8. An advantage of our modeling approach is that it offers a reason that these detectors should give the best fits out of all the combinations Cho et al.

tested. By interpreting sequential effects as due to learning, our theory shows how first- and second-degree REs can arise from the same simple learning mechanism. The only difference between the two REs is that they operate on different representations of each event (physical identity or repetition/alternation status) and thus learn different statistics of the sequence (base rate or repetition rate, respectively). As the experiments and analyses reported below demonstrate, this unified treatment allows for modeling of interactions between the two learning mechanisms, which can explain further patterns of sequential effects not explained by the two REs treated separately.

Separating Stimulus- and Response-Based Sequential Effects

Thus far, we have not distinguished between the stimulus and response on each previous trial, in terms of their role in driving sequential effects. In tasks where stimulus-response mappings are one-to-one, stimuli and responses are highly confounded, and perfectly so on correct trials. Nevertheless, there is the question of which underlies sequential effects. For example, does the sequence of past stimuli affect perception of the current stimulus, or does the sequence of past responses affect execution of the current response? There is a long history of research attempting to answer this question by a variety of methods, including analyzing trials after errors (Laming, 1968; Rabbitt & Rogers, 1977), using many-to-one stimulus-response mappings (Bertelson, 1965; Pashler & Baylis, 1991; Soetens, 1998), varying the stimulus-response mapping across subjects (M. Jones, 2009), using quasicontinuous stimulus and response values (Jesteadt, Luce, & Green, 1977; Petzold,

1981), and providing stochastic feedback (M. Jones, 2009; M. Jones et al., 2006).

The two-component theory advanced here suggests a novel possibility, that stimulus and response processing could be responsible for different types of sequential effects. That is, if first- and second-degree REs are due to separate psychological mechanisms, then it is possible they could be dissociated on the basis of stimulus versus response processing. It turns out that existing data are well explained by such a dissociation, specifically that first-degree REs arise primarily in response processing, whereas second-degree REs arise primarily in stimulus processing. As noted above, we refer to this as the separate-stages hypothesis. This proposal at first appears arbitrary (e.g., the reverse hypothesis is just as sensible a priori), but we demonstrate in this section that it does a surprisingly good job of explaining past findings, and it is further supported by the new experiments reported below.

Consider first the 2AFC data of Jentsch and Sommer (2002, Experiment 1). The sequential effects in RT from that experiment are shown in Figure 1B. The parallel-learning model produces good fits to these data (95.35% variance explained; not shown), but here we focus on a different aspect of their data. Jentsch and Sommer also recorded ERPs during their task, which they used to compute the LRP, a hemispheric asymmetry of activation in motor cortex that shows a spike of activity before a motor response is executed. If each response is executed by one of the two hands (each primarily controlled by the opposite hemisphere), the onset of the LRP spike can be used to divide total RT into separate stimulus- and response-processing stages (Coles, 1989; de Jong, Wierda, Mulder, & Mulder, 1988). The time from stimulus onset to LRP onset is assumed to correspond to stimulus identification and stimulus-response mapping, and the time from LRP onset to the response is assumed to correspond to execution of the response motor plan. When the stimulus-response mapping is trivial (as was the case in Jentsch & Sommer's, 2002, Experiment 1), the stimulus-LRP time (sLRP) can be taken as a measure of stimulus processing time, with the LRP-response time (LRPr) taken as a measure of response processing time.

Figure 3 (solid lines) shows the LRPr and sLRP times obtained by Jentsch and Sommer (2002), as a function of 4-deep trial history. The patterns of sequential effects in these two stages of processing are quite different, with LRPr more consistent with first-degree REs and sLRP more consistent with second-degree REs. To test this dissociation, the parallel-learning model was divided into separate first- and second-degree components, and each component was fit to both LRPr and sLRP.

The prediction of LRPr by the first-degree component is given by omitting the second-degree component from Equation 7, or equivalently by setting $\beta_{rep} = 0$:

$$LRPr = \beta_{LRPr} - \beta_{base} w_{base} E_n, \tag{9}$$

where β_{LRPr} is an intercept term corresponding to some portion of the intercept β_0 in the full model. This first-degree model explains 87.03% of the variance in LRPr across the 16 trial histories, on 3 free parameters (β_{LRPr} , β_{base} , and ϵ_{base} ; see dashed line in Figure 3A). When the corresponding second-degree model is fit to the same data (i.e., Equation 7 with $\beta_{base} = 0$), it explains only 20.98% of the variance (see Table 1). The full model (i.e., with both first- and second-degree components) can explain 88.26%, only a small improvement over the first-degree model with 2 extra

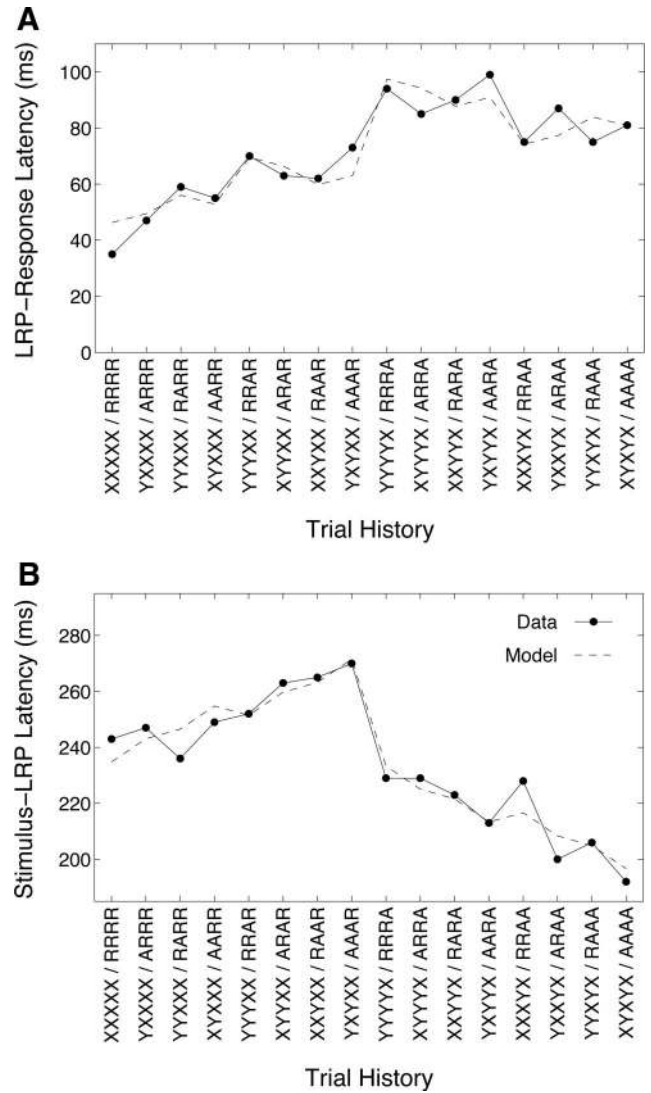


Figure 3. Fits of the parallel-learning model to electroencephalography data of Jentsch and Sommer (2002, Experiment 1), under the assumption that first-degree recency effects operate only in response processing, and second-degree recency effects operate only in stimulus processing. A: Latency from lateralized readiness potential (LRP) onset to response. Only the first-degree component of the model is used to fit these data. B: Latency from stimulus onset to LRP onset. Only the second-degree component of the model, together with a bias term representing an alternation advantage, is used to fit these data.

free parameters (β_{rep} and ϵ_{rep}). Therefore, sequential effects in response processing time appear to be well-explained by first-degree REs alone.

The analysis for sLRP is complicated by the large difference between repetition and alternation trials (compare left and right halves of Figure 3B). Stimulus processing appears to be significantly faster for alternation trials. We refer to this phenomenon, as well as a similar pattern in RT observed in Experiments 1 and 2 below, as an *alternation advantage*. We argue below, via further modeling analyses and data fits, that the alternation advantage is localized to stimulus processing and has important implications for

understanding the learning mechanisms producing sequential effects. However, for the moment, we set this issue aside by introducing an artificial bias parameter to the parallel-learning model, δ , which is added to repetition RTs and subtracted from alternation RTs (cf. Equation 7):

$$RT = \beta_0 - \beta_{\text{base}} w_{\text{base}} E_n - \beta_{\text{rep}} w_{\text{rep}} E_{n-1} E_n + \delta E_{n-1} E_n. \quad (10)$$

As noted above, $E_n E_{n-1}$ equals 1 on repetition trials and -1 on alternation trials, so that $\delta > 0$ produces an alternation advantage. The δ term is similar to the term involving w_{rep} except that δ is constant, meaning its contribution is the same for all repetition trials and for all alternation trials, whereas w_{rep} depends on the full trial history. We henceforth use Equation 10 instead of Equation 7 in defining the parallel-learning model.² We stress that the bias parameter is a wholly ad hoc assumption; the joint-learning model presented later in this article explains the alternation advantage without this assumption.

Because the alternation advantage is claimed to reside in stimulus processing, we include the bias term for modeling sLRP (but not LRP; see Equation 9). Thus, the prediction for sLRP by the second-degree model is given by

$$sLRP = \beta_{\text{sLRP}} - \beta_{\text{rep}} w_{\text{rep}} E_{n-1} E_n + \delta E_{n-1} E_n. \quad (11)$$

This model explains 93.92% of the variance in sLRP across the 16 trial histories, on 4 free parameters (β_{sLRP} , β_{rep} , ϵ_{rep} , and δ ; see dashed line in Figure 3B). When the first-degree model is fit to the same data (with bias term included), it explains 85.93% of the variance (see Table 1). However, the latter model achieves this fit from the bias term alone, together with degenerate parameters under which the first-degree RE partially mimics a second-degree RE.³ Critically, when the full model of Equation 10 is fit to the sLRP data, the optimal value of β_{base} is zero. Thus, the full model reduces to the second-degree model; allowing first-degree effects does not improve the fit. Therefore, sequential effects in stimulus processing time appear to be best explained by second-degree REs alone, together with an alternation advantage.

To be clear, the β_{base} , w_{base} , β_{rep} , and w_{rep} variables used in predicting LRP latencies (Equations 9 and 11) are meant as the same variables as those used in predicting RT (Equation 10). The same is true of the ϵ_{base} and ϵ_{rep} parameters that underlie the dynamics of w_{base} and w_{rep} . Indeed, we demonstrate in Experiment 1 below that these parameters can be estimated from RT data and directly used to predict LRP latencies.

Further evidence for the separate-stages hypothesis comes from reanalysis of data from Maloney et al. (2005). Subjects in these experiments were shown ambiguous motion quartets, consisting of brief visual display of a pair of dots followed immediately by brief display of a second pair. All four dots lay on the circumference of an unseen circle, and together they produced an apparent motion effect in which a single pair of dots appeared to rotate about the center of that circle. Whether the perceived rotation was clockwise or counterclockwise depended on the angle between the pairs' positions on the circle, measured clockwise from the first pair to the second. Angles below 90 degrees tended to induce perceived clockwise motion, and angles above 90 degrees tended to induce perceived counterclockwise motion.

In their first experiment, Maloney et al. (2005) presented subjects with a series of ambiguous motion quartets, varying in angle,

and had subjects report the direction of perceived motion following each. Trials were analyzed according to the sequence of responses on trials $n - 4$ through $n - 1$. For each of these 16 histories, standard psychophysical techniques were used to estimate each subject's threshold on trial n , defined as the angle that led both directions of motion to be perceived equally often. This threshold was found to depend strongly on the trial history. For example, when the previous 4 trials were all perceived as clockwise, there was a strong bias for subjects to perceive clockwise motion on the current trial, such that the angle could be as large as 100 degrees and subjects would still respond clockwise about half the time.

Mean thresholds as a function of stimulus history are shown in Figure 4A (solid line). Maloney et al. (2005) did not collapse over the symmetry between the two responses, so the first-degree history representations in this figure are in terms of the actual responses—c for clockwise and C for counterclockwise—as opposed to the abstract XY labeling used in Figures 2 and 3. Also unlike the previous figures, these histories do not include the current trial; the first-degree representations indicate trials $n - 4$ through $n - 1$, and the second-degree representations indicate trials $n - 3$ through $n - 1$. The function being plotted can be interpreted as an index of subjects' expectancy on trial n , with larger values indicating a stronger expectancy or bias for clockwise motion.

Maloney et al. (2005) interpreted these results as reflecting tendencies to continue patterns of repetition or alternation; that is, as second-degree REs. For example, subjects showed strong biases toward c following cccc and toward C following CCCC, consistent with expectancy of a repetition following many repetitions. They also showed biases toward C following CcCc and toward c following cCcC, consistent with expectancy of an alternation following many alternations. However, we show here that the data are better explained by a combination of first- and second-degree REs, using the parallel-learning model.

To apply the parallel-learning model to this task, we assumed the threshold following each trial history is linearly (negatively) related to the subject's expectancy for counterclockwise motion, with expectancies derived from both first- and second-degree learning. Following the same reasoning that led to Equation 10, and writing $E_n = -1$ for clockwise and 1 for counterclockwise, yields an expression for the model's predictions in terms of the trial history:

² Allowing a bias in fitting the data from Cho et al. (2002) slightly improves the fit to 99.07% variance explained, on 6 free parameters. As will become clear in the analysis of the joint-learning model (see the Cue Competition Explanation for the Alternation Advantage section), the parallel-learning model's need for a bias term depends on the relative magnitudes of β_{base} and β_{rep} , which can vary across experiments.

³ Specifically, setting ϵ_{base} very close to 1 and β_{base} very large makes the $k = 1$ term of the first-degree RE very large, the $k = 2$ term moderate, and longer-lag terms negligible (see Equation 8). The contribution of the $k = 1$ term is exactly opposite that of the bias term. The large value of δ thus cancels out the $k = 1$ term, leaving a net bias of $\delta - \beta_{\text{base}} \epsilon_{\text{base}} = 19.07$ ms, in close agreement with the fit of the second-degree model. The contribution of the $k = 2$ term is identical to that of a second-degree RE at $k = 1$ (because trials $n - 2$ and n are a first-degree match if and only if trials $n - 1$ and n are a second-degree match). Therefore, under these parameters, the model amounts to an alternation advantage plus a second-degree RE restricted to lag 1.

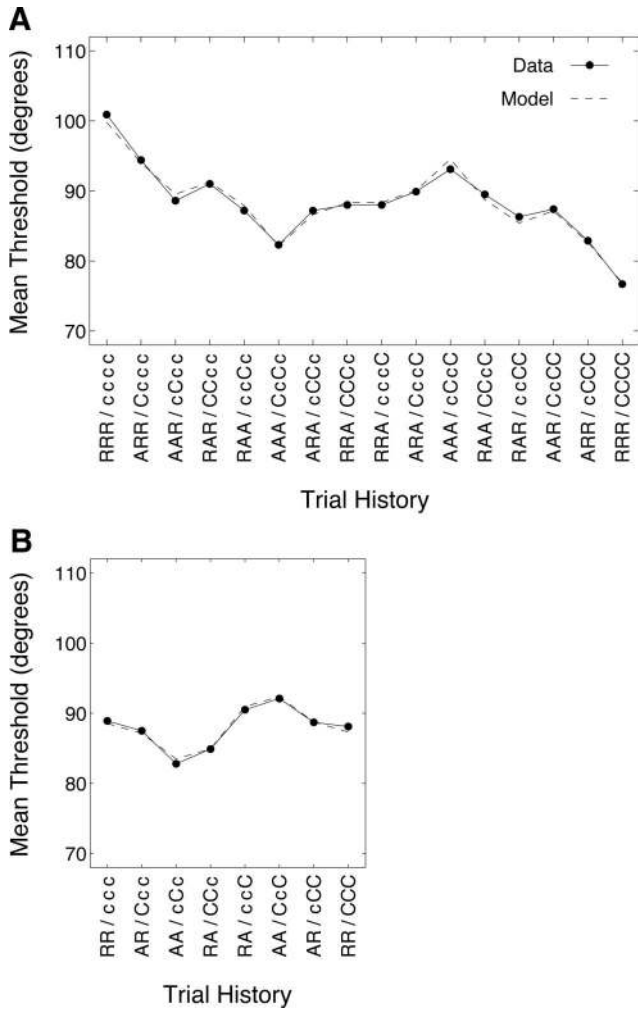


Figure 4. Fits of parallel-learning model to data of Maloney et al. (2005). Trial histories are indicated both as sequences of clockwise (c) and counterclockwise (C) responses for trials $n - 4$ through $n - 1$, and as sequences of repetitions (R) and alternations (A) for trials $n - 3$ through $n - 1$. Threshold is the stimulus angle estimated to produce equal responding of c and C on trial n . Larger values indicate biases to respond c. A: Fit of parallel-learning model to their Experiment 1. B: Fit to their Experiment 2, in which no response was given during the previous three trials. Only the second-degree component of the model, together with the alternation bias term, is used to fit the Experiment 2 data.

$$Threshold = \beta_0 - \beta_{base}w_{base} - \beta_{rep}w_{rep}E_{n-1} + \delta E_{n-1}. \quad (12)$$

This equation embodies the same principles as the model’s predictions for RT tasks (Equation 10). Expectancy for the upcoming trial is a combination first- and second-degree REs, both determined by exponentially weighted sums of past trials (Equations 2 and 6). The first-degree expectancy, w_{base} , is based on the individual identities of past trials. The second-degree expectancy, $w_{rep}E_{n-1}$, is based on whether past trials were repetitions or alternations, and its direction depends on the identity of the most recent trial (E_{n-1}). The final term in Equation 12 produces a bias for alternation of magnitude δ , by shifting the threshold toward E_{n-1} .

The model’s fits to Experiment 1 of Maloney et al. (2005) are shown in Figure 4A (dashed line). The model explains 98.34% of the variance across the 16 trial histories, with 6 free parameters. The fit is nearly unchanged without the alternation-bias parameter (98.23% variance explained with 5 free parameters), but we include it for consistency with the fit of Maloney et al.’s second experiment, presented shortly. Critically, the fit is markedly worse if either the first- or the second-degree component of the model is removed (see Table 1). Although formal model comparison is not possible without individual subjects’ data, we take these fits to indicate that sequential effects in this task are best explained as a combination of both types of REs.

In Experiment 2 of Maloney et al. (2005), subjects only responded on every fourth trial. Stimuli for the intervening trials had sufficiently extreme angles that the directions of perceived motion could be taken as known, even though no responses were given. Thus, a response trial could be preceded by any of eight possible sequences of three no-response trials, according to whether each induced clockwise or counterclockwise motion. Each subject’s threshold for equal responding between c and C was estimated following all eight trial histories. The results are shown in Figure 4B (solid line).

Under the separate-stages hypothesis, this experiment procedure should eliminate first-degree REs, because there are no recent responses to generate a base-rate expectancy. Therefore, we fit the second-degree model (i.e., Equation 12 with $\beta_{base} = 0$) to the data from this experiment. Figure 4B (dashed line) shows the fit, which explains 97.33% of the variance across the 8 trial histories. Fitting the full model produces the same fit, with an optimal value of β_{base} equal to zero. Fitting the first-degree model (Equation 12 with $\beta_{rep} = 0$) produces a significantly worse fit, with degenerate parameters that partially mimic the second-degree model (see Footnote 3). Therefore, the sequential effects in this experiment appear to be best explained by second-degree REs alone (together with an alternation bias).

Our final line of prior evidence for the separate-stages hypothesis comes from Experiment 2 of Wilder et al. (2013). Complementary to Maloney et al.’s (2005) elimination of recent responses, this experiment effectively eliminated stimuli. On each trial, the subject grasped the handle of a large robotic arm and moved it 15 cm forward in the axial plane, while the robot applied a perpendicular perturbing force randomly to either the left or the right. Critically, this force was uncued; the subject experienced it only by its direct influence on the movement. The subject’s task was to move directly forward, compensating as much as possible for the perturbation. On each trial, Wilder et al. measured the maximal lateral deviation (in centimeters) of the subject’s trajectory in the direction of the force. For example, on trials with a rightward force, this measure equaled the maximum distance that the movement deviated to the right from a straight midline path. This deviation measure was then used to investigate sequential effects. For example, on trials with a rightward force, the trajectory tended to deviate much more to the right if the past several trials had leftward forces than if they had rightward forces.

The focus of Wilder et al.’s (2013) analysis was how long-lasting sequential effects in this task were, but here we are concerned with the relative contributions of first- and second-degree effects. We assumed that the absence of overt stimuli would eliminate stimulus-based sequential effects, which the separate-

stages hypothesis predicts to manifest as an absence of second-degree REs. To test this prediction, we fit the full parallel-learning model as well as the first- and second-degree submodels. The alternation bias parameter (δ) was held to zero for all three models, in line with the conjecture (supported by the new experiments reported below) that the bias resides in stimulus processing. To fit the models, we calculated the mean lateral deviation for each 4-deep history for each of the 20 subjects. We then averaged over subjects and fit the model to the means for the 16 histories by least squares.

As Table 1 shows, the optimal value of β_{rep} in the full model is zero, and thus the full model reduces to the first-degree model. This model explains 97.97% of the variance across the 16 trial histories, whereas the second-degree model explains only 21.24%. As a further test for the presence of second-degree effects, we fit the full model to the data of each individual subject, allowing the β parameters to be negative. The value of β_{base} was positive for every subject (mean and median both 1.28), but β_{rep} was negative for 11 of the 20 subjects, with a median of $-.03$. Wilcoxon's signed-rank test showed no reliable difference from zero, $p = .36$ (a nonparametric test was needed because the model found degenerate solutions for 3 subjects, with ϵ_{rep} near zero and β_{rep} very large). Figure 5 shows the fit of the first-degree model, with the trial histories ordered by their second-degree representations (as in previous figures) and with the histories reordered according to their first-degree representations. The latter ordering clearly shows the purity of the first-degree pattern.

In conclusion, the reanalyses of data from Jentsch and Sommer (2002), Maloney et al. (2005), and Wilder et al. (2013) support the separate-stages hypothesis for first- and second-degree REs. The dissociation is seen both neurophysiologically, in Jentsch and Sommer's ERP data, and behaviorally, through Wilder et al.'s and Maloney et al.'s elimination of stimuli and of recent responses from their respective tasks. The fact that these data come from measures quite different from behavioral RT (viz., ERP latencies, motion-perception thresholds, and deviations in motor trajectories) speaks to the robustness of this dissociation. Thus, there is strong evidence for psychologically distinct learning mechanisms underlying sequential effects in binary tasks, one learning the base rate and the other the repetition rate. Furthermore, the nature of the dissociation leads to a more specific proposal, namely that sequential effects arise from incremental learning of the response base rate and the stimulus repetition rate. We turn now to new experiments testing this hypothesis.

Experiment 1

Experiment 1 was designed to test the theoretical proposals introduced above: (1) that sequential effects in 2AFC can be explained as a combination of first- and second-degree REs; (2) that these effects arise from incremental learning of the base rate and repetition rate; and (3) that first- and second-degree REs, respectively, occur in response and stimulus processing. Subjects performed a speeded visual discrimination task, in which on each trial a dot appeared above or below a horizontal fixation line. Measures of RT and ERP were both collected, to enable analysis of sequential effects in each.

The experimental methods closely followed those of Jentsch and Sommer (2002, Experiment 1), with one exception: Each

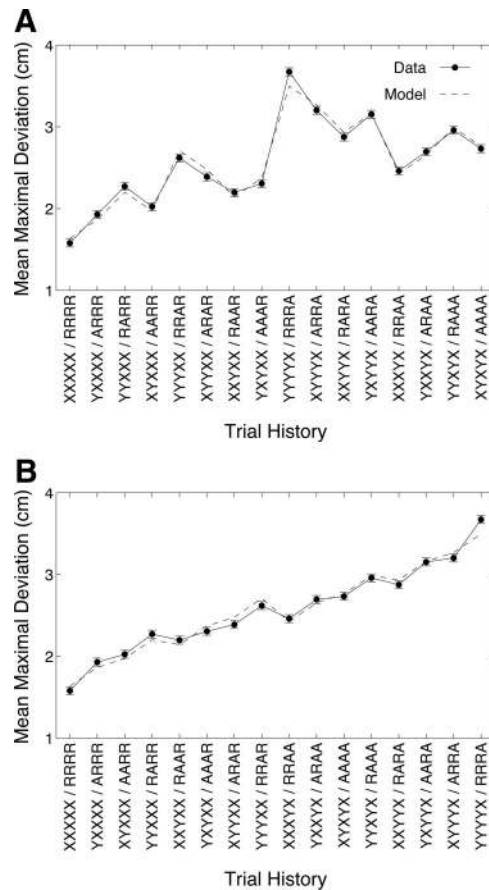


Figure 5. Fits of parallel-learning model to the reaching task of Wilder et al. (2013, Experiment 2). Sequential effects in this task were induced by a random force on each trial perpendicular to the subject's movement; thus, there were no overt stimuli. As predicted by the separate-stages hypothesis, the model fits the data best with no second-degree sequential effects (i.e., $\beta_{\text{rep}} = 0$). A: Mean maximal deviation of the subject's movement trajectory in the direction of the current force, as a function of trial history. Histories are ordered according to their second-degree representations, as in Figures 1–3. B: The same data and model fit, with the histories reordered according to their first-degree representations, to highlight the purity of the first-degree sequential effects. The curve shows perfect monotonicity except in every fourth gap, where multiple past first-degree mismatches (Ys) are replaced by a single more-recent mismatch (a pattern the model captures). Error bars in both graphs indicate within-subject standard errors (Masson & Loftus, 2003).

subject performed two sessions, under different conditions of autocorrelation in the stimulus sequence. In the *positive* condition, the repetition rate was $\frac{2}{3}$, whereas in the *negative* condition it was $\frac{1}{3}$. The sequences had no multi-step dependencies (i.e., they were Markov). The primary purpose of this manipulation was to test how long-term learning of the repetition rate relates to second-degree REs. To foreshadow, subjects showed faster responding to repetitions than to alternations in the positive condition, and the reverse pattern in the negative condition. Thus, they appear to have learned the differential repetition rates in the two sessions. Under the hypothesis that sequential effects and long-term learning both arise from the same incremental mechanism, a single model should be able to explain variations in RT across trial histories as well as between conditions.

ERP analyses focused on LRP and on the P100 component in early visual processing, with the goal of isolating sequential effects in stimulus and response processing. LRP was used to decompose total RT into stimulus- and response-processing times (Coles, 1989; de Jong et al., 1988). Importantly, this approach does not commit to a strictly serial model of stimulus and response processing; response planning could begin while stimulus identification is still incomplete. The only necessary assumptions are that variation in the time from stimulus onset to LRP onset (sLRP) reflects variation in stimulus processing, whereas variation in the time from LRP onset to the response (LRPr) reflects variation in response processing. Therefore, sequential effects in sLRP can be taken to reflect the contribution of stimulus expectancies, and sequential effects in LRPr can be taken to reflect the contribution of response expectancies. The logic behind the P100 analysis is similar: The P100 component is believed to reflect stimulus processing, with a greater amplitude when stimuli are more expected (Mangun & Hillyard, 1991; Spehlmann, 1965). Therefore, sequential effects in P100 amplitude can be taken to reflect stimulus expectancies.

An additional goal of the experiment was to test between the direct and adaptation hypotheses regarding changes in sequential effects from nonstationarity in the task environment. As a reminder, the direct hypothesis states that first-degree REs are unaffected by nonstationarity; sensitivity to autocorrelation arises only because of direct learning of the repetition rate. Thus, the direct hypothesis predicts the only difference between autocorrelation conditions to be in subjects' asymptotic estimates of the repetition rate, which should speed repetition RTs in the positive condition and alternation RTs in the negative condition. The adaptation hypothesis holds that first-degree REs become stronger with positive autocorrelation in the task environment. Thus, it predicts that the contribution of first-degree learning will be stronger in the positive condition. In terms of the parallel-learning model, this adaptation could arise from increased values of β_{base} (i.e., stronger weighting of first-degree expectancies) or from increased values of ϵ_{base} (leading to faster adjustment of base-rate estimates). Experiment 1 affords two means of testing the adaptation hypothesis. First, first-degree REs can be assessed at lags longer than 1, where the estimated repetition rate should have no effect, and compared across conditions. Second, if the ERP analyses supported the separate-stages hypothesis, so that LRPr could be taken as a pure indicator of first-degree learning, then the adaptation hypothesis could also be tested by comparing sequential effects in LRPr across conditions.⁴

Method

Subjects. Twenty-eight undergraduate students participated for monetary compensation.

Design. Each subject performed two sessions, one each in the positive and negative conditions. Sessions were spaced by 2–7 days, and order was counterbalanced between subjects.

Each session comprised 3,744 trials divided into 33 blocks. The first block contained 22 trials and was excluded from all analyses as practice. The remaining 32 blocks contained 116 or 117 trials each. The first 10 trials of each of these blocks were also excluded from analyses, to reduce transient effects following breaks. Thus,

3,402 trials—the *test trials*—were included in the analysis for each session.

Local stimulus histories for test trials were controlled to a depth of six trials in both conditions, as follows. In both conditions, the base rate of the two stimulus locations was exactly 50%. In the positive condition, the repetition rate in the test trials was exactly $\frac{2}{3}$. Each test trial was classified according to the stimulus sequence on the current and previous five trials, yielding 64 history types. The frequency of each history type was exactly as dictated by the overall $\frac{2}{3}$ repetition rate. For example, on $(\frac{2}{3})^5$ of the 3,402 test trials, or 448 trials, the current and previous five stimuli were all identical. These trials were evenly divided between 111111 and 222222 sequences (where 1 and 2 represent the two stimulus locations), reflecting the 50% base rate. On $(\frac{1}{3})^5$ of the test trials, or 14 trials, the current and previous trials followed a perfect alternating pattern. These trials were divided evenly between 121212 and 212121 sequences. Stimulus sequences in the negative condition were defined analogously, using a repetition rate of $\frac{1}{3}$. The sequence of test trials for each subject in each condition was generated pseudo-randomly according to the preceding constraints. The excluded trials were generated pseudo-randomly with the constraint that they conformed to the appropriate repetition rate within each block (i.e., for trials 2–22 in Block 1 and trials 2–10 in the other blocks).

Procedure. The subject's task on each trial was to respond to the location of a white dot, 5 mm in diameter, presented 11 mm above or 12 mm below a 4-mm horizontal white fixation line that was in the center of a computer monitor. Target locations were arranged vertically rather than horizontally to prevent visual laterality effects that might interfere with measurement of ERP components (described below). Responses were made using a button box, which was oriented vertically so as to be spatially compatible with the target locations. The left and right index fingers were assigned to the two buttons, with the assignment counterbalanced across subjects and fixed between sessions for each subject. Stimulus duration was 60 ms. Reaction time was recorded at a resolution of 1 ms. If the subject did not respond within 1,000 ms of stimulus onset, the trial was terminated and a blank response was recorded. A 700-ms response-stimulus interval (RSI) followed each response (or the end of the 1,000-ms response window, when no response was given). This RSI duration was chosen because previous research has shown it to produce expectancy-based sequential effects (see Figure 7 and the General Discussion section). The fixation line was visible throughout the task.

Electrophysiological recordings. Scalp voltages were collected with a 128-channel HydroCel Geodesic Sensor Net connected to AC-coupled, 128-channel, high-input impedance amplifiers (200 M Ω , Net Amps; Electrical Geodesics Inc., Eugene, OR). Amplified analog voltages (0.1–100 Hz bandpass) were digitized at 250 Hz. Individual sensors were adjusted until impedances were less than 50 k Ω . Electroencephalography (EEG) was recorded with respect to a vertex reference (Cz). The EEG was digitally low-pass filtered at 40 Hz prior to ERP analysis. Trials were included in the

⁴ A preliminary analysis of the behavioral data from this experiment is reported in Wilder et al. (2013). The focus there was restricted to second-degree REs and how their magnitude decreases over long lags.

ERP analyses if the current trial and previous 3 trials were all correct with RT between 100 and 1,000 ms. Trials were discarded from analysis if more than 20% of the channels were bad (average amplitude over 100 μV or voltage fluctuations of greater than 50 μV between adjacent samples). These criteria led to exclusion of 31.8% of all trials in peristimulus analyses (i.e., for sLRP and P100) and 31.4% of all trials in periresponse analyses (i.e., for LRPr). Individual bad channels were replaced on a trial-by-trial basis with a spherical spline algorithm (Srinivasan, Nunez, Silberstein, Tucker, & Cadusch, 1996). Eye movements were corrected using an ocular artifact detection algorithm (Gratton, Coles, & Donchin, 1983).

For the LRP analyses, the ERPs were re-referenced to the average of the two mastoid channels, as is standard for LRP analyses. LRPs were measured using the voltage difference between EEG channels 36 (C3) and 104 (C4), which are located over the hand areas of the left and right motor cortex, respectively. The LRP on each trial was defined in the standard way, by subtracting the voltage in the channel ipsilateral to the correct response from the voltage in the contralateral channel (Coles, 1989). To assess sequential effects on the LRP onset, trials were classified into eight 3-deep trial histories (4-deep histories, as were used with RT data, provided insufficient trial counts for the ERP analyses). For each subject, two LRP waveforms were constructed for each history in each condition, one by temporally aligning trials on the stimulus (peristimulus) and the other by aligning on the response (periresponse). All peristimulus waveforms were defined from 200 ms before stimulus onset to 800 ms post-onset and were baseline-corrected using average LRP activity from 0 to 100 ms after onset. All periresponse waveforms were defined from 600 ms before the response to 400 ms after and were baseline-corrected using average LRP activity from 300 to 200 ms before the response. Thus, 896 total waveforms were produced, corresponding to 28 subjects, crossed with 2 conditions, crossed with 8 trial histories, crossed with peristimulus versus periresponse.

The onset of the LRP spike was estimated using three different criteria in both peristimulus and periresponse analyses. For periresponse waveforms (used for estimating LRPr), LRP onset was estimated as the time at which LRP dropped below $-0.8 \mu\text{V}$, $-1 \mu\text{V}$, or 50% of the minimum for each waveform (i.e., 50% of the peak negative voltage). These criteria were taken from Jentsch and Sommer (2002). For peristimulus waveforms (used for estimating sLRP), LRP onset was estimated as the first time point at which the LRP dropped below $-0.4 \mu\text{V}$, $-0.5 \mu\text{V}$, or 50% of the minimum. The milder absolute criteria were necessary in the peristimulus analysis because LRP waveforms were broader with less extreme peaks than in the periresponse analysis. Initial analyses reported below (see the LRP section) tested each criterion separately and then chose the criteria for sLRP and for LRPr that provided the most statistically reliable pattern of sequential effects for purposes of modeling. For all criteria in both peristimulus and periresponse analyses, linear interpolation was used between the first time point satisfying the criterion and the preceding time point, to obtain a continuous-valued estimate of the moment at which the threshold was crossed (time points were separated by 4 ms, due to the 250-Hz digitization of all ERP data).

For each history and condition, a pair of aggregate LRP waveforms (periresponse and peristimulus) was created by averaging the waveforms of individual subjects. The aggregate waveforms

were used for estimating LRP onset, and hence sLRP and LRPr, averaged over all subjects. In addition, for each aggregate waveform, a set of 28 jackknife waveforms was created by omitting each of the 28 subjects and averaging over the other 27. LRP onsets from the jackknife waveforms were used in statistical tests reported below, following Miller, Patterson, and Ulrich (1998; see also Miller, Ulrich, & Schwarz, 2009; Ulrich & Miller, 2001). The logic of this procedure is that the residual variance of the jackknife-based onsets is approximately equal to the residual variance of onsets from individual subjects, divided by $(n - 1)^2$, where n is the number of subjects. Because LRP onsets cannot be reliably estimated from individual subjects (the data are too noisy and noise has a nonlinear effect on onset estimates), one can perform statistics on jackknife-based onsets and then correct for the change in residual variance. The correction amounts to multiplying F statistics by $(n - 1)^2$ or t statistics by $n - 1$ before computing p -values. All inferential statistics reported below (see the LRP section) are values after this correction has been made.

For the P100 analyses, the EEG was re-referenced to an average reference, the voltage difference between that channel and the average of all channels, to minimize the effects of reference site activity and to improve estimates of electrical field topography (Dien, 1998). The average reference was corrected for the polar average reference effect (Junghöfer, Elbert, Tucker, & Braun, 1999). The ERPs were baseline corrected with respect to the 200-ms pre-stimulus interval. P100 latency and amplitude were measured over two bilateral electrode clusters, selected by identifying the bilateral electrodes where the P100 was maximal across all subjects and conditions (left channel number 66, between standard 10–10 locations P1 and O1; right channel number 84, between locations P2 and O2). To account for topographic variability, analyses were conducted on the mean ERPs across electrode clusters including each of those maximal locations and the 6 immediately surrounding electrodes within each hemisphere. Based on visual inspection of the ERPs within those clusters averaged across all subjects and conditions, a temporal analysis window of 96–156 ms after stimulus onset was selected. Voltage amplitudes were computed as the mean amplitude within that window. To assess sequential effects in P100, trials were classified according to 3-deep trial history as in the LRP analysis, and the mean P100 amplitude was estimated for each subject within each history and condition. P100 amplitudes were estimated for individual subjects rather than from jackknife waveforms and were analyzed using standard statistical tests (in contrast to LRP onsets).

Results: Reaction Time

Trials were included in RT analyses if the subject responded within the 1,000-ms deadline, the RT was at least 100 ms, and the current and previous four responses were all correct. These criteria resulted in exclusion of an average of 32.5% of the 3402 test trials in each session, primarily due to the correct-response requirement (only 1.3% of responses fell outside 100–1,000 ms). The rationale for this stringent requirement is that we are interested in sequential effects arising from both stimulus and response processes, and hence stimuli and responses needed to be jointly controlled. An alternative analysis that required only the current response and one previous response to be correct, which excluded 15.4% of the test trials, yielded the same qualitative results, with only minor differ-

ences for sequential effects at higher lags. The button box failed for the 20th block of the negative condition for one subject, so those trials were excluded as well.

Trials were classified according to the history of the current and previous five stimuli, as follows. First, 64 groups of trials were defined according to the sequence of physical stimulus locations (top or bottom) on trials $n - 5$ through n . Second, these groups were combined into 32 groups, by collapsing over the physical identity of stimulus n . These 32 groups thus correspond to the possible 5-deep histories. Each such history is specified by whether trials $n - 5$ through $n - 1$ are first-degree matches or mismatches with trial n (e.g., XYXXXX, where the final X indicates trial n), or equivalently by whether trials $n - 4$ through n are repetitions or alternations (e.g., AARRR).

Mean RT was calculated for each 5-deep history for each subject in both conditions. These mean RTs were then averaged in pairs (collapsing over trial $n - 5$ in the first-degree labeling, or equivalently collapsing over trial $n - 4$ in the second-degree labeling) to obtain estimated mean RTs for all sixteen 4-deep histories. This two-step procedure served to control for the first-degree correlation between trial $n - 4$ and all earlier trials, which differed between conditions. In other words, this procedure ensured that the RT estimate for each history reflected an equal balance between trials on which stimuli $n - k$ and n were first-degree matches and trials on which they were mismatches, for all $k > 4$. This balance also facilitates interpretation of the parallel-learning model (see Appendix A).

Figure 6 shows the mean RTs for each history and condition, averaged over subjects. A 2×16 (Condition \times History) repeated-measures analysis of variance (ANOVA) shows a strong main effect of history, $F(1.75, 42.29) = 34.44, p < .0001$, using Greenhouse–Geisser (GG) sphericity correction with $\epsilon = .117$. This effect follows a similar pattern to that seen in the RT data of




Figure 6. Mean response time (RT) by trial history for Experiment 1. The 4-deep histories are labeled by both their first- and second-degree representations, as in previous figures. Error bars correspond to the within-subjects standard error taken across all 32 observations, following Masson and Loftus (2003), to facilitate comparisons across histories and conditions. Positive = positive condition; Negative = negative condition.

Cho et al. (2002) and Jentzsch and Sommer (2002) (see Figure 1). In addition, there is a strong history-condition interaction, $F(5.10, 137.70) = 102.58, p < .0001$, GG $\epsilon = .340$, such that RTs on repetition trials are faster in the positive condition, and RTs on alternation trials are faster in the negative condition. The main effect of condition is nonsignificant, $F(1, 27) < 1$.

A core assumption of our modeling approach is that variation in mean RT across histories and conditions is due to varying expectancies prior to each trial (see Equation 10). This assumption leads to a “cost-benefit” prediction, whereby faster responding to one stimulus is always perfectly balanced by slower responding to the other stimulus (Audley, 1973; Laming, 1968). For example, Figure 6 shows the fastest mean RT occurs in the negative condition with history XYXYX/AAAA. The assumption is that this fast responding is due to strong expectancy of X/A following XYXY/AAA, which implies expectancy of Y/R must be very low. This reasoning leads to the correct prediction that mean RT will be especially long for history XYXY/AAAR (which by our convention is rewritten YXYXX/AAAR) in the negative condition.

A standard way to test the cost-benefit prediction is with an exchange plot, a scatterplot of mean RT for both possible stimuli following each history (Audley, 1973). To create such a plot, the trial histories are combined into pairs that differ only in the value of the current trial (e.g., pairing XYXYX/AAAA and YXYXX/AAAR). This pairing is the same under first- or second-degree labeling of trial histories. For each pair, mean RT for the repetition case (e.g., YXYXX/AAAR) is plotted against mean RT for the alternation case (e.g., XYXYX/AAAA). The values being plotted are the same as in a standard plot of RT by trial history (e.g., Figures 1 and 6), but they are arranged in a way that makes the implications of the expectancy assumption more apparent. If the assumption is correct, then the cost-benefit pattern should manifest as all points lying on a common negative diagonal. This prediction can be contrasted with a “pure facilitation” pattern that is often observed in tasks with very short RSIs, whereby RT is faster or slower for both members of a history pair, leading to points lying on a positive diagonal (Soetens et al., 1985).

Figure 7 shows the exchange plot for Experiment 1. Across all 16 history pairs (8 for each condition), the correlation between mean RT on repetition and alternation trials is $-.986$. The plot thus shows a strong cost-benefit pattern, supporting the expectancy assumption. Also evident is the effect of the autocorrelation manipulation, in that RT is faster for repetitions in the positive condition and faster for alternations in the negative condition. The fact that the data from the two conditions differ in this way but still lie on the same diagonal shows that long-term learning of the repetition rate is mediated by expectancy just as sequential effects are. This conclusion is consistent with the assumption that sequential effects and long-term learning arise from the same mechanism.

In the next two subsections, we show that the parallel-learning model fits the data from Experiment 1 well overall, but it cannot account for two key phenomena. These phenomena lead to a new model—the joint-learning model—that learns the same two sequence statistics as the parallel-learning model but with a more sophisticated learning mechanism. This new model is then applied to the ERP data to further evaluate the separate-stages hypothesis.

This document is copyrighted by the American Psychological Association or one of its allied publishers. This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

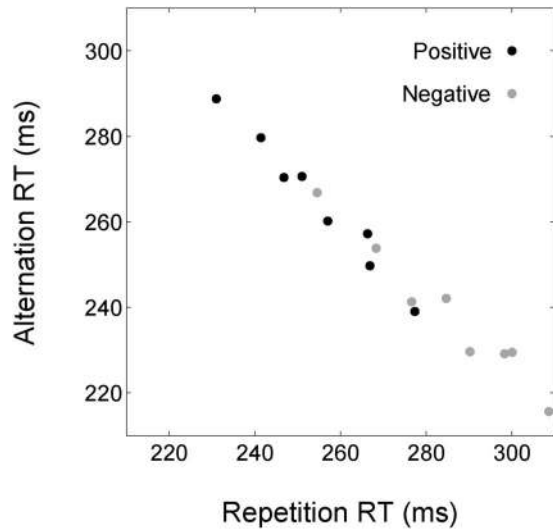


Figure 7. Exchange plot for Experiment 1. Each point represents a pair of trial histories differing only in the value of the current trial, one a repetition and the other an alternation. The values plotted are the same as in Figure 6. RT = response time; Positive = positive condition; Negative = negative condition.

Parallel-Learning Model Fits

The parallel-learning model (Equation 10) was fit separately for each subject by least squares to the 32 mean RTs from all 4-deep histories in both conditions. Model predictions were derived from Equation 10, as described in detail in Appendix A. Model parameters for each subject were held constant between conditions. Figure 8 shows the model fits, averaged over subjects and compared to the mean RTs from Figure 6. Table 2 shows the median parameter values across subjects. To evaluate the fit of the model, it was fit to the aggregate data (i.e., averaged over subjects) using a single set of parameters. This aggregate model explains 96.67% of the variance in mean RT across the 32 trial types (16 histories in 2 conditions), on 6 free parameters.

One important contribution of the parallel-learning model is that it explains sequential effects and long-term learning by the same mechanism. Specifically, the model assumes that incremental learning of the repetition rate (Equation 5) is responsible both for second-degree REs and for differential learning of the repetition rate in each condition. The latter effect is seen in the fact that RT for every history ending in a repetition is faster in the positive condition (left half of Figure 8), and RT for every history ending in an alternation is faster in the negative condition (right half of Figure 8). According to the model, this long-term learning effect is due to the accumulated contributions of trials $n - 5$ and earlier (i.e., those not controlled in our partitioning of trial histories). As shown in Appendix A (see Equation A11), the magnitude of this predicted long-term learning effect can be directly calculated as $\frac{4}{3}\beta_{\text{rep}}(1 - \epsilon_{\text{rep}})^4$. Because β_{rep} and ϵ_{rep} are the same parameters responsible for the second-degree RE, there is a strong constraint linking the model's predictions for second-degree REs and for the size of the long-term learning effect. The fact that the model simultaneously provides good fits of both phenomena supports the link between REs and incremental learning.

Lag Profiles

A more fine-grained picture of the parallel-learning model's fit to the RT data comes from assessing the first- and second-degree REs separately at each lag. We quantify REs at each lag, k , as the reduction in RT due to a match between trials n and $n - k$. This is done for both first- and second-degree matches. Once REs are calculated at each lag, one can analyze the first- and second-degree lag profiles, which show the strengths of the first- and second-degree REs as a function of lag.

The first-degree lag profile, $RE_1(k)$, represents the RT facilitation (if $RE_1 > 0$) or slowing (if $RE_1 < 0$) due to a first-degree match between trials n and $n - k$. To calculate this effect, we separated the sixteen 4-deep histories into the eight for which trials n and $n - k$ mismatch and the eight for which they match. Mean RTs were averaged over both groups of histories, and then the difference was taken. Averaging the means for the histories, rather than directly averaging individual trials, leads to all histories being weighted equally, which simplifies the model-based analysis.

The second-degree lag profile, $RE_2(k)$, was computed analogously to the first. The histories were grouped according to whether there is a second-degree match between trials n and $n - k$. Mean RTs were averaged across groups of histories and the difference taken between groups.

When starting from mean RTs based on 4-deep histories, $RE_1(k)$ can be defined up to a lag of $k = 4$, whereas $RE_2(k)$ can be defined up to $k = 3$. $RE_1(2)$ and $RE_2(1)$ are logically equal and hence unidentifiable, because trials n and $n - 2$ are a first-degree match if and only if trials n and $n - 1$ are a second-degree match (i.e., both repetitions or both alternations). Thus, the lag profiles provide six independent measures of sequential effects. Each measure corresponds to a balanced contrast among the mean RTs for the sixteen 4-deep histories, and it is easily verified that all six contrasts are orthogonal.

The lag profiles are especially useful in the context of the parallel-learning model, because they essentially invert that mod-

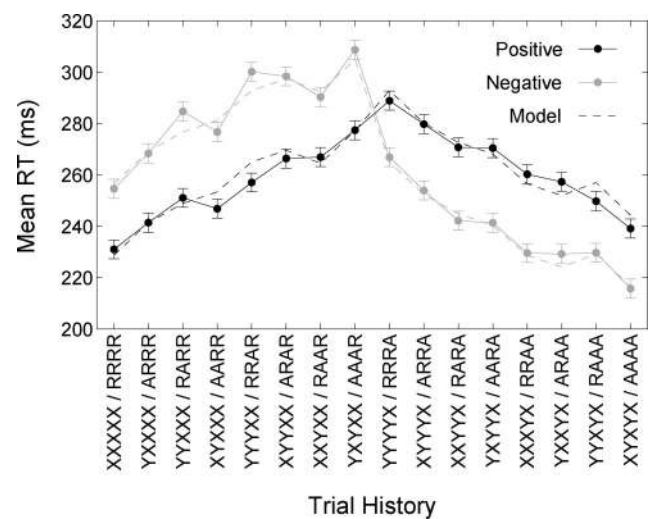


Figure 8. Fit of the parallel-learning model to response time (RT) data of Experiment 1. Error bars correspond to the within-subjects standard error taken across all 32 observations, following Masson and Loftus (2003). Positive = positive condition; Negative = negative condition.

Table 2
Fits of Models to Response Time Data of Experiment 1

Model	Parameter						
	β_0	β_{base}	ϵ_{base}	β_{rep}	ϵ_{rep}	δ	Var
Parallel learning	260.19	67.45	.551	62.13	.091	39.13	.9667
Joint learning	260.72	24.85	.190	80.46	.099	—	.9713

Note. Parameter values are medians across fits to individual subjects. Var column indicates proportion of variance explained across all 32 of the 4-deep trial histories (16 in each condition), based on fit to group data. Dash indicates parameter held to zero.

el’s predictions, isolating the effects of both learning components at all lags. The model’s predictions for both lag profiles can be exactly solved at all lags in terms of the model parameters, as derived in Appendix B. The results are shown in Table 3. The predicted REs at higher lags ($RE_1(k)$ for $k \geq 3$ and $RE_2(k)$ for $k \geq 2$) are pure measures of the respective learning components, decaying exponentially as a function of k with magnitude $2\beta\epsilon$ and decay rate $1 - \epsilon$. $RE_1(2)$ and $RE_2(1)$ are not especially diagnostic because they are equal and depend on both learning components, as explained above. Finally, $RE_1(1)$ corresponds to (the negative of) the alternation advantage. It reflects both a first-degree RE and long-term second-degree learning, the latter being the only way in which the model predicts the two conditions to differ (note the opposite signs in rows 1–2 of Table 3). The artificial bias term is also reflected in $RE_1(1)$, shifting it negatively (when $\delta > 0$) to capture the net alternation advantage between the two conditions.

Figure 9 (dashed lines) shows the joint-learning model’s predictions for the second-degree lag profile, $RE_2(k)$. These predictions are derived from the same fits used in Figure 8 and are averaged across the predictions for individual subjects. The model predicts a second-degree lag profile that decays with increasing lag and that is identical for the two conditions. Solid lines show the empirical second-degree lag profile, computed for individual subjects and then averaged over subjects. The data are consistent with the model’s prediction of no difference between conditions, as supported by a paired-samples t -test at each of the three lags, as well as on the average across all three lags ($p > .2$ for each of the four comparisons). We regard the second-degree lag profile to be consistent with the predictions of the parallel-learning model, as well as of the joint-learning model considered below (fits not shown), and we do not consider it further.

Table 3
Predictions of Parallel-Learning Model for Lag Profiles

Measure	Value
$RE_1(1)$, Positive condition	$2\beta_{base} \epsilon_{base} + \frac{2}{3}\beta_{rep}(1 - \epsilon_{rep})^4 - 2\delta$
$RE_1(1)$, Negative condition	$2\beta_{base} \epsilon_{base} - \frac{2}{3}\beta_{rep}(1 - \epsilon_{rep})^4 - 2\delta$
$RE_1(2)$	$2\beta_{base} \epsilon_{base} (1 - \epsilon_{base}) + 2\beta_{rep}\epsilon_{rep}$
$RE_1(3)$	$2\beta_{base} \epsilon_{base} (1 - \epsilon_{base})^2$
$RE_1(4)$	$2\beta_{base} \epsilon_{base} (1 - \epsilon_{base})^3$
$RE_2(1)$	$2\beta_{base} \epsilon_{base} (1 - \epsilon_{base}) + 2\beta_{rep}\epsilon_{rep}$
$RE_2(2)$	$2\beta_{rep} \epsilon_{rep} (1 - \epsilon_{rep})$
$RE_2(3)$	$2\beta_{rep} \epsilon_{rep} (1 - \epsilon_{rep})^2$

Note. All measures but $RE_1(1)$ are predicted to be the same for the two conditions.

The first-degree lag profile turns out to be more theoretically informative. Figure 10A shows the parallel-learning model’s prediction for the first-degree lag profile, based on averaging the predictions from the fits of individual subjects. The model predicts the only difference between conditions to lie in $RE_1(1)$, due to long-term learning of the repetition rate (see also Table 3). These predictions correspond to the direct hypothesis regarding the influence of autocorrelation on sequential effects. Recall that this hypothesis states that differences in first-degree REs due to autocorrelation come only from direct contributions of second-degree learning (i.e., learning the repetition rate), rather than from adaptation of first-degree learning. The model as presented in Figure 10A embodies this hypothesis, because there is no change in model parameters (in particular, β_{base} or ϵ_{base}) between conditions.

The alternative possibility considered here is the adaptation hypothesis, which posits that first-degree learning adapts to environmental statistics. For example, β_{base} could increase in the positive condition, because such an adaptation would contribute to more correct expectancies. This possibility is illustrated in Figure 10B. To generate representative predictions, each subject’s estimated β_{base} was scaled by an arbitrary constant factor (1.25) to

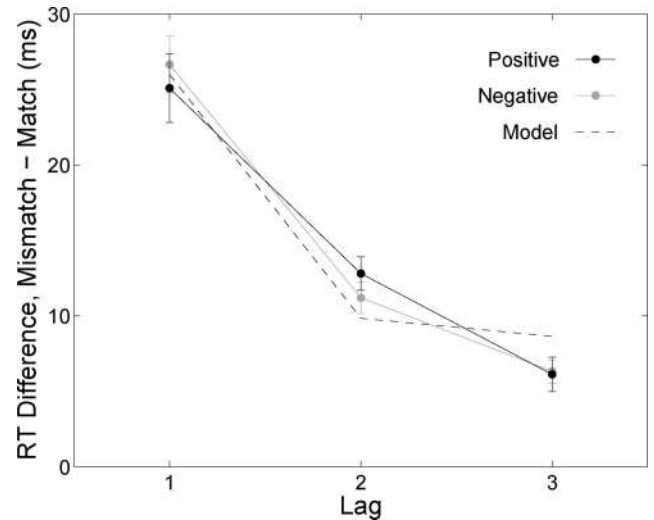


Figure 9. Second-degree lag profiles from Experiment 1, with predictions from the parallel-learning model. Model’s predictions are identical for the two conditions. Error bars indicate standard error between subjects. RT = response time; Positive = positive condition; Negative = negative condition.

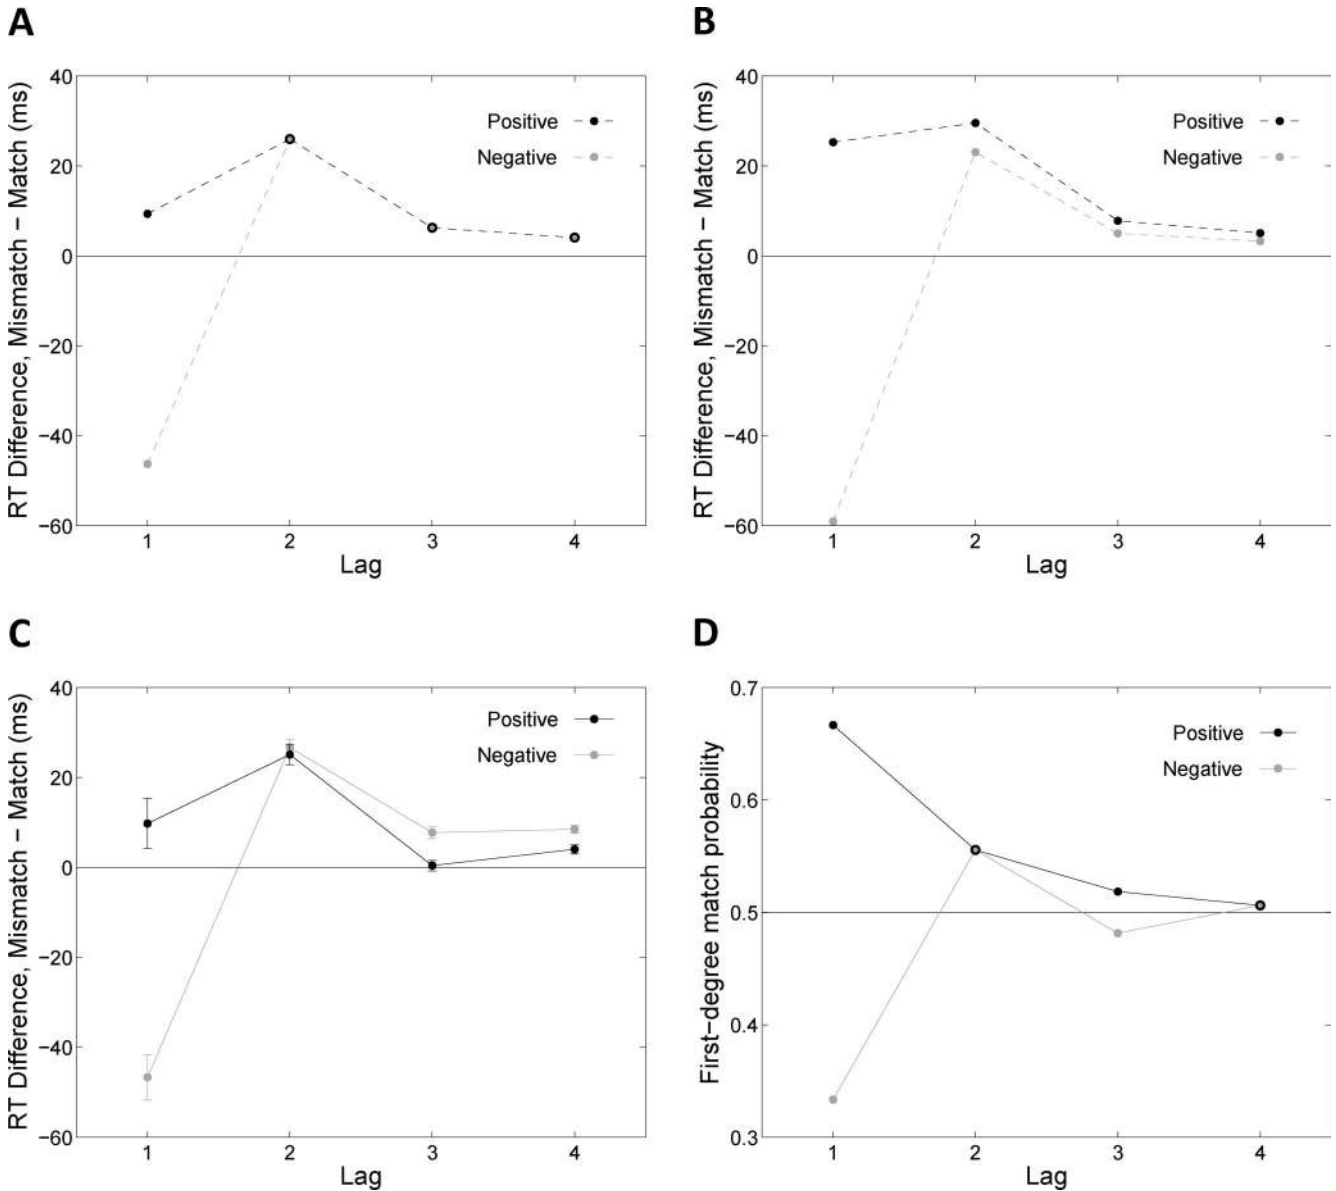


Figure 10. A: Predictions of the parallel-learning model for the first-degree lag profile in Experiment 1. These predictions correspond to the direct hypothesis, in that the only difference between conditions is due to direct learning of the repetition rate. Predictions for the two conditions are identical for lags greater than 1. B: Example predictions of the adaptation hypothesis, derived by rescaling the magnitude of the first-degree component of the model (β_{base}), upward in the positive condition and downward in the negative condition. C: Empirical first-degree lag profiles from Experiment 1. Error bars indicate standard error between subjects. D: Probability, as determined by the experiment design, of a first-degree match between the current trial and previous trials as a function of lag. Probabilities for even lags are the same for both conditions. RT = response time; Positive = positive condition; Negative = negative condition.

predict the positive condition, and the inverse scaling was applied in the negative condition. The result is an adaptation of RE_1 at all lags. A similar prediction can be obtained by adjusting ϵ_{base} between conditions (not shown).

Figure 10C shows the empirical first-order lag profiles, averaged over the profiles derived from individual subjects. Most relevant to the contrast between the direct and adaptation hypotheses are the differ-

ences in RE_1 at lags 2 and greater. Contrary to both hypotheses, $RE_1(k)$ for $k \geq 2$ is larger in the negative condition than in the positive condition. This difference is nonsignificant at lag 2 ($p > .5$), but it is highly reliable at lag 3 ($M_{\text{positive}} = 0.42$ ms, $M_{\text{negative}} = 7.73$ ms, paired $t(27) = -5.94$, $p < .0001$) and lag 4 ($M_{\text{positive}} = 4.04$ ms, $M_{\text{negative}} = 8.49$ ms, paired $t(27) = -4.21$, $p < .001$). We refer to this finding as a *reverse adaptation effect*, because it is opposite

the effect predicted by the adaptation hypothesis (whereas the direct hypothesis predicts no effect in either direction). The reverse adaptation effect is a counterintuitive result because it goes against the statistics built into the trial sequences. Specifically, the probability that the current trial will physically match the trial at lag k is always greater in the positive condition (for k odd) or else equal in both conditions (for k even), as shown in Figure 10D.

Figure 10C also highlights a second challenge to the parallel-learning model: the alternation advantage, which the model can only explain using the artificial bias term. The large difference between conditions at lag 1 is explained by the model as long-term second-degree learning, as discussed above. However, the data also show that the negative value in the negative condition is stronger than the positive value in the positive condition; that is, $RE_1(1)$ is reliably negative when averaged over the two conditions ($M = -18.45$ ms, $t(27) = -3.65$, $p < .01$). This effect represents the net alternation advantage: overall faster responding on alternation trials. The analytic predictions in Table 3 demonstrate that the parallel-learning model logically cannot explain the net alternation advantage without an explicit bias term (add the predictions in rows 1 and 2, with $\delta = 0$). The bias term was included in the model for generating the predictions in Figures 8 and 10A (it has no impact on Figure 9), but it is entirely ad hoc. It was added to enable presentation of the various aspects of the data that the model does explain, while postponing real consideration of the alternation advantage until this point in the article.

In summary, the lag profiles do not answer the question between the direct and adaptation hypotheses, but they do leave us with two critical phenomena that cannot be explained by the parallel-learning model: the reverse adaptation effect and the net alternation advantage. We turn now to a new model that suggests these findings are both signatures of the learning mechanism underlying sequential effects.

Joint-Learning Model

Both the reverse adaptation effect and the net alternation advantage have a potential explanation rooted in joint error correction. Joint error

correction is a well-established principle of learning theory, whereby learning is driven by the prediction error of all available cues combined (Rescorla & Wagner, 1972), as opposed to learning each cue-outcome contingency separately (Estes, 1950; Hebb, 1949). The process of joint error reduction leads to interactive dynamics among different association weights (here w_{base} and w_{rep}), in particular producing cue competition, whereby multiple cues can compete to predict the outcome (Gluck & Bower, 1988; Rescorla & Wagner, 1972). Cue competition is central to most theories of classic learning phenomena such as blocking (Kamin, 1968), in which little is learned about one cue that is redundant to another cue that has already been learned. Theories of joint error correction and cue competition have previously been applied only to cases of multiple cues present on a single trial (e.g., a tone and a light presented concurrently), but a similar mechanism may be in play with sequential effects, where the “cues” are aspects of the trial sequence.

Here, we develop a joint-learning analog to the parallel-learning model analyzed above and show how it can explain the reverse adaptation effect and net alternation advantage. We begin with a simple connectionist interpretation of the parallel-learning model, shown in Figure 11A. Under this interpretation, the parallel-learning model learns two separate predictions for the upcoming trial (E_n), one based on a constant cue for learning the base rate, and the other using the previous trial as a cue, for learning the repetition rate. The update rules of Equations 1 and 5 can be viewed as Hebbian learning, where each association weight (w_{base} and w_{rep}) is adjusted toward the product of its input and target output. Alternatively, they can be viewed as error-correction rules, in which each weight is adjusted by a proportion of the prediction error, multiplied by the value of the input (i.e., gradient descent). The latter interpretation is straightforward for w_{base} (see Equation 1), and it can be seen for w_{rep} by rewriting Equation 5 as

$$\Delta w_{rep} = \epsilon_{rep}(E_n - w_{rep}E_{n-1})E_{n-1}. \tag{13}$$

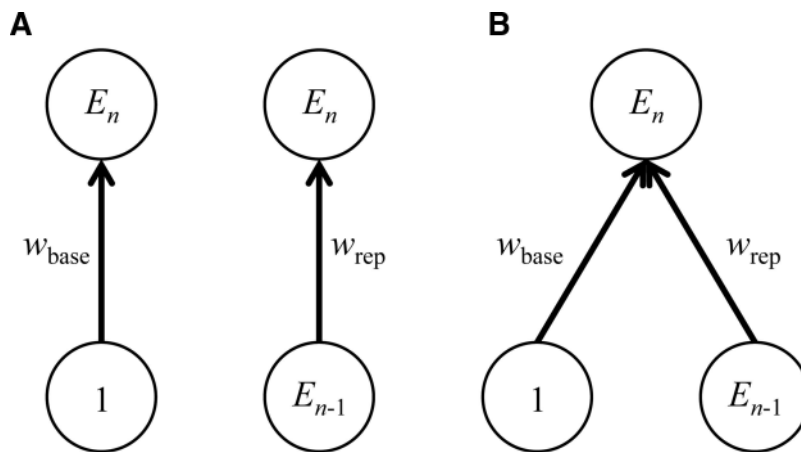


Figure 11. Illustration of the difference between parallel-learning and joint-learning models. A: In the parallel-learning model, the association weights w_{base} and w_{rep} generate separate predictions for each upcoming trial (E_n). A constant cue (1) is used for w_{base} , leading it to estimate the base rate. The previous trial (E_{n-1}) acts as a cue for w_{rep} , leading it to estimate the repetition rate. Each association weight is updated by Hebbian learning, or equivalently based on its own prediction error. B: In the joint-learning model, the two weights combine to generate a joint prediction, and each is updated based on the joint error.

Table 4
Comparison of Parallel- and Joint-Learning Models

Assumption	Parallel learning	Joint learning
Update rules	$\Delta w_{\text{base}} = \epsilon_{\text{base}} (E_n - w_{\text{base}})$ $\Delta w_{\text{rep}} = \epsilon_{\text{rep}} (E_{n-1} E_n - w_{\text{rep}})$	$\Delta w_{\text{base}} = \epsilon_{\text{base}} (E_n - w_{\text{base}} - E_{n-1} w_{\text{rep}})$ $\Delta w_{\text{rep}} = \epsilon_{\text{rep}} (E_{n-1} E_n - w_{\text{rep}} - E_{n-1} w_{\text{base}})$
RT	$RT = \beta_0 - \beta_{\text{base}} w_{\text{base}} E_n - \beta_{\text{rep}} w_{\text{rep}} E_{n-1} E_n - \delta E_{n-1} E_n$	$RT = \beta_0 - \beta_{\text{base}} w_{\text{base}} E_n - \beta_{\text{rep}} w_{\text{rep}} E_{n-1} E_n$
LRP–Response latency	$LRPr = \beta_{\text{LRPr}} - \beta_{\text{base}} w_{\text{base}} E_n$	$LRPr = \beta_{\text{LRPr}} - \beta_{\text{base}} w_{\text{base}} E_n$
Stimulus–LRP latency	$sLRP = \beta_{\text{sLRP}} - \beta_{\text{rep}} w_{\text{rep}} E_{n-1} E_n - \delta E_{n-1} E_n$	$sLRP = \beta_{\text{sLRP}} - \beta_{\text{rep}} w_{\text{rep}} E_{n-1} E_n$

Note. The models differ in two ways. The joint-learning model learns from joint prediction error, leading to cue-competition terms in its update rules. The parallel-learning model includes an alternation-bias term (δ) in predicting response time (RT) and stimulus processing time, whereas the joint-learning model does not. These differences between the models are underlined in the table. LRPr = latency from lateralized readiness potential (LRP) onset to response; sLRP = latency from stimulus onset to LRP onset.

The joint-learning model assumes that, instead of learning two separate predictions, w_{base} and w_{rep} combine to generate a single prediction,

$$\hat{E}_n = w_{\text{base}} + w_{\text{rep}} E_{n-1}, \quad (14)$$

as illustrated in Figure 11B. Updating of each weight is then based on the error of this joint prediction:

$$\begin{aligned} \Delta w_{\text{base}} &= \epsilon_{\text{base}} (E_n - \hat{E}_n) \\ &= \epsilon_{\text{base}} (E_n - w_{\text{base}} - E_{n-1} w_{\text{rep}}) \end{aligned} \quad (15)$$

and

$$\begin{aligned} \Delta w_{\text{rep}} &= \epsilon_{\text{rep}} (E_n - \hat{E}_n) E_{n-1} \\ &= \epsilon_{\text{rep}} (E_n - w_{\text{base}} - E_{n-1} w_{\text{rep}}) E_{n-1} \\ &= \epsilon_{\text{rep}} (E_{n-1} E_n - w_{\text{rep}} - E_{n-1} w_{\text{base}}). \end{aligned} \quad (16)$$

The third line in Equation 16 is meant to show the symmetry between the learning rules for w_{base} and w_{rep} , as well as their close similarity to the learning rules of the parallel-learning model (Equations 1 and 5). The only difference between the learning rules for the two models is the joint-learning model's inclusion of the cue competition terms, $-E_{n-1} w_{\text{rep}}$ and $-E_{n-1} w_{\text{base}}$, by which each weight affects updating of the other. As shown in the next two subsections, these two cue competition effects turn out to predict the reverse adaptation effect and net alternation advantage, respectively.

Table 4 summarizes the differences between the parallel- and joint-learning models. In addition to learning w_{base} and w_{rep} from joint prediction error, the joint-learning model abandons the artificial alternation-bias term for predicting RT. Because cue competition can explain the alternation advantage (as is shown below), there is no longer a need for this ad hoc assumption. Thus the joint-learning model's RT prediction reverts to Equation 7 instead of Equation 10. We also drop the alternation bias for predicting ERP components, in the ERP analyses reported below.

Cue competition explanation for the alternation advantage.

The joint-learning model explains the net alternation advantage as a consequence of how w_{base} modulates the learning dynamics of w_{rep} , via the cue competition term $-E_{n-1} w_{\text{base}}$ in Equation 16. This cue competition effect turns out to bias w_{rep} toward an expectation of alternation, as follows.

Incremental learning leads w_{base} to be biased toward the values of recent trials (relative to the true long-term base rate). This is the

basic explanation of the first-degree RE. In particular, at the start of trial n , the update following trial $n - 1$ has shifted w_{base} toward E_{n-1} . Thus the cue competition term, $-E_{n-1} w_{\text{base}}$, in the update equation for w_{rep} (Equation 16) will on average be negative. Therefore, cue competition exerts a negative bias on w_{rep} , shifting it toward expectancy of alternation.

More intuitively, the first-degree RE produced by incremental learning of w_{base} biases the model's expectancy toward repetition. Because learning is driven by joint error, w_{rep} will learn to compensate for this bias, by shifting its value in the direction of alternation. Thus, w_{rep} will approximately learn whatever the repetition rate happens to be, but it will always be biased away from the true repetition rate, in the direction of alternation, as a result of cue competition from w_{base} .

Whether the model as a whole produces a repetition advantage or an alternation advantage depends on the relative influences of first- and second-degree expectancies on RT (see Equation 7). It can be shown mathematically that an alternation advantage arises if and only if $\beta_{\text{rep}} > \beta_{\text{base}}$. The fits presented below of the joint-learning model confirm this is the case for Experiment 1. Experiment 2 tests the further prediction that selectively increasing β_{rep} leads to a stronger alternation advantage.

Cue-competition explanation of the reverse adaptation effect. The complementary cue competition mechanism, by which w_{rep} modulates the learning dynamics of w_{base} , logically leads to the reverse adaptation effect found in the first-degree lag profiles. The explanation for this conclusion is illustrated in Figure 12.

The starting point for this explanation is that, on average, w_{rep} will be greater in the positive condition than in the negative condition, due to long-term learning of the repetition rate in each condition. This difference in w_{rep} leads to a difference in the model's prediction for the upcoming trial, such that it is more in the direction of repetition in the positive condition and alternation in the negative condition. Imagine first that the current trial is a repetition (see Figure 12, top). Because repetition is more expected in the positive condition, the prediction error will tend to be smaller.⁵ Therefore learning will shift w_{base} toward E_n , but by a relatively small amount. In contrast, the prediction error in the negative condition will tend to be larger, leading to a bigger update, shifting w_{base} more strongly toward E_n . Thus, following learning on trial n , w_{base} will be closer to E_n and E_{n-1} (which were equal by assumption) in the negative condition than in the positive condition. Therefore any future instances of that same

⁵ The situation is different if the outcome is overpredicted, that is, $|\hat{E}_n| > 1$, but the conclusion is the same.

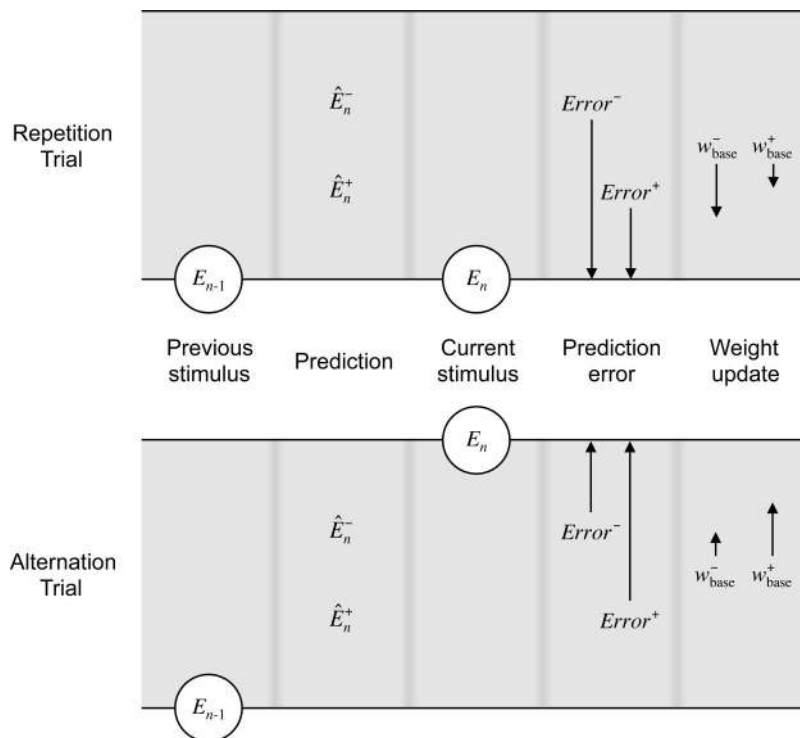


Figure 12. Illustration of cue-competition explanation for reverse adaptation effect. Upper panel shows a repetition trial, and lower panel shows an alternation trial. In both diagrams, time flows from left to right, horizontal lines correspond to the two possible stimulus or response values, and the space between them represents the range of possible predictions (i.e., strengths of expectancy between the two outcomes). The vertical positions of circles labeled E_{n-1} and E_n indicate the values of trials $n - 1$ and n . Superscript + and - indicate positive and negative conditions, respectively. The contribution of the repetition weight (w_{rep} , not shown) leads the current prediction, \hat{E}_n , to be further from E_{n-1} in the negative condition than in the positive condition. Regardless of whether trial n matches (upper panel) or mismatches (lower panel) trial $n - 1$, the prediction error is more in the direction of E_{n-1} in the negative compared to the positive condition. Thus, the update to the base-rate weight, w_{base} , is more in the direction of E_{n-1} in the negative condition. Consequently, there is greater facilitation in the negative condition when trials $n + 1$ and later match trial $n - 1$.

stimulus on the next several trials will produce faster responding in the negative condition. That is, the negative condition will show a stronger facilitation effect from subsequent first-degree matches to E_{n-1} .

Next consider the case of an alternation trial, $E_n \neq E_{n-1}$ (see Figure 12, bottom). In this case, the prediction error will tend to be smaller in the negative condition, and the update will tend to be greater in the positive condition. Therefore, w_{base} will shift more toward E_n —and further from E_{n-1} —in the positive condition than in the negative condition. Thus once again the negative condition will benefit more from future matches to E_{n-1} .

Therefore, regardless of whether trial n is a repetition or an alternation, w_{base} tends to end up closer to E_{n-1} in the negative condition than in the positive condition. The negative condition thus exhibits a greater benefit of matches to E_{n-1} on trials $n + 1$ and later. In other words, the first-order RE is stronger in the negative condition than in the positive condition, at lags 2 and greater. This corresponds exactly to the reverse adaptation effect that was observed in the empirical data.

The joint-learning model’s explanation of the reverse-adaptation effect can also be seen more formally, as a consequence of the cue

competition term $-E_{n-1}w_{rep}$ in Equation 15. In the positive condition, long-term learning of the repetition rate leads to a positive expected value of w_{rep} . Consequently, the cue competition term biases w_{base} away from E_{n-1} at the conclusion of trial n . Thus, the expectancy for trial $n + 1$ is biased away from the outcome of trial $n - 1$, reducing the first-degree RE at lag 2. Because the bias in w_{base} persists (only decaying due to subsequent updates), the first-degree RE at longer lags is reduced as well. In the negative condition, the effect is reversed. On average w_{rep} is negative, so the cue competition term biases w_{base} toward E_{n-1} . The result is an increase in the first-degree RE at lags 2 and greater. Thus, cue competition leads $RE_1(k)$ for $k \geq 2$ to be larger in the negative condition than in the positive condition, in agreement with the reverse adaptation effect.

Joint-learning model fits to RT data. The joint-learning model’s explanation of the alternation advantage and reverse adaptation effect suggest it is a more accurate model of sequential effects than is the parallel-learning model. To test this assertion, the joint-learning model was fit to the mean RTs of each subject for all 4-deep trial histories, by least squares. Model predictions were generated by simulating the model on the actual stimulus sequence experienced by each subject. Exact simulation was used because the joint-learning

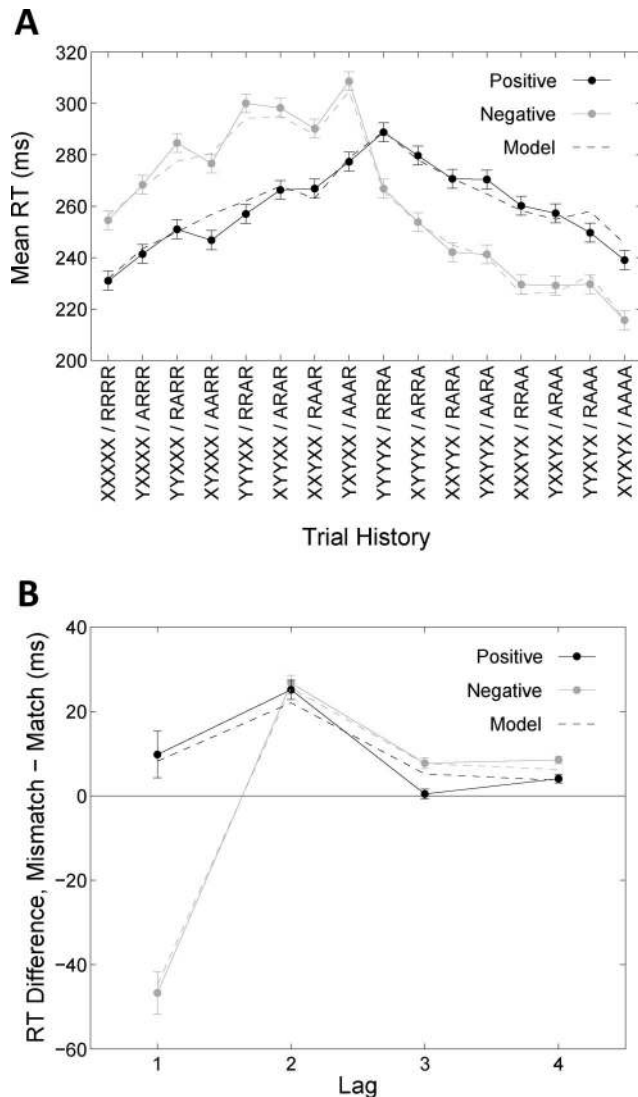


Figure 13. Fit of joint-learning model to response time (RT) data from Experiment 1. A: Mean RT for each 4-deep trial history in each condition. Error bars indicate within-subjects standard error. B: First-degree lag profile, predicted from fit in Figure 13A. Error bars indicate between-subjects standard error. Positive = positive condition; Negative = negative condition.

model does not admit simple analytic derivations in the way the parallel-learning model does. Mean predicted RTs were calculated from the same trials that were included in the subject analyses (e.g., excluding trials with recent errors), using the same procedure of balancing over 5-deep histories to obtain means for 4-deep histories.⁶ Model parameters were constrained to be equal for both conditions within each subject.

Figure 13A shows the joint-learning model fits averaged over subjects. These fits were converted to first-degree lag profiles, and the resulting predictions are shown in Figure 13B. The joint-learning produces the basic pattern of first- and second-degree REs just as the parallel-learning model does, but the joint-learning model's predictions are superior in two ways. First, it captures the

net alternation advantage, in that the first-degree RE at lag 1 is negative when averaged between conditions, and it does this without the ad hoc bias parameter needed by the parallel-learning model. Second, it captures the reverse adaptation effect, in that the first-degree RE at lags 2 and greater is larger in the negative condition. A further sign that the joint-learning model provides a better explanation of the data than the parallel-learning model is that it explains a greater proportion of variance when fit to the mean RTs across subjects (see Table 2), even though it has one less free parameter (5 vs. 6).

As noted above, the model predicts an alternation advantage only if $\beta_{\text{rep}} > \beta_{\text{base}}$. A Wilcoxon signed-ranks test on the difference $\beta_{\text{rep}} - \beta_{\text{base}}$ estimated from individual subjects showed a reliable difference (median difference = 52.5 ms, $z = 2.39$, $p < .05$), with 24 of the 28 subjects showing $\beta_{\text{rep}} > \beta_{\text{base}}$. This result supports the prediction that repetition expectancy has a more pronounced effect on RT than does base-rate expectancy, which is elemental to the model's explanation of the alternation advantage.

As a more stringent test of whether the model fully captures the alternation advantage in the data, a biased joint-learning model was fit by adding the same $-\delta E_{n-1} E_n$ term to the model's RT prediction that was added to the parallel-learning model (i.e., using Equation 10). If the joint-learning model's explanation for the alternation advantage were insufficient (or oversufficient), then it would be expected to take advantage of this extra flexibility by yielding fitted values of the bias parameter (δ) that systematically differ from zero. This was not the case. A Wilcoxon signed-rank test on the fitted δ values for all subjects showed no reliable difference from zero (median $\delta = 8.2$ ms, $z = 0.4$, $p = .68$), and a t -test produced the same conclusion ($p = .49$). Therefore, the joint-learning model provides a simultaneous explanation of both the alternation advantage and the more detailed sequential effects found in this experiment.

Summary of RT Modeling

The analyses presented above reveal a great deal about the psychological mechanisms underlying sequential effects in 2AFC. Table 5 lists the critical phenomena found in the RT data, as well as which of the two models explains each and what the explanation is.

The parallel-learning model, which assumes independent learning of the base rate and repetition rate in the trial sequence, successfully explains the dependence of RT on trial history as a combination of first- and second-degree REs. It also explains the primary difference between conditions, in that repetition RTs are faster in the positive condition and alternation RTs are faster in the negative condition, as a consequence of long-term learning of the repetition rate. The fact that short- and long-term effects can be explained by a common mechanism supports the basic premise of sequential effects as arising from incremental learning. However, the parallel-learning model fails to account for the overall alternation advantage (requiring an ad hoc bias parameter to match the

⁶ On trials when the subject made an error (i.e., stimulus and response were not the same), we arbitrarily used the stimulus to define the feedback for learning (i.e., E_n). This choice should have little impact on model predictions because of the exclusion of trials with current or recent errors in calculating mean RTs.

Table 5
Model Explanations of Key Phenomena in Experiment 1 Response Time Data

Phenomenon	Explained by models		Explanation
	Parallel learning	Joint learning	
First-degree recency effect	Yes	Yes	Incremental learning of base rate
Second-degree recency effect	Yes	Yes	Incremental learning of repetition rate
Lag-1 difference between conditions	Yes	Yes	Long-term learning of repetition rate
Reverse adaptation	No	Yes	Cue competition on base rate
Alternation advantage	No	Yes	Cue competition on repetition rate

data), or the reverse adaptation effect whereby first-degree REs at higher lags are stronger in the negative condition.

The joint-learning model accounts for all of the phenomena explained by the parallel-learning model, and it also explains the alternation advantage and the reverse adaptation effect, as signatures of cue competition from joint error correction. Cue competition from w_{base} on the learning dynamics of w_{rep} produces an alternation advantage, as w_{rep} compensates for the positive first-degree RE in w_{base} (assuming second-degree expectations have a stronger effect on RT, that is, $\beta_{rep} > \beta_{base}$ as was found in the model fits). Cue competition from w_{rep} on the learning dynamics of w_{base} produces the reverse adaptation effect, because long-term learning of w_{rep} leads expectancies to be further from the previous trial in the negative condition, and thus error correction pushes w_{base} further toward the previous trial in that condition. The joint-learning model thus produces an excellent fit to the full pattern of RT data, including sequential effects, the lag profiles, and long-term differences between the conditions.

We consider it a remarkable result that the two mechanisms introduced by the joint-learning model (i.e., the two types of cue competition) align exactly with the two empirical patterns that the parallel-learning model cannot explain. Neither of these phenomena was anticipated, and the experiment was designed as a focused test of the parallel-learning model (together with the direct and adaptation hypotheses). In conclusion, the sequential effects in this experiment turned out to be informative not just about the representations on which learning operates (i.e., physical stimulus or response identities, and repetitions vs. alternations) but also about the nature of the learning mechanism itself (i.e., that it is driven by joint error correction).

ERP Results

Having established the joint-learning model as a comprehensive and parsimonious explanation for the various phenomena observed in the RT data, we now test whether it can predict the ERP data. ERP analyses focused on LRP, to estimate separate effects on the durations of stimulus and response processing, and on P100, to estimate stimulus expectancies. In both cases, we tested the separate-stages hypothesis by evaluating whether these indices of response and stimulus processing could be separately predicted by the first- (w_{base}) and second-degree (w_{rep}) components of the model, respectively. These were parameter-free tests, in that the model as fit to the RT data was used directly to make predictions for the ERP data, without further parameter tuning (other than linear scaling parameters).

One concern with this approach might be that the RT and ERP analyses used slightly different sets of trials. As reported above, the proportions of excluded trials were very similar for RT (32.5%) and ERP (31.8% for peristimulus analyses and 31.4% for perireponse), but the exclusionary criteria were different. Most important for the goal of predicting ERP from RT is that the number of trials within each history and autocorrelation condition do not markedly differ between RT and ERP analyses. Therefore, we determined the number of retained trials for each 3-deep history within each condition. The RT and ERP trial counts differed by less than 10% in every case.

LRP. The LRP analysis is founded on the assumption that variation in sLRP and in LRPr, respectively, reflects the contributions of stimulus and response expectancies. Specifically, the sLRP interval should be shorter when the stimulus is more expected, and the LRPr interval should be shorter when the response is more expected. Thus, the separate-stages hypothesis predicts that sLRP and LRPr will be well fit by the second- and first-degree components of the joint-learning model, respectively.

Figure 14 shows aggregate LRP waveforms for each condition and 2-deep history (obtained by averaging pairs of waveforms for 3-deep histories). LRP is characterized by a spike beginning 100–200 ms after the stimulus onset and 50–100 ms before the response. The figure indicates several effects of history and condition on sLRP and LRPr. These and higher-order predictions of the joint-learning model are formally tested below.

As noted in the Method section, three different criteria were used to define the time of LRP onset in both peristimulus (sLRP) and perireponse (LRPr) analyses. To define a target for modeling, we wanted to use the onset criteria that provided the greatest signal-to-noise ratio across the histories and conditions, which is equivalent to having the greatest *F* ratio in a one-way repeated-measures ANOVA over the 16 history-condition pairs. The optimal criteria by this measure were 50% of peak negative activation for sLRP and $-0.8 \mu\text{V}$ for LRPr. All results reported in this section are based on these onset criteria, although the other criteria lead to the same qualitative conclusions.

Figure 15 shows sLRP and LRPr as a function of 3-deep history and condition, derived using the onset criteria just given. To test the separate-stages hypothesis, the joint-learning model was used to generate separate predictions for sLRP and LRPr, with the first-degree component predicting LRPr and the second-degree component predicting sLRP, as shown in Table 4. The model parameters, β_{base} , ϵ_{base} , β_{rep} , and ϵ_{rep} , as well as the trial-by-trial values of w_{base} and w_{rep} , were taken directly from the fits to the RT data. Therefore, the only free parameters in fitting the LRP data

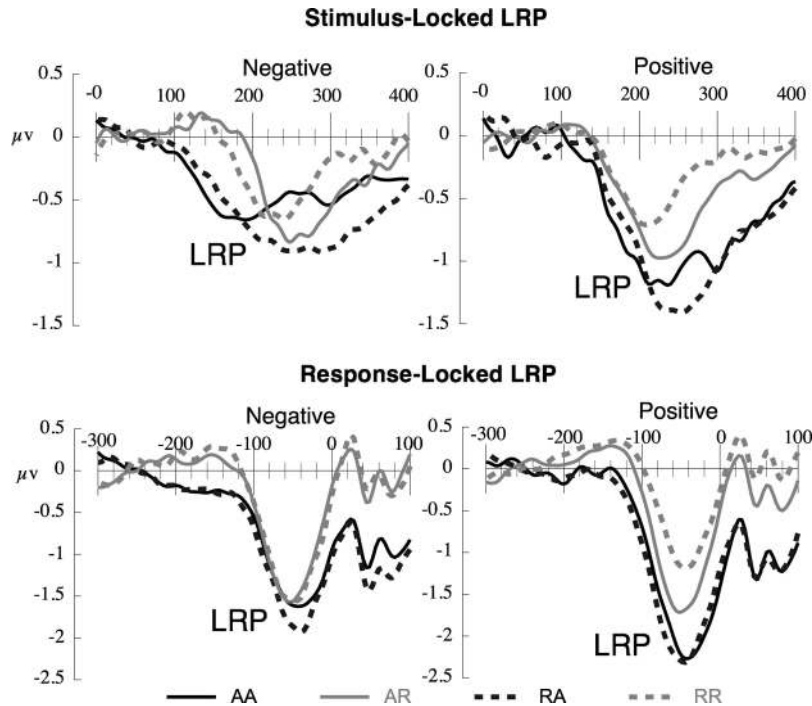


Figure 14. Average lateralized readiness potential (LRP) waveforms from Experiment 1, for all 2-deep histories in both conditions. Positive = positive condition; Negative = negative condition.

were the intercept terms, β_{LRPr} and β_{sLRP} , which partition the β_0 intercept in the RT model (although the fits were not constrained to satisfy $\beta_{LRPr} + \beta_{sLRP} = \beta_0$).

The model predictions are shown as dashed lines in Figure 15. The correlation between data and model taken across the 16 condition-history combinations is .921 for sLRP ($p < .0001$) and .746 for LRPr ($p < .01$). The predictions for sLRP are more accurate than those for LRPr (see Footnote 8 below for one possible explanation), but we view both as impressive considering they are based on only the behavioral data. To test the purity of the dissociation, the opposite correlations were also computed. The correlation between empirical sLRP and the model's first-degree component is $-.059$ ($p = .83$), and the correlation between empirical LRPr and the model's second-degree component is $-.089$ ($p = .74$).

These results provide strong support for the separate-stages hypothesis, extending the reanalysis of Jentsch and Sommer's (2002) LRP data presented in the Separating Stimulus- and Response-Based Sequential Effects section in three ways. First, the model accounts for the effects of the autocorrelation manipulation, which differ between LRPr and sLRP. Second, it fits the data without the alternation bias used in the parallel-learning model.⁷ Third, the fits represent genuine predictions from the RT data, with no free parameters except for the intercept terms (which do not affect the model-data correlations).

A useful way to gain insight into the sequential effects in the LRP data, as well as how the joint-learning model explains them, is by consideration of lag profiles. As explained in the Lag Profiles section, lag profiles measure first- and second-degree REs separately at all lags. Lag profiles for sLRP and LRPr were defined in

the same way as for RT. For first-degree REs, $RE_1^{sLRP}(k)$ was defined at lags $k = 1, 2, 3$, as the difference in mean sLRP between the four 3-deep trial histories with first-degree mismatches between trials $n - k$ and n , and the mean sLRP for the four histories with first-degree matches between trials $n - k$ and n . For second-degree REs, $RE_2^{sLRP}(k)$ was defined at lags $k = 1$ and 2 , in the same way as $RE_1^{sLRP}(k)$ but based on second-degree matches. The corresponding measures for LRPr were defined similarly, as $RE_1^{LRPr}(k)$ and $RE_2^{LRPr}(k)$. All four lag profiles were calculated separately for the positive and negative autocorrelation conditions.

Four comparisons of the LRP lag profiles are theoretically informative for the present study. Table 6 summarizes the results of these comparisons and their implications. Statistical tests of these comparisons, as reported below, used the jackknife procedure described in the Electrophysiological Recordings section, by calculating lag profiles from the jackknife waveforms and applying the appropriate statistical corrections.

We describe first the results for higher lags, specifically $RE_1(3)$ and $RE_2(2)$, because they most transparently distinguish between first- and second-degree effects. These two measures are closely related, because as Table 7 shows, they represent contrasts over the same minimal pairs of histories, each pair differing in first-degree representation only at lag 3 and in second-degree representation only at lag 2. The two effects agree in direction for half of these

⁷ Fits of the joint-learning model to Jentsch and Sommer's (2002) data are not reported here, but the results are comparable to those reported in the Separating Stimulus- and Response-Based Sequential Effects section for the parallel-learning model, on one less free parameter.

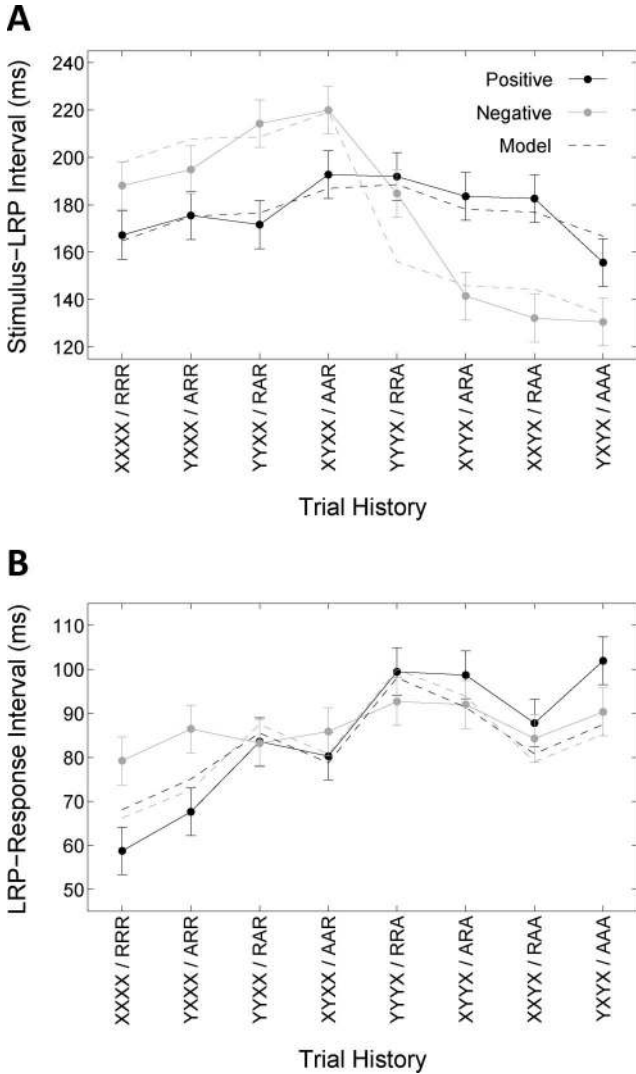


Figure 15. Estimates of (A) stimulus processing time, sLRP, and (B) response execution time, LRPr, based on lateralized readiness potential (LRP) from electroencephalography recordings in Experiment 1. Empirical values were obtained from the aggregate event-related potential waveforms averaged over subjects. Error bars indicate within-subjects standard error and are estimated using the jackknife procedure of Miller et al. (1998). Model predictions are derived from fits of the joint-learning model to the response time data, under the separate-stages hypothesis. Thus, predictions for sLRP are based only on the second-degree component of the model, and predictions for LRPr are based only on the first-degree component. Positive = positive condition; Negative = negative condition.

pairs, and they disagree for the other half. Thus first- and second-degree REs both predict differences for all of these pairs, and assessment of $RE_1(3)$ and $RE_2(2)$ can determine which mechanism is responsible for these differences.

Averaging over both autocorrelation conditions, the mean value of $RE_1^{LRPr}(3)$ is reliably positive, $t(27) = 2.07, p < .05$. That is, response processing time averaged 5.02 ms faster following a first-degree match at lag 3. Mean $RE_1^{sLRP}(3)$ was not significantly different from zero ($p > .5$). Analysis of $RE_2(2)$ shows the oppo-

site pattern. Specifically, $RE_2^{rLRP}(2)$ is not reliably different from zero ($p > .5$), but $RE_2^{sLRP}(2)$ is reliably positive, $t(27) = 3.10, p < .01$, indicating that stimulus processing was 10.15 ms faster following a second-degree match at lag 2. In summary, higher-order sequential effects in LRPr are driven by first-degree matches, whereas higher-order sequential effects in sLRP are driven by second-degree matches. These results lend strong support to the separate-stages hypothesis, in agreement with the modeling analysis above.

The remaining tests of the LRP lag profiles address questions of learning mechanisms, based on $RE_1(1)$, which measures the difference in processing time between alternation and repetition trials. The following analyses are interpreted under the separate-stages hypothesis, which we now take as firmly supported by the data.

We first address the value of $RE_1(1)$ averaged between conditions. The joint-learning model predicts $RE_1^{rLRP}(1) > 0$, corresponding to a repetition advantage in response processing, and $RE_1^{sLRP}(1) < 0$, corresponding to an alternation advantage in stimulus processing. The former prediction arises from incremental learning of the base rate (i.e., a simple first-degree RE), and the latter prediction arises as a byproduct of the former, due to cue competition. The data support both predictions: mean $RE_1^{LRPr}(1)$ is reliably positive, $t(27) = 3.66, p < .01$, and mean $RE_1^{LRPr}(1)$ is reliably negative, $t(27) = -4.62, p < .001$. These effects can be seen in Figure 14, where the time from stimulus to LRP onset is greater for repetition trials (upper panels, grey curves compared to black), and the time from LRP onset to response is greater for alternation trials (lower panels). These results support the joint-learning model's explanation for the net alternation advantage in the RT data, as arising from cue competition in second-degree learning, which compensates for the repetition advantage inherent in first-degree learning.

Finally, we consider the difference in $RE_1(1)$ between conditions. The RT data show a large difference, with a repetition advantage in the positive condition and a (stronger) alternation advantage in the negative condition. A primary question motivating this study was whether this difference is due to adaptation of the mechanism underlying first-degree REs or to a direct effect of second-degree learning. These hypotheses are indistinguishable from the RT data at lag 1, but using LRP and the separate-stages hypothesis we can break the difference between conditions into first- and second-degree effects. The direct hypothesis predicts the difference to lie in sLRP, due to long-term learning of the stimulus repetition rate. The adaptation hypothesis predicts an additional effect in LRPr, due to strengthening (positive condition) or weakening (negative condition) of the sequential effects produced by incremental learning of the base rate (i.e., of first-degree REs).

The RE profiles from the LRP data support both predictions. In stimulus processing, there is a large and reliable difference in $RE_1^{sLRP}(1)$ between conditions ($M_{Positive} = -19.35$ ms, $M_{Negative} = -55.84$ ms, $t(27) = 3.60, p < .01$), consistent with the direct hypothesis. In response processing, there is a smaller but still reliable difference in $RE_1^{LRPr}(1)$ between conditions ($M_{positive} = 24.04$ ms, $M_{negative} = 5.60$ ms, $t(27) = 3.66, p < .01$), indicating adaptation of the strength of first-degree REs. These effects can both be seen in Figure 14, where the alternation advantage in sLRP (earlier LRP onset for alternation trials) is greater in the negative condition, and the repetition advantage in LRPr (shorter time from LRP to response in repetition trials) is greater in the positive condition. Note

Table 6
Critical Measures of Recency Effects in Event-Related Potential Components From Experiment 1

Effect	LRPr (ms)	sLRP (ms)	P100 (μ V)	Conclusion
$RE_1(3)$	5.02*	-.35	.008	First-degree effects lie in response processing
$RE_2(2)$	1.24	10.15*	.075*	Second-degree effects lie in stimulus processing
$RE_1(1)$	14.82*	-37.60*	-.155*	Alternation advantage is due to cue competition
$RE_1(1), P - N$	18.44*	36.49*	.208*	Support for both direct and adaptation hypotheses

Note. LRPr = latency from lateralized readiness potential (LRP) onset to response; sLRP = latency from stimulus onset to LRP onset. Positive values for LRPr and sLRP indicate shorter processing time following matches to recent trials. Positive values for P100 indicate greater amplitudes following matches to recent trials. $P - N$ indicates difference between positive and negative autocorrelation conditions. First three rows show averages across the two conditions.

* $p < .05$.

that the evidence against adaptation in the RT data came from the reverse adaptation effect at lags 3 and 4, which was attributed to cue competition, but it is possible that adaptation and cue competition are both present at higher lags, with the latter exerting a stronger effect.⁸

In summary, the model-based and statistical analyses of the LRP data both provide excellent support for the joint-learning model and the separate-stages hypothesis. First- and second-degree sequential effects appear to be separately localized to response and stimulus processing, respectively. The data also support the joint-learning model's account of the alternation advantage in RT, as due to a cue competition effect on learning the stimulus repetition rate coupled with slower speed of stimulus versus response processing (i.e., $\beta_{\text{rep}} > \beta_{\text{base}}$). Finally, if one accepts the separate-stages hypothesis, then the comparison of $RE_1^{\text{LRPr}}(1)$ between conditions gives the first evidence for the adaptation hypothesis that is not confounded by direct contributions of second-degree learning. This finding could be incorporated in the model by allowing differential β_{base} or ϵ_{base} between conditions, but it remains a challenge for future work to determine the learning mechanism from which such differential values might emerge.

P100. Figure 16 shows P100 waveforms averaged across subjects for each 2-deep history in each condition. The P100 component is the positive peak between 100 and 150 ms after the stimulus onset. The waveforms show changes in P100 amplitude as a function of history and condition. These and higher-order predictions of the joint-learning model are formally tested below.

The primary prediction regarding P100 amplitude was that it would correlate with the second-degree component of the joint-learning model. This prediction was based on the evidence above that second-degree sequential effects are located in stimulus processing, together with the standard interpretation of P100 as arising in visual processing (e.g., Spehlmann, 1965). To test this hypothesis, the RT predictions of the model were separated into first- and second-degree components as was done for the LRP predictions. Because P100 amplitude is greater for expected than for unexpected stimuli (Mangun & Hillyard, 1991), we defined the degree to which the current stimulus matches the model's first- and second-degree expectations as

$$\begin{aligned} P_1 &= w_{\text{base}}E_n \\ P_2 &= w_{\text{rep}}E_{n-1}E_n \end{aligned} \quad (17)$$

These predictions are identical to those made above for LRPr and sLRP (see Table 4), except for a reversal of sign and omission of

the scaling parameters (β s). Because P100 is defined on a different scale than RT and the LRP intervals (i.e., voltage rather than time), we simply evaluated the correlations of P_1 and P_2 with P100 amplitude. As with the LRP predictions, the values of w_{base} and w_{rep} were taken directly from the RT fits. The hypothesis given above, that P100 reflects the second-degree component of the model, translates to the specific prediction that empirical P100 should correlate positively with P_2 and not with P_1 .

To test this prediction, P_1 , P_2 , and the P100 amplitude were averaged across subjects for each 3-deep history in both conditions. Over the 16 history-condition pairs, the correlation between P100 and P_2 is .706 ($p < .01$), whereas the correlation with P_1 is .071 ($p > .5$). Therefore, P100 appears to selectively reflect the second-degree learning component of the model. This result is illustrated in Figure 17, which shows P100 amplitude as a function of condition and history, together with the model predictions from the second-degree component. These predictions involve no free parameters, other than the two degrees of freedom for linearly scaling them into voltage units. The predictions are not as accurate as they were for the sLRP data, but they are still quite good. The fact that RT can be used to generate good predictions of a variable as different as P100 amplitude (which does not even involve timing, as LRP does) suggests the model is capturing a fundamental psychological process underlying both measures.

As a model-free test of the connection between P100 amplitude and second-degree sequential effects, the same statistical comparisons of first- and second-degree lag profiles that were performed on the LRP data were repeated with P100. Lag profiles for P100 were defined up to $k = 3$ for $RE_1^{\text{P100}}(k)$ and up to $k = 2$ for $RE_2^{\text{P100}}(k)$, following essentially the same definition used for RT and LRP. The only difference was a reversal of sign: The mean P100 amplitude over all histories having a (first- or second-degree) mismatch at lag k was subtracted from mean P100 over histories having a match at lag k . Thus, the values of RE_1^{P100} and RE_2^{P100} reflect the increase in amplitude due to a match to each past trial.

The same four comparisons tested with the LRP lag profiles were repeated on P100 (see Table 6). The predictions were the

⁸ The evidence for the adaptation effect explains why the joint-learning model predicts sLRP better than it does LRPr. The model embodies the direct hypothesis, in that it assumes equal values of β_{base} and ϵ_{base} in both conditions. Thus, its first-degree expectancies (i.e., its predictions for LRPr) are nearly identical between conditions, whereas empirically the adaptation effect leads to slightly different patterns of LRPr between conditions (see Figure 15).

Table 7
Contrast Between $RE_1(3)$ and $RE_2(2)$

Recency effect	History			
	XXXX/RRR vs. YXXX/ARR	YYXX/RAR vs. XYXX/AAR	YYXX/RRA vs. XYYX/ARA	XXYX/RAA vs. YXYX/AAA
First-degree	<	>	>	<
Second-degree	<	<	>	>

Note. Each entry indicates which of the two trial histories is predicted to show faster processing on the current trial, due to either first- or second-degree recency effects alone. For each pair, differences in first-degree effects arise only from trial $n - 3$, and differences in second-degree effects arise only from trial $n - 2$. The table shows $RE_1(3)$ and $RE_2(2)$ are orthogonal across the full set of 3-deep histories.

same as those made above for sLRP, because both of these measures are assumed to reflect stimulus processing. Starting with REs at higher lags, which are fairly pure measures of first- and second-degree learning, $RE_1^{P100}(3)$ was not significantly different from zero ($p > .5$), whereas $RE_2^{P100}(2)$ was reliably positive, $t(27) = 2.55$, $p < .05$. Thus, P100 amplitude was increased by a second-degree match two trials back, but not by a first-degree match three trials back. This pattern matches that seen in sLRP and is opposite that seen in LRP, and it supports the hypothesis that P100 reflects second- but not first-degree learning.

The tests of $RE_1^{P100}(1)$ also replicated the findings from the LRP analysis. Averaged over both autocorrelation conditions, $RE_1^{P100}(1)$ was reliably negative, $t(27) = -3.44$, $p < .01$. This effect can be seen in Figure 16 as the overall greater P100 amplitude on alternation trials. This finding reflects an overall alternation advantage in stimulus processing that is explained in the joint-learning model by cue competition. Comparing conditions, $RE_1^{P100}(1)$ was reliably more negative in the negative condition ($M_{\text{positive}} = -0.051 \mu\text{V}$, $M_{\text{negative}} = -0.259 \mu\text{V}$, $t(27) = 2.29$, $p < .05$). This effect can be seen in Figure 16 in that the amplitude difference between alternation and repetition trials is greater in the Negative condition. This finding is explained in the model as long-term learning of the repetition rate.

In summary, both the model-based and statistical analyses of P100 replicate the LRP findings, supporting the hypothesis that stimulus processing reflects second-degree but not first-degree sequential effects. Specifically, the P100 amplitudes show evidence of second-degree REs from incremental learning of the repetition rate, differences between conditions due to long-term

learning, and cue competition from learning from joint prediction error.

Experiment 1 Conclusions

The RT and ERP data from Experiment 1, together with the model-based analyses, provide converging evidence for a remarkably consistent picture of the psychological mechanisms underlying sequential effects in 2AFC, summarized in the joint-learning model and the separate-stages hypothesis.

The RT data show evidence of first- and second-degree REs, as well as long-term learning of the differential repetition rates between conditions. These findings are consistent with the parallel-learning model, which holds that sequential effects arise from separate incremental learning of the sequence base rate and repetition rate. However, the RT data also show a net alternation advantage and a reverse adaptation effect (i.e., long-lag first-degree REs are stronger in the negative condition), which the parallel-learning model cannot explain. These two phenomena turn out to align exactly with the two cue-competition effects introduced by the joint-learning model, which holds that first- and second-degree learning are driven by a common prediction-error signal. The joint-learning model thus provides an excellent fit of the RT data.

The LRP and P100 data also support this model, together with the separate-stages hypothesis. Sequential effects in LRP are well-explained by first-degree REs, with no evidence of second-degree effects. Conversely, sLRP and P100 show no evidence of first-degree effects and are well-explained by second-degree REs,




Figure 16. Average P100 waveforms in Experiment 1, for all 2-deep histories in both conditions. Positive = positive condition; Negative = negative condition.

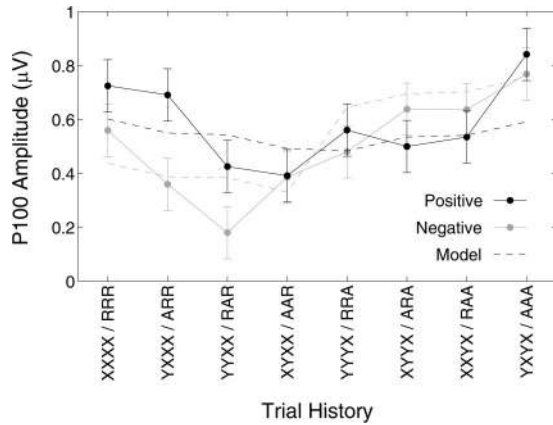


Figure 17. Mean P100 amplitude by history and condition. Error bars indicate within-subjects standard error. Dashed lines show predictions from the second-degree component of the joint-learning model. Positive = positive condition; Negative = negative condition.

long-term learning of the repetition rate, and cue competition. The latter two measures thus support the present theory in the same ways, despite being quite different measures of stimulus processing (timing in motor cortex vs. amplitude in visual cortex). Furthermore, when the joint-learning model is fit to RT and its predictions are separated into first- and second-degree expectancies, LRPr is well predicted by the first-degree component alone, and sLRP and P100 are well predicted by the second-degree component alone. The fact that the model makes such good parameter-free predictions across such different variables suggests that it captures fundamental psychological processes that underlie all of these measures.

Although the ERP predictions were based on the group data, it would be interesting in future work to test whether individual differences could be predicted in the same way. That is, subjects with stronger first-degree effects in RT might show stronger sequential effects in LRPr, and likewise subjects with stronger second-degree effects in RT might show stronger sequential effects in sLRP and P100.⁹

Finally, the results show mixed support for the adaptation hypothesis, that first-degree REs become stronger with positive autocorrelation in the environment. The RT data appeared to show a reverse adaptation effect, but this was explained in the joint-learning model as cue competition. On the other hand the LRPr data do show positive adaptation. To the extent that LRPr is affected only by first-degree learning, this is the first evidence for adaptation of REs that is not confounded by second-degree learning (i.e., by the direct hypothesis). It is possible that adaptation is present in RT as well but that it is masked by the cue-competition effect. Thus, the data as a whole are consistent with the adaptation hypothesis, but more research will be needed to better answer this question.

Experiment 2

Experiment 2 sought to separately manipulate the speed of stimulus and response processing and to test the impact on sequential effects. Stimuli were random-dot kinematograms, and the

subject's task was to discriminate between leftward and rightward motion. Motion coherence in the stimulus was varied between blocks, as a way of manipulating the rate of stimulus processing (e.g., Gold & Shadlen, 2000). In addition, each subject was assigned to respond throughout the experiment with the middle and index fingers of either the dominant or the nondominant hand. The dominant hand was predicted to exhibit faster response execution.

These manipulations provide further tests of the joint-learning model and the separate-stages hypothesis, as follows. According to this theory, overall RT depends on the summed contributions of a stimulus-processing component that is influenced by second-degree expectancies and a response-processing component that is influenced by first-degree expectancies (see Table 4). The timing of these two components is determined by the parameters β_{rep} and β_{base} , respectively. Slowing either processing component should manifest in an increase of the corresponding β parameter. Because RT depends on the product of each β and the corresponding expectancy, w , slowing a component (i.e., increasing its β) should also yield an amplification of expectancy effects on RT. In other words, slowing stimulus processing should magnify the effects of second-degree expectancies, and slowing response processing should likewise magnify the effects of first-degree expectancies.

We therefore predicted that, provided the stimulus and response manipulations in this experiment affected overall RT, they would also moderate the effects of second- and first-degree expectancies, respectively. Specifically, lower stimulus coherence should increase the magnitude of second-degree REs and increase the alternation advantage, and responding with the nondominant hand should increase the magnitude of first-degree REs and decrease the alternation advantage (or produce a repetition advantage). Because autocorrelation was not manipulated in this experiment, predicted effects on long-term learning and the reverse adaptation effect were not tested.

Method

Participants. One hundred eighty-one undergraduate students participated for partial course credit. Each subject was randomly assigned to respond using the dominant ($n = 90$) or nondominant ($n = 91$) hand. Dominant hand was determined by self-report, with 21 subjects (12 in the dominant condition) reporting being left-handed.

Stimuli. The stimulus on each trial was a random-dot kinematogram, displayed on a black background on a 20-in. (50.8 cm) computer monitor with 60-Hz refresh rate. The stimulus region was circular, in the center of the monitor, with radius subtending approximately 2.5 degrees of visual angle. Dots were white with diameter approximately .6 mm.

Five dots were present on each refresh frame of the monitor. Three interleaved sequences of dots were shown, so that the dots on each frame corresponded to the dots from three frames earlier. On each frame, every dot had an independent probability of moving coherently relative to its position three frames earlier, versus jumping randomly to a new point in the stimulus region. This probability, referred to as the coherence of the stimulus, was 40% in low-coherence blocks and 80% in high-coherence blocks.

⁹ We thank an anonymous reviewer for this suggestion.

The stimulus on each trial was either a rightward or a leftward stimulus. For rightward stimuli, every time a dot moved coherently, it shifted 3.5 mm to the right from its previous position (with no vertical change). This shift corresponds to motion at approximately 6.2 degrees of visual angle per second. If the motion took the dot off the right edge of the stimulus region, it was replaced on the corresponding point on the left edge. Leftward stimuli were constructed similarly.

Design. Each subject performed five blocks of trials in a single session. The first block comprised 35 trials at high coherence and was excluded from all analyses as practice. Of the remaining 4 blocks, 2 used high coherence and 2 used low coherence, with the order counterbalanced across subjects (high-low-high-low or low-high-low-high). Each of these blocks comprised 201 trials, the first 9 of which were excluded from analysis to reduce transient effects following breaks. Thus, there were 192 test trials in each of the four test blocks.

Local stimulus histories for test trials were controlled to a depth of 6 trials, with each of the 64 possible sequences of leftward and rightward stimuli occurring exactly three times in each test block. The stimulus sequence for each test block was generated pseudo-randomly according to this constraint. The excluded trials in all five blocks were generated pseudo-randomly under the constraints that the repetition rate in each block was exactly 50%, and the base rate was as close to 50% as possible (each block had an odd number of trials).

Each subject responded with either the dominant or the nondominant hand throughout the experiment. This assignment was counterbalanced across subjects.

Procedure. The subject's task on each trial was to respond to the motion of the dots, by pressing G for left or H for right on a standard keyboard. Each subject was instructed to keep the index and middle fingers of the assigned hand on these two keys throughout the experiment.

Each trial began with a 700-ms RSI, during which only a 1-mm white fixation dot was present in the center of the monitor. The stimulus was then displayed and continued until the subject made a response. Following error responses, the word "Wrong" was displayed in red in the center of the monitor throughout the RSI, in place of the fixation dot. Subjects were given self-paced breaks between blocks. The entire experiment lasted about 30 min.

Results

To assess sequential effects, test trials were classified according to 4-deep histories, and the mean RT was computed for each history in each coherence condition for every subject. Because the repetition rate was exactly 50% in all blocks, there was no need to balance over 5-deep histories as was done in Experiment 1. Trials were excluded if the present RT was less than 100 ms or greater than three standard deviations above the mean for the test trials on the current block (the 1,000-ms criterion used in Experiment 1 excluded too many trials), or if the present or any of the three previous responses was incorrect. These criteria led to exclusion of 24.1% of the test trials. One subject in the nondominant condition had no valid trials for one cell (history ARRA with low coherence) and was eliminated from all subsequent analyses. Thus, 90 subjects were analyzed in each condition. An alternative analysis that required only the current and one previous responses to be correct,

which excluded 14.9% of the test trials and included all subjects, yielded the same qualitative results as those reported here.

Table 8 shows mean RT as a function of stimulus coherence and response hand, averaged over trial histories and subjects. The manipulation of stimulus coherence had a large effect (198.9 ms), whereas the effect of response hand was much less and opposite the predicted direction (11.6 ms faster for the nondominant hand). To assess the influences of the stimulus and response manipulations on RT and on sequential effects, a mixed-effects ANOVA was applied to the mean RTs, with 4-deep history and coherence as within-subjects factors and dominance as a between-subjects factor. Results showed highly reliable effects of coherence, $F(1, 178) = 555.7, p < .0001$; history, $F(9.03, 1607.63) = 47.2, GG \epsilon = .602, p < .0001$; and their interaction, $F(10.86, 1933.83) = 9.77, GG \epsilon = .724, p < .0001$, indicating the coherence manipulation affected both overall RT and the pattern of sequential effects. However, neither the main effect nor any interaction involving dominance approached significance, even without sphericity corrections ($ps > .2$). Therefore, it appears that the response manipulation had no effect in this experiment. All subsequent analyses collapse over this variable.

Figure 18 (solid lines) shows mean RT as a function of trial history and stimulus coherence, calculated separately for each subject and averaged over subjects. As evident in the figure, there is a large alternation advantage in the low-coherence condition, with mean RT being 37.10 ms faster on alternation trials, and a weaker alternation advantage of 4.35 ms in the high-coherence condition. The difference in alternation advantage between conditions is reliable by paired *t*-test, $t(179) = 7.135, p < .0001$. This result supports the prediction that slowing stimulus processing will increase the alternation advantage. Moreover, the effect on the alternation advantage is not simply a byproduct of overall slowing of RT, because the 37% increase in RT between conditions is much less than the 8.5-fold increase in the alternation advantage. Similarly, the range of sequential effects in the low-coherence condition (131 ms, between the fastest history, AAAA, and the slowest, AAAR) is only 85% greater than the range of sequential effects in the high-coherence condition (71 ms). Therefore, the coherence manipulation affected the alternation advantage much more than it did sequential effects in general.

Turning to sequential effects from lags greater than 1, it is evident from Figure 18 that these effects are stronger in the low-coherence condition (because each half of the high-coherence curve is flatter than the corresponding half of the low-coherence curve). As with the alternation advantage, one possibility to consider is that this difference is just a byproduct of the overall slowing of RT. However, the separate-stages hypothesis makes a more specific prediction, namely that the difference between con-

Table 8
Mean Response Time (in Milliseconds) by Condition in Experiment 2

Response hand	Stimulus coherence	
	High	Low
Dominant	536.5	744.4
Nondominant	533.9	723.8




Figure 18. Mean response time (RT) by trial history for Experiment 2, together with fits of the joint-learning model. Error bars correspond to the within-subjects standard error taken across all 32 observations. Low = low stimulus coherence; High = high stimulus coherence.

ditions should be restricted to second-degree effects. To test this prediction, the joint-learning model was fit to the mean RTs of each subject for all 4-deep trial histories, by least squares. Model predictions were generated by simulating the model on the actual stimulus sequence experienced by each subject, with mean predicted RTs calculated from the same trials that were included in the empirical analysis (e.g., excluding error trials). For each subject, the model was fit separately to the low- and high-coherence conditions, to test how the stimulus manipulation affected parameter values. To allow unbiased comparisons of β_{base} and β_{rep} between conditions, these parameters were not constrained to be nonnegative (as they were in the fits of Experiment 1).

Median parameter values for both conditions are reported in Table 9, and model fits averaged across subjects are shown as dashed lines in Figure 18. Wilcoxon signed-ranked comparisons were used to compare each parameter between conditions, because of heavy-tailed distributions for the β parameters (t -tests on the ϵ parameters led to the same conclusions as reported here). These tests showed significant differences for β_0 ($z = 11.62$, $p < .0001$) and β_{rep} ($z = 4.36$, $p < .0001$), but not for β_{base} , ϵ_{base} , or ϵ_{rep} ($ps > .1$). Thus, the influence of stimulus coherence on sequential effects appears to have been localized to the second-degree component of the model, as predicted by the separate-stages hypothesis. Specifically, the differences in β_0 and β_{rep} are both consistent with a slowing of stimulus processing, leading to an increased dependence of RT on second-degree expectancies. As is evident in Figure 18, the change in scaling of the model's second-degree component does a good job of explaining the difference between conditions both in the alternation advantage and in the detailed pattern of sequential effects.

Experiment 2 Conclusions

Although the manipulation of response hand appeared to have no effect, the manipulation of stimulus coherence produced a large

difference in overall RT. The slowed responding in the low-coherence condition presumably resulted from a reduction in the rate of sensory evidence accumulation, as suggested by a large body of behavioral and neurophysiological evidence with these stimuli (Gold & Shadlen, 2007).

According to the separate-stages hypothesis, this slowing of stimulus processing should selectively increase the contributions of second-degree expectancies (i.e., expectancies of repetition vs. alternation) to sequential effects in RT. This prediction was confirmed by fits of the joint-learning model. Comparison of estimated model parameters between conditions showed a large and reliable difference in β_{rep} , which determines the influence of second-degree expectancies on RT. There was no detectable difference between conditions in the influence of first-degree expectancies (β_{base}) or either learning rate (ϵ_{rep} , ϵ_{base}). Fits of the parallel-learning model (not reported) showed the same pattern of results.

The results also support the joint-learning model (over the parallel-learning model) and its explanation of the alternation advantage. The joint-learning model holds that the alternation advantage arises from a bias in the learned repetition rate (w_{rep}), due to cue competition from the base rate. Increasing the influence of second-degree expectancies (β_{rep}) should increase the impact of this bias on RT. This prediction was confirmed in the data, as the alternation advantage was significantly stronger in the low-coherence condition.

In conclusion, the results of Experiment 2 are in excellent agreement with the present theory. Changing the speed of stimulus processing selectively influenced the contributions of second-degree expectancies, as predicted by the separate-stages hypothesis, and this change carried with it a change in the alternation advantage, a signature of the cue competition predicted by the joint-learning model.

General Discussion

The results of the present study yield a remarkably consistent picture of the mechanisms underlying sequential effects in binary choice tasks. Specifically, sequential effects appear to result from simultaneous learning of the base rate and the repetition rate of the trial sequence. The base rate is primarily represented in terms of responses, whereas the repetition rate is primarily represented in terms of stimuli. Furthermore, those aspects of the data that cannot be explained by these two mechanisms separately are predicted by their interaction, in learning from joint prediction error.

These conclusions are supported by several converging lines of evidence from the present experiments and from previous data. First, sequential effects in RT are well explained by incremental learning of the base rate and repetition rate. When

Table 9

Median Parameter Estimates From Fits of Joint-Learning Model to Experiment 2

Coherence	β_0^*	β_{base}	ϵ_{base}	β_{rep}^*	ϵ_{rep}
Low	716.48	10.74	.128	88.20	.209
High	517.50	15.91	.105	45.70	.246

* Significant difference between conditions ($p < .0001$).

this incremental learning is formalized in the joint-learning model, it gives excellent simultaneous fits to first- and second-degree REs, long-term learning of the repetition rate, and the effects of cue competition (i.e., the alternation advantage and reverse adaptation effect). Second, this model as fit to RT yields good parameter-free predictions of sequential effects in ERP components, and moreover these effects show a clean dissociation into stimulus and response processing, in agreement with the separate-stages hypothesis. Third, manipulation of stimulus discriminability selectively influences the contribution of second-degree expectancies (i.e., of repetition vs. alternation), as seen both in second-order REs and in the alternation advantage. Fourth, past studies that eliminated stimuli (Wilder et al., 2013) or recent responses (Maloney et al., 2005) exhibit sequential effects characterized by first- or second-degree REs alone.

A strength of the present theory is that it provides a unified explanation of sequential effects in quite different dependent measures: timing (RT and LRP latency), P100 amplitude, motor trajectories, and perceptual thresholds. We suggest the common underlying construct is expectancy, which speeds or slows RT (Laming, 1968), influences the magnitude of the P100 component (Mangun & Hillyard, 1991), affects motoric compensation for external forces (Wilder et al., 2013), and biases subjects' perceptions of ambiguous stimuli (Maloney et al., 2005). This explanation is fundamental to our interpretation of sequential effects as reflecting learning of the trial sequence. The models tested here embody this assumption, using past stimuli and responses to estimate sequence statistics that generate expectancies for each upcoming trial. In this way, we have been able to use sequential effects to draw conclusions about the representations of stimulus and response sequences and the mechanisms by which those representations are learned.

The literature on sequential effects in 2AFC is over 50 years old (e.g., Bertelson, 1961), and it has produced numerous models of varying complexity and coverage of empirical phenomena (e.g., Gao, Wong-Lin, Holmes, Simen, & Cohen, 2009; Squires et al., 1976; Yu & Cohen, 2009). The present model explains nearly all of the previously reported phenomena, as well as new findings (the reverse adaptation effect and the various dissociations between stimulus and response processing), and it does so with a single learning rule operating on two simple statistics (base rate and repetition rate). However, one notable exception that the model does not address is the influence of varying the response-stimulus interval (RSI). With RSIs of about a half-second or longer, sequential effects conform to the patterns found here and are well-fit by the joint-learning model. With shorter RSIs, in the range of 50–100 ms, the pattern of sequential effects becomes very different. RTs on alternation trials are slower than on repetition trials, and the alternation portion of the plot of RT by trial history (i.e., the right half of the graphs in Figures 1, 6, and 18) has an overall upward instead of downward slope (Soetens et al., 1985; Vervaeck & Boer, 1980). One explanation offered for this pattern is that conflict monitoring during previous alternation trials increases strategic control, thus slowing RT on the current trial (Gao et al., 2009; A. D. Jones, Cho, Nystrom, Cohen, & Braver, 2002).¹⁰

Regardless of the explanation for sequential effects under short RSIs, they do not bear on the questions of the present article. Under our fundamental assumption that RT indexes expectancy, sequential effects should exhibit a tradeoff between the two pos-

sible stimuli following each history (e.g., fast for RRRR and slow for RRRA). This is the case in most 2AFC tasks with long RSIs (Audley, 1973; Soetens et al., 1985), and it holds in the present experiments as well, as shown by the strong negative correlation in the exchange plot of Figure 7. With short RSIs, this cost-benefit pattern is absent. Instead, the exchange plot shows a correlation near +1, referred to in the literature as a benefit-only (or cost-only) pattern (Soetens et al., 1985). Thus, RT varies by history regardless of the current stimulus (e.g., fast for both RRRR and RRRA). The correct explanation for the benefit-only pattern is an important question, as is the question of why it completely replaces the cost-benefit pattern under short RSIs. The critical point for present purposes is that the two patterns are orthogonal. Therefore sequential effects with short RSIs do not reflect expectancy, or learning of sequence statistics, in any direct way. Nevertheless, further work on understanding the transition from the benefit-only to the cost-benefit pattern might shed light on the within-trial dynamics of how the expectancy is generated.

Implications for Learning and Representation

We argue that sequential effects merit more attention in psychology and in particular in cognitive modeling, because they offer powerful means for identifying perceptual and cognitive representations, and for determining the learning mechanisms that operate on those representations. This approach is founded on the interpretation of (most) sequential effects as reflecting trial-by-trial learning and updating of representations in response to feedback.

The present work suggests that sequential effects in 2AFC are driven by learning of stimulus-response sequences, which in turn is grounded in two simple statistics: the base rate and the repetition rate. First- and second-degree REs can thus be thought of as signatures of these two representations. That is, the sequence or its generating process is represented in terms of estimates of these two parameters. This in turn implies that individual trials are represented both as concrete events (specifically, by the physical response that was executed) and more abstractly as repetitions or alternations. These representations correspond to two possible strategies for performing the task: One can identify the stimulus and give the appropriate response, or one can identify whether the stimulus changed from the previous trial and accordingly change or repeat the previous response (Fletcher & Rabbitt, 1978).

Of course, sequence learning is not limited in general to base rates and repetition rates. Large bodies of research on artificial grammar learning and serial reaction time show that, when a stimulus or stimulus-response sequence conforms to more complex (stochastic or deterministic) regularities, people often readily

¹⁰ A conceptual difficulty with this explanation is that those models implement control as reducing the input to a diffusion process representing evidence accumulation (i.e., changing the drift rate). This assumption goes against a large body of research fitting diffusion models to speeded choice data, which concludes the drift rate is selectively influenced by the strength of stimulus information, whereas strategic control processes selectively influence decision thresholds (e.g., Ratcliff & Smith, 2004; however, for dissenting views, see M. Jones & Dzhafarov, in press; Starns, Ratcliff, & McKoon, 2012). An alternative model in which control modulates thresholds could not explain the short-RSI data, because threshold variation produces a speed-accuracy tradeoff. Sequential effects with short RSIs are not characterized by a speed-accuracy tradeoff; RT and error rates vary in the same direction across trial histories (e.g., Soetens et al., 1985).

learn them (e.g., Cleeremans & McClelland, 1991; Nissen & Bullemer, 1987; Reber, 1967). A complete theory of sequence learning might best be founded on a Bayesian or simplicity principle, whereby the brain seeks simple structure by default and more complex structure when warranted by experience (Chater & Vitányi, 2003; Feldman, 2000). In random or near-random sequences like the ones studied here, base rate and repetition rate appear to provide a nearly complete account.

The present approach of using sequential effects to uncover sequence representation builds on earlier work focused on stimulus and category representations. M. Jones et al. (2006) analyzed first-degree REs in choice probability in a perceptual categorization task, and found the tendency to respond with the previous feedback (i.e., with the correct category response from the previous trial) is determined by the similarity between the present and previous stimuli. REs in perceptual categorization can thus be interpreted as reflecting generalization of category knowledge between successive stimuli (Shepard, 1987). By assessing the strength of the first-degree RE as a function of the present and previous stimuli, one can directly obtain a map of the similarity structure of the stimulus space. M. Jones, Maddox, and Love (2005) used this approach to test how stimulus representations change with learning, specifically with learned shifts of attention among stimulus dimensions. They found REs are stronger when successive stimuli differ on unattended dimensions than when they differ on attended dimensions. This finding supports classic theories of dimensional attention, which hold that similarity is more sensitive to attended dimensions (Medin, Goldstone, & Gentner, 1993; Nosofsky, 1986; Sutherland & Mackintosh, 1971). M. Jones (2009) applied a similar approach to the question of category representation, showing the pattern of REs in category learning matches the predictions of exemplar models but is inconsistent with prototype or rule-based (decision bound) models. Finally, Collins and Frank (2013) show the pattern of REs in RT can be used to infer a hierarchical rule or control structure that a subject spontaneously adopts in a simple association task with multidimensional stimuli.

Turning from representation to learning mechanisms, our results show the base rate and repetition rate in 2AFC are learned via error correction based on their joint prediction, as embodied in the joint-learning model. This is in contrast to the parallel-learning model, in which each of these statistics is learned separately. An alternative interpretation of the parallel-learning model is as simple priming, whereby expectancies are based on adding the decaying traces of all past events (see Equations 2 and 6). Although the parallel-learning model provides good overall fits to the data, it cannot account for two key phenomena: the alternation advantage and the reverse adaptation effect. These phenomena were found to be signatures of the cue competition posited by the joint-learning model.

The principle of cue competition via joint error correction is well established in associative learning (Rescorla & Wagner, 1972), but the present model applies this principle in a novel way, treating events from previous trials as latent or internal cues that can contribute to learning of sequence statistics (rather than of concrete stimulus-response associations). Moreover, this application of joint error correction produces a more complex type of cue competition effects. In previous findings of cue competition, such as blocking (Kamin, 1968) or base-rate neglect (Gluck & Bower,

1988), the mean or steady-state value of one association weight biases the mean value of another. In the present context, cue competition manifests in a bidirectional interaction between error correction and the very sequential effects that error correction produces. In the case of the alternation advantage, the bias on w_{rep} comes not from the mean value of w_{base} but from sequential effects in w_{base} . In the case of the reverse adaptation effect, cue competition from w_{rep} produces a bias not in the mean value of w_{base} but in its sequential effects (i.e., its relationship to past trials). Sakamoto, Jones, and Love (2008) demonstrate a similar phenomenon in learning the variability of time-varying perceptual categories, whereby sequential effects in the estimated category mean produce (via cue competition) a systematic downward bias in the estimated category variance.

One challenge to the present theory concerns the manner in which sequential effects decrease as a function of lag, as captured by the lag profile (e.g., Figure 9). Models based on error correction must predict lag profiles that, at least asymptotically, follow an exponential decrease (see Equation 8 for the parallel-learning model; a similar result holds for the joint-learning model). In contrast to this prediction, Wilder et al. (2013) find that sequential effects are better described by a power-law lag profile, with significantly nonzero REs extending more than 500 trials back. Power-law decay can easily be incorporated into trace models, as is common in many models of memory (Anderson et al., 2004; Wickelgren, 1974), but trace models cannot explain the cue competition effects observed here because they do not use error correction. Thus, an important topic for future research is to investigate how joint error correction and power-law decay might be reconciled. Because power law decay is well approximated by a mixture of exponentials, one potential solution is a multiscale model (e.g., Mozer, Pashler, Cepeda, Lindsey, & Vul, 2009) in which multiple copies of the joint-learning model operate over a range of learning rates. Variability in the learning rates would yield exponential decay on multiple time scales, which when combined would result in approximately power-law decay.

Relation to Other Models

Some previous models of sequential effects in 2AFC have explained first- and second-degree REs, but they cannot explain the full set of present results because they do not adopt the incremental learning framework used here. For example, Squires et al. (1976) proposed a model in which expectancies are based on weighted averages of past stimulus values, thus producing first-degree REs. Their model also includes a special expectancy of alternation when at least two of the last three trials were alternations. This assumption produces second-degree REs, but of a different (and weaker) form than those predicted by either the parallel-learning or joint-learning model developed here. Moreover, their heuristic framing of the alternation expectancy misses the symmetry between first- and second-degree REs, as arising from the same learning mechanism operating on different representations. As a further consequence of not casting both mechanisms as incremental learning, Squires et al.'s model cannot explain the cue competition effects predicted by the joint-learning model.

Cho et al. (2002) propose a model in which separate detectors develop expectancies based on past trials, which bias the decision-

making process on the current trial. They tested an array of detectors to determine which combination gives the best fit to data (including the data shown in Figure 1A). Cho et al. interpreted the two detectors in their best-fitting model as a repetition detector and an alternation detector, but in fact they correspond to first- and second-degree detectors. The first detector (named IR1) builds up expectancy for each stimulus based on past occurrences of that stimulus, such that the net effect (i.e., the difference between expectancies for the two stimuli) is formally identical to w_{base} in the parallel-learning model. The second detector (SA2) builds up expectancy for alternation based on each past alternation, and its contribution can be shown to be linearly related to w_{rep} in the parallel-learning model (specifically, $\frac{w_{\text{rep}} - 1}{2}$). Because this latter mechanism detects only alternations, it predicts an alternation advantage in addition to second-degree REs. However, there is no a priori reason why there should be a detector for alternations and not for repetitions. Indeed, Cho et al. also defined a detector that builds up expectancy for repetition based on each past repetition (SR2), but they did not test that detector in conjunction with IR1. Moreover, their model cannot explain the reverse adaptation effect, because its detectors operate independently. The joint-learning model explains both the alternation advantage and the reverse adaptation effect parsimoniously as consequences of the same learning principle (i.e., cue competition from joint error correction).

Wilder et al. (2010) present a Bayesian model of sequential effects in 2AFC, named DBM2, that explicitly assumes environmental nonstationarity in the base rate and repetition rate. The model builds on Yu and Cohen's (2009) finding that Bayesian inference of discrete change-points in environmental statistics leads to approximately exponentially decaying REs. DBM2 thus predicts first- and second-degree REs in close agreement with the predictions of the parallel-learning model. However, by casting learning at a normative level, DBM2 fails to anticipate the cue competition effects that result from joint error correction, and thus it does not predict the alternation advantage or reverse adaptation effect.

Cue competition thus plays a critical role in distinguishing the joint-learning model both from simpler accounts based on decaying traces of past events and from rational accounts based on expectations of nonstationarity. These three frameworks for explaining sequential effects—decaying traces, incremental learning, and rational inference—are difficult to distinguish when applied to a single task variable (e.g., response base rate), because in that case they all merely predict decaying influences of past events. However, their predictions diverge when sequential effects are driven by multiple variables simultaneously (as seen here with base rate and repetition rate). Among the various explanations for sequential effects offered in the literature, learning from joint prediction error is unique in predicting the cue competition effects that were observed here.

A final issue for sequential effects in RT concerns the within-trial dynamics that produce them. In this article, we have abstracted over the within-trial dynamics of the decision process, because that level of detail is unnecessary within our framework for explaining sequential effects. Previous models (Cho et al., 2002; Gao et al., 2009) have implemented sequential effects as

biases acting on competitive leaky diffusion processes (based on Usher & McClelland, 2001). It would be straightforward to elaborate the joint-learning model in the same way, using w_{base} and w_{rep} to determine these biases. An advantage of such an approach is that it might enable prediction of error rates and RT distributions, in addition to mean RT. Importantly, the bias parameters could determine the starting values or input levels to evidence-accumulation processes, but they cannot be identified with those processes. The model of Gao et al. (2009) incorporates the latter assumption, taking the starting evidence values on trial n to be determined by residual activation of the evidence processes on trial $n - 1$. This mechanism contributes a first-degree RE at a lag of one trial (i.e., a simple repetition advantage), but as Gao et al. note, it cannot produce sequential effects at higher lags because the decision dynamics during trial n erase the residual information from trial $n - 1$, preventing that information from carrying over to trial $n + 1$.

Further implications for modeling within-trial dynamics come from the separate-stages hypothesis. Most current models of speeded choice involve a single decision-making stage, in which evidence for each option (or net evidence between two options) evolves until some threshold is met (Brown & Heathcote, 2008; Ratcliff & Smith, 2004; Usher & McClelland, 2001). The present results suggest that a more accurate model would comprise two such processes, one for stimulus identification and another for response selection. Sequential effects would arise from w_{rep} acting on the first process and w_{base} on the second (e.g., by determining initial evidence levels at the start of each trial). The two stages might be discrete, with one beginning once the other reaches threshold, or the activation level of the first might act continuously to determine the drift rate of the second (McClelland, 1979). This view is consistent with recent neuroimaging research showing that sensory evidence accumulation and response selection are implemented in different brain regions (Filimon, Philiastides, Nelson, Kloosterman, & Heekeren, 2013). It would be an interesting question for future research to determine whether a two-stage evidence-accumulation model fits benchmark RT data better than extant one-stage models, and whether stimulus and response manipulations would be found to selectively affect model parameters of the corresponding stages (extending the present results from Experiment 2 with the joint-learning model).

Separating Stimulus- and Response-Based Sequential Effects

The question of whether sequential effects are due to past stimuli or past responses has been a focus of research in several domains (DeCarlo & Cross, 1990; Jesteadt et al., 1977; M. Jones et al., 2006; Pashler & Baylis, 1991; Soetens, 1998; Ward & Lockhead, 1971; Willingham, Wells, Farrell, & Stemwedel, 2000). The theory proposed here offers a new perspective on this question: If sequential effects arise from two separate learning mechanisms, then these mechanisms may be differentially grounded in stimulus and response processing. The present experiments and reanalysis of previous data support this idea. In particular, they substantiate what we refer to as the separate-stages hypothesis, that sequential effects due to incremental learning of the base rate are rooted in response processing, whereas sequential effects from

learning the repetition rate are rooted in stimulus processing. Support for this hypothesis comes from three lines of evidence.

First, the ERP data from Experiment 1 and from Jentzsch and Sommer (2002) show a strong dissociation between sequential effects in stimulus and response processing. For response processing, the time from LRP onset to the response (LRPr) exhibited first-degree REs and not second-degree REs. For stimulus processing, the time from stimulus onset to LRP onset (sLRP) followed the opposite pattern. The second ERP measure of stimulus processing in Experiment 1, the amplitude of the P100 component, also showed only second-degree REs. Moreover, the alternation advantage, which the joint-learning model explains as a bias in the repetition rate (via cue competition), was present in sLRP and P100, but not in LRPr. When the RT fits of the joint-learning model are divided into first- and second-degree components, they give impressive parameter-free (except for linear scaling) predictions of the ERP response and stimulus measures, respectively. Therefore, sequential effects in response processing are well-explained by first-degree learning, and sequential effects in stimulus processing are well-explained by second-degree learning.

Second, the manipulation of stimulus coherence in Experiment 2 selectively influenced the contribution of second-degree learning. This was shown by the fits of the joint-learning model, which revealed a large difference between conditions in the dependence of RT on second-degree expectancies (β_{rep}), as well as in the intercept (β_0), but not in any other parameters. This difference manifested in a larger alternation advantage in the low-coherence condition, and in more-pronounced higher-order sequential effects, which the model attributed entirely to second-degree REs. Therefore, slowing stimulus processing appears to selectively magnify these two phenomena, implying stimulus expectancies are determined by the learned repetition rate.

Third, experiments that have eliminated stimuli or past responses produce sequential effects characterized by pure first- or second-degree REs, respectively. As shown in the Separating Stimulus- and Response-Based Sequential Effects section, comparison of Experiments 1 and 2 of Maloney et al. (2005) shows that omission of previous responses selectively eliminates first-degree REs. Their first experiment required responses on every trial, and we found the data were well fit by a model that learns both the base rate and the repetition rate. In their second experiment, subjects gave no responses in the three-trial history leading up to each response trial, and we found the data were best fit by a model that learns only the repetition rate. Therefore, first-degree expectancies appear to produce sequential effects only when recent trials include responses that can drive base-rate learning. A complementary result comes from Experiment 2 of Wilder et al. (2013). Sequential effects in this task arose from random forces acting directly on the subject's movement, and they conformed to pure first-order REs, with no indication of second-degree effects. Therefore, second-degree sequential effects appear to be absent when the task contains no (visual) stimuli.

The separate-stages hypothesis may appear arbitrary, apart from its support by the results just listed, but it does have precedent in previous work attempting to separate stimulus- and response-based sequential effects. In more complex tasks such as absolute identification and magnitude estimation, where the previous stimulus and previous response are not perfectly confounded, one can assess

the effect of each while controlling for the other. Research using this approach (Jestadt et al., 1977; Petzold, 1981) has found that the present response is biased away from the previous stimulus (i.e., stimulus contrast) and toward the previous response (i.e., response assimilation). This approach suffers problems of interpretation due to the possibilities of autocorrelated response errors and relative judgment strategies (DeCarlo & Cross, 1990; Holland & Lockhead, 1968; Luce & Green, 1974), but studies that have attempted to avoid these problems have come to the same conclusions (M. Jones, 2009; M. Jones et al., 2006; M. Jones & Sieck, 2003).

M. Jones (2009; see also M. Jones et al., 2006) formalized these conclusions in a reinforcement-and-contrast model, in which incremental learning produces assimilation to the correct responses on previous trials (as in Equation 1), and stimulus percepts are biased away from previous stimuli, producing contrast (for similar proposals, see Brown, Marley, Donkin, & Heathcote, 2008; Treisman & Williams, 1984). The explanation for response assimilation is identical to that for first-degree REs in the present study: Incremental learning of response probabilities produces a bias toward recent correct responses. In identification and categorization, this effect is moderated by stimulus and response generalization (Shepard, 1957), but these mechanisms are inconsequential in binary tasks such as 2AFC. Regarding stimulus-based sequential effects, we suggest the unifying principle is that stimuli are perceived relatively to preceding stimuli, rather than absolutely (Lockhead, 2004; Stewart et al., 2002; Stewart, Brown, & Chater, 2005). In identification and categorization tasks with many stimuli, this relative perception manifests in contrast effects. In the binary tasks studied here, relative perception leads stimuli to be perceived as repetitions or alternations, which in conjunction with incremental learning of stimulus probabilities leads to second-degree REs.

Conclusions

Sequential effects are often viewed as simple phenomena, involving a bias toward recent feedback or better memory for recent events. They are generally explained implementationally or algorithmically by residual activity or incremental learning, or computationally by adaptation to a dynamic environment. When applied to a single task variable (e.g., response probability), these explanations indeed predict a simple recency bias in the weighting of past events (Estes, 1957; Yu & Cohen, 2009). However, more complex patterns can arise when REs are based on abstract or latent variables or on multiple variables at once (M. Jones, Mozer, & Kinoshita, 2009; Sakamoto et al., 2008; Treisman & Williams, 1984).

We have shown here that the complex, nonadditive pattern of sequential effects in 2AFC can be explained by two simple learning mechanisms and their interaction. This leads to a quite specific conclusion about sequence learning in these tasks: It is based on the response base rate and the stimulus repetition rate. More complex mechanisms certainly operate in richer tasks, for example with more stimuli or responses or with nonrandom sequences, but these appear to be the default mechanisms that dominate in simple tasks with unstructured sequences.

Setting aside the specifics of 2AFC, the larger point is that sequential effects offer a valuable window into many aspects of cognition, due to their close connection to learning. Not all se-

quential effects are due to learning, but nearly all learning is likely to produce sequential effects. Biological learning is generally characterized by incremental or iterated updating of knowledge, and different representations will naturally be updated in different ways. Therefore, careful consideration of sequential effects can reveal a great deal about learning mechanisms, the representations on which they operate, and their neurophysiological underpinnings.

References

- Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, *111*, 1036–1060. doi:10.1037/0033-295X.111.4.1036
- Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science*, *2*, 396–408. doi:10.1111/j.1467-9280.1991.tb00174.x
- Audley, R. J. (1973). Some observations on theories of choice reaction time: Tutorial review. In S. Kornblum (Ed.), *Attention and performance IV* (pp. 509–545). New York, NY: Academic Press.
- Bertelson, P. (1961). Sequential redundancy and speed in a serial two-choice responding task. *Quarterly Journal of Experimental Psychology*, *13*, 90–102. doi:10.1080/17470216108416478
- Bertelson, P. (1965, April 10). Serial choice reaction-time as a function of response versus signal-and-response repetition. *Nature*, *206*, 217–218. doi:10.1038/206217a0
- Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, *57*, 153–178. doi:10.1016/j.cogpsych.2007.12.002
- Brown, S. D., Marley, A. A. J., Donkin, C., & Heathcote, A. (2008). An integrated model of choices and response times in absolute identification. *Psychological Review*, *115*, 396–425. doi:10.1037/0033-295X.115.2.396
- Chater, N., & Vitányi, P. (2003). Simplicity: A unifying principle in cognitive science? *Trends in Cognitive Sciences*, *7*, 19–22. doi:10.1016/S1364-6613(02)00005-0
- Cho, R. Y., Nystrom, L. E., Brown, E. T., Jones, A. D., Braver, T. S., Holmes, P. J., & Cohen, J. D. (2002). Mechanisms underlying dependencies of performance on stimulus history in a two-alternative forced-choice task. *Cognitive, Affective, & Behavioral Neuroscience*, *2*, 283–299. doi:10.3758/CABN.2.4.283
- Cleeremans, A., & McClelland, J. L. (1991). Learning the structure of event sequences. *Journal of Experimental Psychology: General*, *120*, 235–253. doi:10.1037/0096-3445.120.3.235
- Coles, M. G. H. (1989). Modern mind-brain reading: Psychophysiology, physiology, and cognition. *Psychophysiology*, *26*, 251–269. doi:10.1111/j.1469-8986.1989.tb01916.x
- Collins, A. G., & Frank, M. J. (2013). Cognitive control over learning: Creating, clustering, and generalizing task-set structure. *Psychological Review*, *120*, 190–229. doi:10.1037/a0030852
- Cuthill, I. C., Kacelnik, A., Krebs, J. R., Haccou, P., & Iwasa, Y. (1990). Starlings exploiting patches: The effect of recent experience on foraging decisions. *Animal Behaviour*, *40*, 625–640. doi:10.1016/S0003-3472(05)80692-X
- DeCarlo, L. T., & Cross, D. V. (1990). Sequential effects in magnitude scaling: Models and theory. *Journal of Experimental Psychology: General*, *119*, 375–396. doi:10.1037/0096-3445.119.4.375
- de Jong, R., Wierda, M., Mulder, G., & Mulder, L. J. M. (1988). Use of partial stimulus information in response processing. *Journal of Experimental Psychology: Human Perception and Performance*, *14*, 682–692. doi:10.1037/0096-1523.14.4.682
- Dien, J. (1998). Issues in the application of the average reference: Review, critiques, and recommendations. *Behavior Research Methods, Instruments & Computers*, *30*, 34–43. doi:10.3758/BF03209414
- Engler, J. (1958). Marginal and conditional stimulus and response probabilities in verbal conditioning. *Journal of Experimental Psychology*, *55*, 303–317. doi:10.1037/h0040170
- Estes, W. K. (1950). Toward a statistical theory of learning. *Psychological Review*, *57*, 94–107. doi:10.1037/h0058559
- Estes, W. K. (1957). Theory of learning with constant, variable, or contingent probabilities of reinforcement. *Psychometrika*, *22*, 113–132. doi:10.1007/BF02289048
- Feldman, J. (2000, October 5). Minimization of Boolean complexity in human concept learning. *Nature*, *407*, 630–633. doi:10.1038/35036586
- Filimon, F., Philiastrides, M. G., Nelson, J. D., Kloosterman, N. A., & Heekeren, H. R. (2013). How embodied is perceptual decision making? Evidence for separate processing of perceptual and motor decisions. *Journal of Neuroscience*, *33*, 2121–2136. doi:10.1523/JNEUROSCI.2334-12.2013
- Fletcher, B., & Rabbitt, P. M. A. (1978). The changing pattern of perceptual analytic strategies and response selection with practice in a 2-choice reaction-time task. *Quarterly Journal of Experimental Psychology*, *30*, 417–427. doi:10.1080/00335557843000025
- Flood, M. M. (1954). Environmental non-stationarity in a sequential decision-making experiment. In R. M. Thrall, C. H. Coombs, & R. L. Davis (Eds.), *Decision processes* (pp. 287–299). New York, NY: Wiley.
- Gambino, B., & Myers, J. L. (1967). Role of event runs in probability learning. *Psychological Review*, *74*, 410–419. doi:10.1037/h0024880
- Gao, J., Wong-Lin, K. F., Holmes, P., Simen, P., & Cohen, J. D. (2009). Sequential effects in two-choice reaction time tasks: Decomposition and synthesis of mechanisms. *Neural Computation*, *21*, 2407–2436. doi:10.1162/neco.2009.09-08-866
- Garner, W. R. (1953). An informational analysis of absolute judgments of loudness. *Journal of Experimental Psychology*, *46*, 373–380. doi:10.1037/h0063212
- Gluck, M. A., & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, *117*, 227–247. doi:10.1037/0096-3445.117.3.227
- Gold, J. I., & Shadlen, M. N. (2000, March 23). Representation of a perceptual decision in developing oculomotor commands. *Nature*, *404*, 390–394. doi:10.1038/35006062
- Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, *30*, 535–574. doi:10.1146/annurev.neuro.29.051605.113038
- Gratton, G., Coles, M. G., & Donchin, E. (1983). A new method for off-line removal of ocular artifact. *Electroencephalography & Clinical Neurophysiology*, *55*, 468–484. doi:10.1016/0013-4694(83)90135-9
- Hebb, D. O. (1949). *The organization of behavior*. New York, NY: Wiley.
- Hogarth, R. M., & Einhorn, H. J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology*, *24*, 1–55. doi:10.1016/0010-0285(92)90002-J
- Holland, M. K., & Lockhead, G. R. (1968). Sequential effects in absolute judgments of loudness. *Perception & Psychophysics*, *3*, 409–414. doi:10.3758/BF03205747
- Jentzsch, I., & Sommer, W. (2002). Functional localization and mechanisms of sequential effects in serial reaction time tasks. *Perception & Psychophysics*, *64*, 1169–1188. doi:10.3758/BF03194765
- Jesteadt, W., Luce, R. D., & Green, D. M. (1977). Sequential effects in judgments of loudness. *Journal of Experimental Psychology: Human Perception and Performance*, *3*, 92–104. doi:10.1037/0096-1523.3.1.92
- Jones, A. D., Cho, R. Y., Nystrom, L. E., Cohen, J. D., & Braver, T. S. (2002). A computational model of anterior cingulate function in speeded response tasks: Effects of frequency, sequence, and conflict. *Cognitive, Affective, & Behavioral Neuroscience*, *2*, 300–317. doi:10.3758/CABN.2.4.300
- Jones, M. (2009). A reinforcement-and-generalization model of sequential

- effects in identification learning. In N. Taatgen & H. van Rijn (Eds.), *Proceedings of the 31st Annual Meeting of the Cognitive Science Society* (pp. 1180–1185). Austin, TX: Cognitive Science Society.
- Jones, M., & Dzhaferov, E. N. (in press). Unfalsifiability and mutual translatability of major modeling schemes for choice reaction time. *Psychological Review*.
- Jones, M., Love, B. C., & Maddox, W. T. (2006). Recency effects as a window to generalization: Separating decisional and perceptual sequential effects in category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*, 316–332. doi:10.1037/0278-7393.32.3.316
- Jones, M., Maddox, W. T., & Love, B. C. (2005). Stimulus generalization in category learning. In B. G. Bara, L. Barsalou, & M. Bucciarelli (Eds.), *Proceedings of the 27th Annual Meeting of the Cognitive Science Society* (pp. 1066–1071). Alpha, NJ: Sheridan.
- Jones, M., Mozer, M. C., & Kinoshita, S. (2009). Optimal response initiation: Why recent experience matters. In D. Koller, D. Schuurmans, Y. Bengio, & L. Bottou (Eds.), *Advances in neural information processing systems* (pp. 785–792). Retrieved from http://matt.colorado.edu/jones-mozer-kinoshita_nips08.pdf
- Jones, M., & Sieck, W. R. (2003). Learning myopia: An adaptive recency effect in category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*, 626–640. doi:10.1037/0278-7393.29.4.626
- Junghöfer, M., Elbert, T., Tucker, D. M., & Braun, C. (1999). The polar average reference effect: A bias in estimating the head surface integral in EEG recording. *Clinical Neurophysiology*, *110*, 1149–1155. doi:10.1016/S1388-2457(99)00044-9
- Kamin, L. J. (1968). "Attention-like" processes in classical conditioning. In M. R. Jones (Ed.), *Miami Symposium on the Prediction of Behavior: Aversive Stimulation* (pp. 9–31). Coral Gables, FL: University of Miami Press.
- Kirby, N. H. (1976). Sequential effects in two-choice reaction time: Automatic facilitation or subjective expectancy? *Journal of Experimental Psychology: Human Perception and Performance*, *2*, 567–577. doi:10.1037/0096-1523.2.4.567
- Laming, D. R. J. (1968). *Information theory of choice-reaction times*. Oxford, England: Academic Press.
- Lockhead, G. R. (2004). Absolute judgments are relative: A reinterpretation of some psychophysical ideas. *Review of General Psychology*, *8*, 265–272. doi:10.1037/1089-2680.8.4.265
- Luce, R. D., & Green, D. M. (1974). The response ratio hypothesis for magnitude estimation. *Journal of Mathematical Psychology*, *11*, 1–14. doi:10.1016/0022-2496(74)90008-X
- Maloney, L. T., Dal Martello, M. F., Sahn, C., & Spillmann, L. (2005). Past trials influence perception of ambiguous motion quartets through pattern completion. *Proceedings of the National Academy of Sciences, USA*, *102*, 3164–3169. doi:10.1073/pnas.0407157102
- Mangun, G. R., & Hillyard, S. A. (1991). Modulations of sensory-evoked brain potentials indicate changes in perceptual processing during visual-spatial priming. *Journal of Experimental Psychology: Human Perception and Performance*, *17*, 1057–1074. doi:10.1037/0096-1523.17.4.1057
- Masson, M. E. J., & Loftus, G. R. (2003). Using confidence intervals for graphically based data interpretation. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale*, *57*, 203–220. doi:10.1037/h0087426
- McClelland, J. L. (1979). On the time relations of mental processes: An examination of systems of processes in cascade. *Psychological Review*, *86*, 287–330. doi:10.1037/0033-295X.86.4.287
- Medin, D. L., Goldstone, R. L., & Gentner, D. (1993). Respects for similarity. *Psychological Review*, *100*, 254–278. doi:10.1037/0033-295X.100.2.254
- Miller, J., Patterson, T., & Ulrich, R. (1998). Jackknife-based method for measuring LRP onset latency differences. *Psychophysiology*, *35*, 99–115. doi:10.1111/1469-8986.3510099
- Miller, J., Ulrich, R., & Schwarz, W. (2009). Why jackknifing yields good latency estimates. *Psychophysiology*, *46*, 300–312. doi:10.1111/j.1469-8986.2008.00761.x
- Mozer, M. C., Pashler, H., Cepeda, N., Lindsey, R., & Vul, E. (2009). Predicting the optimal spacing of study: A multiscale context model of memory. In Y. Bengio, D. Schuurmans, J. Lafferty, C. K. I. Williams, & A. Culotta (Eds.), *Advances in neural information processing systems 22* (pp. 1321–1329). La Jolla, CA: NIPS Foundation.
- Myers, J. L. (1976). Probability learning. In W. K. Estes (Ed.), *Handbook of learning and cognitive processes: Vol. 3. Approaches to human learning and motivation* (pp. 171–205). Hillsdale, NJ: Erlbaum.
- Nissen, M. J., & Bullemer, P. (1987). Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology*, *19*, 1–32. doi:10.1016/0010-0285(87)90002-8
- Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, *115*, 39–57. doi:10.1037/0096-3445.115.1.39
- Pashler, H., & Baylis, G. C. (1991). Procedural learning: 2. Intertrial repetition effects in speeded choice tasks. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*, 33–48. doi:10.1037/0278-7393.17.1.33
- Petzold, P. (1981). Distance effects on sequential dependencies in categorical judgments. *Journal of Experimental Psychology: Human Perception and Performance*, *7*, 1371–1385. doi:10.1037/0096-1523.7.6.1371
- Posner, M. I., & Cohen, Y. (1984). Components of visual orienting. In H. Bouma & D. Bouwhuis (Eds.), *Attention and performance X* (pp. 531–556). Hillsdale, NJ: Erlbaum.
- Rabbitt, P. M. A., & Rodgers, B. (1977). What does a man do after he makes an error? An analysis of response programming. *The Quarterly Journal of Experimental Psychology*, *29*, 727–743. doi:10.1080/14640747708400645
- Ratcliff, R., & Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Review*, *111*, 333–367. doi:10.1037/0033-295X.111.2.333
- Real, L. A. (1991, August 30). Animal choice behavior and the evolution of cognitive architecture. *Science*, *253*, 980–986. doi:10.1126/science.1887231
- Reber, A. S. (1967). Implicit learning of artificial grammars. *Journal of Verbal Learning and Verbal Behavior*, *6*, 855–863. doi:10.1016/S0022-5371(67)80149-X
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64–99). New York, NY: Appleton-Century-Crofts.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986, October 9). Learning representations by back-propagating errors. *Nature*, *323*, 533–536. doi:10.1038/323533a0
- Sakamoto, Y., Jones, M., & Love, B. C. (2008). Putting the psychology back into psychological models: Mechanistic versus rational approaches. *Memory & Cognition*, *36*, 1057–1065. doi:10.3758/MC.36.6.1057
- Schultz, W., Dayan, P., & Montague, P. R. (1997, March 14). A neural substrate of prediction and reward. *Science*, *275*, 1593–1599. doi:10.1126/science.275.5306.1593
- Shepard, R. N. (1957). Stimulus and response generalization: A stochastic model relating generalization to distance in psychological space. *Psychometrika*, *22*, 325–345. doi:10.1007/BF02288967
- Shepard, R. N. (1987, September 11). Toward a universal law of generalization for psychological science. *Science*, *237*, 1317–1323. doi:10.1126/science.3629243

- Soetens, E. (1998). Localizing sequential effects in serial choice reaction time with the information reduction procedure. *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 547–568. doi:10.1037/0096-1523.24.2.547
- Soetens, E., Boer, L. C., & Hueting, J. E. (1985). Expectancy or automatic facilitation? Separating sequential effects in two-choice reaction time. *Journal of Experimental Psychology: Human Perception and Performance*, *11*, 598–616. doi:10.1037/0096-1523.11.5.598
- Spehlmann, R. (1965). The average electrical responses to diffuse and to patterned light in the human. *Electroencephalography & Clinical Neurophysiology*, *19*, 560–569. doi:10.1016/0013-4694(65)90241-5
- Squires, K. C., Wickens, C., Squires, N. K., & Donchin, E. (1976, September 17). The effect of stimulus sequence on the waveform of the cortical event-related potential. *Science*, *193*, 1142–1146. doi:10.1126/science.959831
- Srinivasan, R., Nunez, P. L., Silberstein, R. B., Tucker, D. M., & Cadusch, P. J. (1996). Spatial sampling and filtering of EEG with spline-Laplacians to estimate cortical potentials. *Brain Topography*, *8*, 355–366. doi:10.1007/BF01186911
- Starns, J. J., Ratcliff, R., & McKoon, G. (2012). Evaluating the unequal-variability and dual-process explanations of zROC slopes with response time data and the diffusion model. *Cognitive Psychology*, *64*, 1–34. doi:10.1016/j.cogpsych.2011.10.002
- Stewart, N., Brown, G. D. A., & Chater, N. (2002). Sequence effects in categorization of simple perceptual stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *28*, 3–11. doi:10.1037/0278-7393.28.1.3
- Stewart, N., Brown, G. D. A., & Chater, N. (2005). Absolute identification by relative judgment. *Psychological Review*, *112*, 881–911. doi:10.1037/0033-295X.112.4.881
- Sutherland, N., & Mackintosh, N. (1971). *Mechanisms of animal discrimination learning*. New York, NY: Academic Press.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Treisman, M. C., & Williams, T. C. (1984). A theory of criterion setting with an application to sequential dependencies. *Psychological Review*, *91*, 68–111. doi:10.1037/0033-295X.91.1.68
- Ulrich, R., & Miller, J. (2001). Using the jackknife-based scoring method for measuring LRP onset effects in factorial designs. *Psychophysiology*, *38*, 816–827. doi:10.1111/1469-8986.3850816
- Usher, M., & McClelland, J. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, *108*, 550–592. doi:10.1037/0033-295X.108.3.550
- Vervaeck, K., & Boer, L. (1980). Sequential effects in two-choice reaction time: Subjective expectancy and automatic after-effects at short response-stimulus intervals. *Acta Psychologica*, *44*, 175–190. doi:10.1016/0001-6918(80)90066-9
- Ward, L. M., & Lockhead, G. R. (1971). Response system processes in absolute judgment. *Perception & Psychophysics*, *9*, 73–78. doi:10.3758/BF03213031
- Wickelgren, W. A. (1974). Single-trace fragility theory of memory dynamics. *Memory & Cognition*, *2*, 775–780. doi:10.3758/BF03198154
- Wilder, M. H., Jones, M., Ahmed, A. A., Curran, T., & Mozer, M. C. (2013, February 21). The persistent impact of incidental experience. *Psychonomic Bulletin & Review*. Advance online publication. doi:10.3758/s13423-013-0406-3
- Wilder, M., Jones, M., & Mozer, M. C. (2010). Sequential effects reflect parallel learning of multiple environmental regularities. In Y. Bengio, D. Schuurmans, J. Lafferty, C. K. I. Williams, & A. Culotta (Eds.), *Advances in neural information processing systems 23* (pp. 2053–2061). La Jolla, CA: NIPS Foundation.
- Willingham, D. B., Wells, L. A., Farrell, J. M., & Stemwedel, M. E. (2000). Implicit motor sequence learning is represented in response locations. *Memory & Cognition*, *28*, 366–375. doi:10.3758/BF03198552
- Yu, A., & Cohen, J. (2009). Sequential effects: Superstition or rational behavior? In Y. Bengio, D. Schuurmans, J. Lafferty, C. K. I. Williams, & A. Culotta (Eds.), *Advances in neural information processing systems 22* (pp. 1873–1880). La Jolla, CA: NIPS Foundation.

(Appendices follow)

Appendix A

Derivation of Parallel-Learning Model Predictions for Experiment 1

To derive predictions of the parallel-learning model for Experiment 1, we start by using Equation 10 to predict the mean response time (RT) for each 5-deep trial history in each condition. Let h be any 5-deep history. Specifying h means that $M_1(n-1, n), \dots, M_1(n-5, n)$ and $M_2(n-1, n), \dots, M_2(n-4, n)$ are fixed, but $M_1(n-k, n)$ for $k \geq 6$ and $M_2(n-k, n)$ for $k \geq 5$ are variable. Therefore, to evaluate Equation 10, we need the expected values of these higher-lag matches conditioned on h .

In the positive condition, trials $n-6$ and $n-5$ will be a first-degree match $\frac{2}{3}$ of the time. Because matches are coded as 1 and mismatches as -1 , the expected value of $M_1(n-6, n-5)$ equals $\frac{1}{3}$. More generally,

$$E[M_1(n-k, n-5)] = \frac{1}{3^{k-5}} \text{ for } k \geq 6. \quad (\text{A1})$$

In the negative condition, the expected value of $M_1(n-6, n-5)$ equals $-\frac{1}{3}$, and more generally

$$E[M_1(n-k, n-5)] = \frac{1}{(-3)^{k-5}} \text{ for } k \geq 6. \quad (\text{A2})$$

Because $M_1(n-k, n) = M_1(n-5, n) \cdot M_1(n-k, n-5)$, the expected value of the first-degree match between trials $n-k$ and n following any history for $k \geq 6$ is given by

$$E[M_1(n-k, n) | h] = \begin{cases} \frac{M_1(n-5, n)}{3^{k-5}} & \text{Positive condition} \\ \frac{M_1(n-5, n)}{(-3)^{k-5}} & \text{Negative condition.} \end{cases} \quad (\text{A3})$$

The calculation for second-degree matches is easier. In the positive condition, trial $n-k$ is a repetition $\frac{2}{3}$ of the time. If trial n is a repetition, then $M_2(n-k, n)$ will equal 1 with probability $\frac{2}{3}$, and if trial n is an alternation, then the probability is $\frac{1}{3}$. These values are reversed in the negative condition. Therefore, the expected value of the second-degree match between trials $n-k$ and n for $k \geq 5$ following any history is given by

$$E[M_2(n-k, n) | h] = \begin{cases} \frac{M_1(n-1, n)}{3} & \text{Positive condition} \\ -\frac{M_1(n-1, n)}{3} & \text{Negative condition.} \end{cases} \quad (\text{A4})$$

Equations A3 and A4 enable calculation of the average contributions of first- and second-degree REs from trials earlier than the specified history length. For first-degree REs, this contribution comes from all summands with $k \geq 6$ from the first-degree component of the model (see Equation 8), given by

$$E\left[-\beta_{\text{base}} \sum_{k \geq 6} \epsilon_{\text{base}} (1 - \epsilon_{\text{base}})^{k-1} M_1(n-k, n) | h\right] = \begin{cases} -\frac{\beta_{\text{base}} \epsilon_{\text{base}} (1 - \epsilon_{\text{base}})^5 M_1(n-5, n)}{2 + \epsilon_{\text{base}}} & \text{Positive condition} \\ \frac{\beta_{\text{base}} \epsilon_{\text{base}} (1 - \epsilon_{\text{base}})^5 M_1(n-5, n)}{4 - \epsilon_{\text{base}}} & \text{Negative condition.} \end{cases} \quad (\text{A5})$$

(Appendices continue)

For second-degree REs, the contribution from earlier trials comes from all summands with $k \geq 5$ from the second-degree component of the model, given by

$$E\left[-\beta_{\text{rep}} \sum_{k \geq 5} \epsilon_{\text{rep}}(1 - \epsilon_{\text{rep}})^{k-1} M_2(n - k, n) \mid h\right] = \begin{cases} -\frac{\beta_{\text{rep}}(1 - \epsilon_{\text{rep}})^4 M_1(n - 1, n)}{3} & \text{Positive condition} \\ \frac{\beta_{\text{rep}}(1 - \epsilon_{\text{rep}})^4 M_1(n - 1, n)}{3} & \text{Negative condition.} \end{cases} \quad (\text{A6})$$

In Equations A5 and A6, we make the simplifying assumption that the sums in Equation 8 continue to infinity instead of $n - 1$. This amounts to assuming that $(1 - \epsilon_{\text{base}})^n$ and $(1 - \epsilon_{\text{rep}})^n$ are negligibly small.

With the average contribution of earlier trials worked out, we can now evaluate Equation 10. In the positive condition, mean RT following any 5-deep history h is given by

$$E[RT \mid h] = \beta_0 - \beta_{\text{base}} \left[\sum_{k=1}^5 \epsilon_{\text{base}}(1 - \epsilon_{\text{base}})^{k-1} M_1(n - k, n) + \frac{\epsilon_{\text{base}}(1 - \epsilon_{\text{base}})^5}{2 + \epsilon_{\text{base}}} M_1(n - 5, n) \right] - \beta_{\text{rep}} \left[\sum_{k=1}^4 \epsilon_{\text{rep}}(1 - \epsilon_{\text{rep}})^{k-1} M_2(n - k, n) + \frac{(1 - \epsilon_{\text{rep}})^4}{3} M_1(n - 1, n) \right] + \delta E_n E_{n-1}. \quad (\text{A7})$$

In the negative condition, mean RT following any 5-deep history h is given by

$$E[RT \mid h] = \beta_0 - \beta_{\text{base}} \left[\sum_{k=1}^5 \epsilon_{\text{base}}(1 - \epsilon_{\text{base}})^{k-1} M_1(n - k, n) - \frac{\epsilon_{\text{base}}(1 - \epsilon_{\text{base}})^5}{4 - \epsilon_{\text{base}}} M_1(n - 5, n) \right] - \beta_{\text{rep}} \left[\sum_{k=1}^4 \epsilon_{\text{rep}}(1 - \epsilon_{\text{rep}})^{k-1} M_2(n - k, n) - \frac{(1 - \epsilon_{\text{rep}})^4}{3} M_1(n - 1, n) \right] + \delta E_n E_{n-1}. \quad (\text{A8})$$

Finally, the predictions of pairs of 5-deep histories are averaged to get predictions for 4-deep histories, paralleling the analysis of the data. Of the two 5-deep histories corresponding to any 4-deep history, one has $M_1(n - 5, n) = 1$, and the other has $M_1(n - 5, n) = -1$; also one has $M_2(n - 4, n) = 1$, and the other has $M_2(n - 4, n) = -1$ (it is inconsequential which is which). Averaging their predictions together yields

$$E[RT \mid h] = \beta_0 - \beta_{\text{base}} \sum_{k=1}^4 \epsilon_{\text{base}}(1 - \epsilon_{\text{base}})^{k-1} M_1(n - k, n) - \beta_{\text{rep}} \sum_{k=1}^3 \epsilon_{\text{rep}}(1 - \epsilon_{\text{rep}})^{k-1} M_2(n - k, n) + \left[\delta - \frac{\beta_{\text{rep}}(1 - \epsilon_{\text{rep}})^4}{3} \right] M_1(n - 1, n) \quad (\text{A9})$$

for the positive condition and

$$E[RT \mid h] = \beta_0 - \beta_{\text{base}} \sum_{k=1}^4 \epsilon_{\text{base}}(1 - \epsilon_{\text{base}})^{k-1} M_1(n - k, n) - \beta_{\text{rep}} \sum_{k=1}^3 \epsilon_{\text{rep}}(1 - \epsilon_{\text{rep}})^{k-1} M_2(n - k, n) + \left[\delta + \frac{\beta_{\text{rep}}(1 - \epsilon_{\text{rep}})^4}{3} \right] M_1(n - 1, n) \quad (\text{A10})$$

for the negative condition.

Equations A9 and A10 were used to fit the parallel-learning model to the RT data of Experiment 1 (see Figure 8). These equations show how the method of computing mean RTs for 5-deep histories before averaging to 4-deep histories cancels out the first-degree effects of earlier trials, because the expected value of these earlier effects is proportional to $M_1(n - 5, n)$. The second-degree effects of earlier trials do not cancel out but instead accumulate as an expectancy for repetition ($M_1(n - 1, n) = 1$) in the positive condition or alternation ($M_1(n - 1, n) = -1$) in the negative condition. This long-term learning effect combines with the built-in alternation advantage (δ). The model's predictions for the two conditions are thus identical in all ways except for the alternation advantage, which is predicted to be greater in the negative condition by an amount equal to

$$\frac{4\beta_{\text{rep}}(1 - \epsilon_{\text{rep}})^4}{3}. \quad (\text{A11})$$

(Appendices continue)

Appendix B

Parallel-Learning Model Predictions for Lag Profiles

The predictions of the parallel-learning model for the lag profiles can be derived from its predicted mean response times (RTs) for the 4-deep trial histories, as given by Equations A9 and A10 in Appendix A. Each measure $RE_i(k)$ equals the average RT prediction for histories with $M_i(n-k, n) = -1$ minus the average prediction for histories with $M_i(n-k, n) = 1$ (for $i = 1$ or 2). It is easily verified that, with the exception of $M_1(n-2, n)$ and $M_2(n-1, n)$, the various $M_i(n-k, n)$ are all orthogonal over the sixteen 4-deep histories. That is, if $(i, k) \neq (i', k')$, the eight histories with $M_i(n-k, n) = 1$ are split evenly between $M_{i'}(n-k', n) = 1$ and $M_{i'}(n-k', n) = -1$, as are the eight histories with $M_i(n-k, n) = -1$. Therefore, because Equations A9 and A10 are linear in all M terms, all except $M_i(n-k, n)$ will cancel out when calculating $RE_i(k)$.

This reasoning leads to the following expressions for the model's predictions. In the positive condition, $RE_1(1)$ is found by collecting all $M_1(n-1, n)$ terms in Equation A9 and subtracting their values when $M_1(n-1, n) = 1$ from their values when $M_1(n-1, n) = -1$:

$$\begin{aligned} RE_1(1) &= \left(-\beta_{\text{base}}\epsilon_{\text{base}}(-1) + \left[\delta - \frac{\beta_{\text{rep}}(1 - \epsilon_{\text{rep}})^4}{3} \right](-1) \right) \\ &\quad - \left(-\beta_{\text{base}}\epsilon_{\text{base}} \cdot 1 + \left[\delta - \frac{\beta_{\text{rep}}(1 - \epsilon_{\text{rep}})^4}{3} \right] \cdot 1 \right) \\ &= 2\beta_{\text{base}}\epsilon_{\text{base}} + \frac{2\beta_{\text{rep}}(1 - \epsilon_{\text{rep}})^4}{3} - 2\delta. \end{aligned} \quad (\text{B1})$$

Similarly, in the negative condition,

$$\begin{aligned} RE_1(1) &= \left(-\beta_{\text{base}}\epsilon_{\text{base}}(-1) + \left[\delta + \frac{\beta_{\text{rep}}(1 - \epsilon_{\text{rep}})^4}{3} \right](-1) \right) \\ &\quad - \left(-\beta_{\text{base}}\epsilon_{\text{base}} \cdot 1 + \left[\delta + \frac{\beta_{\text{rep}}(1 - \epsilon_{\text{rep}})^4}{3} \right] \cdot 1 \right) \\ &= 2\beta_{\text{base}}\epsilon_{\text{base}} - \frac{2\beta_{\text{rep}}(1 - \epsilon_{\text{rep}})^4}{3} - 2\delta. \end{aligned} \quad (\text{B2})$$

The other measures are calculated similarly, with results shown in Table 3.

Finally, because $M_1(n-2, n) = M_2(n-1, n)$, $RE_1(2)$ and $RE_2(1)$ will be based on both terms. The result is

$$\begin{aligned} RE_1(2) = RE_2(1) &= (-\beta_{\text{base}}\epsilon_{\text{base}}(1 - \epsilon_{\text{base}})(-1) - \beta_{\text{rep}}\epsilon_{\text{rep}}(-1)) \\ &\quad - (-\beta_{\text{base}}\epsilon_{\text{base}}(1 - \epsilon_{\text{base}}) \cdot 1 - \beta_{\text{rep}}\epsilon_{\text{rep}} \cdot 1) \\ &= 2\beta_{\text{base}}\epsilon_{\text{base}}(1 - \epsilon_{\text{base}}) + 2\beta_{\text{rep}}\epsilon_{\text{rep}}. \end{aligned} \quad (\text{B3})$$

Received February 19, 2013

Revision received April 29, 2013

Accepted April 29, 2013 ■