In Press: PSYCHONOMIC BULLETIN AND REVIEW, 2013.

The Persistent Impact of Incidental Experience

Matthew H. Wilder*, Matt Jones^, Alaa Ahmed#, Tim Curran^, & Michael C. Mozer*

*Dept. of Computer Science,

^Dept. of Psychology and Neuroscience,

#Dept. of Integrative Physiology;

University of Colorado, Boulder, CO 80309

<wildermh@colorado.edu, mcj@colorado.edu, alaa@colorado.edu,

Tim.Curran@colorado.edu, mozer@colorado.edu>

Abstract

As we perform daily activities—driving to work, unlocking the office door, grabbing a coffee cup—our actions seem automatic and preprogrammed. Nonetheless, routine, well-practiced behavior is continually modulated by incidental experience: in repetitive experimental tasks, recent (~4) trials reliably influence performance and action choice. Psychological theories downplay the significance of sequential effects, explaining them as rapidly decaying perturbations of behavior with no long-term consequences. We challenge this traditional perspective in two studies designed to probe the impact of more distant experience, finding evidence for effects spanning up to a thousand intermediate trials. We present a normative theory in which these persistent effects reflect optimal adaptation to a dynamic environment exhibiting varying rates of change. The theory predicts a heavy-tailed decaying influence of past experience, consistent with our data, and suggests that individual incidental experiences are catalogued in a temporally extended memory utilized to optimize subsequent behavior.

Introduction

Throughout our daily lives, we encounter an ongoing barrage of mundane stimuli that demand routine responses. This *incidental* experience forms the fabric of our interaction with the world. Clearly, the sum of this experience determines our behavior, but how long-lasting is the effect of each experience on subsequent behavior?

The effects of recent experience on decision making have been studied via a Two-Alternative Forced-Choice (2AFC) paradigm. On each trial, one of two stimuli is presented. Subjects are asked to press one of two response keys as quickly as possible. Response time (RT) varies reliably as a function of the exact sequence of preceding trials as in Figure 1a (e.g., Cho et al., 2002; Jentzsch & Sommer, 2002; Remington, 1969; Soetens, Boer, & Hueting, 1985). These *sequential dependencies* are not a mere laboratory curiosity but can have a meaningful impact on naturalistic decision making. For example, professional basketball players choose their shot locations based on recent attempts and successes (Meiman & Loewenstein, 2011). Recent braking or acceleration actions of automobile drivers can explain variability in response latencies of up to 100 ms, potentially the difference between a collision and a near miss (Doshi, Tran, Wilder, Mozer, & Trivedi, 2012). Sequential dependencies have also been demonstrated in legal reasoning and jury evidence interpretation (e.g., Furnham, 1986; Hogarth & Einhorn, 1992) and clinical assessments (Mumma & Wilson, 2006).

Sequential dependencies arise naturally from psychological and neurobiological models of incremental learning, including error correction methods (Rescorla & Wagner, 1972), reinforcement learning (Sutton & Barto, 1998) and Hebbian learning (Hebb, 1949). These models yield an exponentially discounted influence of past trials, which explains the inverted-V pattern common to many 2AFC experiments (as in Figure 1a). Similarly, models from optimal control theory for tracking nonstationary environments, such as the Kalman filter (Kalman, 1960), also produce exponential decay. These models are all appealing because the past trial history is captured by a single state variable (or sufficient statistic) that can be maintained and updated from trial to trial.

Models that produce exponential decay of past trials predict sequential dependencies to operate only on short timescales. Moreover, analyses of sequential dependencies have focused on the short timescale, and the design of experiments has not been well suited to measuring longer-range effects. However, several studies hint at the possibility that a single experience can have an influence on behavior persisting minutes (e.g., Link, Kow, Wager, & Mozer, 2011; Wong & Shelhamer, 2011) or even a day (Ward & Lockheed, 1970), consistent with an alternative theoretical perspective in which each experience is stored in long-term memory, and behavior is guided by the cumulative impact of these memories (e.g., Kasif, Salzberg, Waltz, Rachlin, & Aha, 1998; Stanfill & Waltz, 1986).

Instead of an exponential discounting of the past, long-term memory is typically characterized as following a "power law of forgetting" (Anderson et al., 2004; Rubin & Wenzel, 1996; Wixted & Carpenter, 2007; Wixted & Ebbesen, 1997). Power functions are qualitatively different from exponential functions because they can produce a single curve that exhibits *both* rapid decay of the most recent trials (a strong *short-term recency effect*) and slow decay of far-back trials (a long-range residual effect). With exponential decay, long-term effects are vanishingly small, at least with decay rates in the range needed to explain short-term recency.

Our investigation explored the persistence of incidental experience, both in terms of the scope of its influence and the nature of its decay. We began by reanalyzing trial-by-trial data from a typical 2AFC experiment (Jentzsch & Sommer, 2002). We compared two models of sequential effects that assume that subjects form an expectation for the next trial using an average of previous trials that is weighted either exponentially or according to a power function. RT was predicted to be fast when the expectation matched the actual trial and slow when it did not. Throughout the article, each model was fit to the *specific* trial history of *individual* subjects by minimizing the mean squared error across all trials. Both models had a single theoretically relevant free parameter for determining the relative weighting of past trials.

The analysis used in previous investigations, in which RTs are conditioned on the four-back sequence (Figure 1a) , does not gauge the persistence of experience or facilitate discrimination of the two models. Thus, to examine the influence of past trials more closely, we studied how model fits vary as a function of the number of past trials used to form each expectation (the *context horizon*). Out-of-sample fits were obtained for each context horizon by iteratively computing a prediction for each trial using a model that was fit to the preceding trials and constrained to have the desired context horizon. All models have *one* free parameter regardless of horizon size. Accumulative prediction error (Wagenmakers, Grünwald, & Steyvers, 2006) is computed from the out-of-sample fits (Figure 1b). For an error measure, we use the coefficient of determination (R^2) which is derived from the sum of squared residuals error measure recommended in Wagenmakers et al. (2006). Increasing the horizon beyond four trials back yields reliable improvements in fit: across models that use 4 to 1024 past trials, there is a significant main effect of horizon on R^2 ($F(8,72)=3.28$, $p=.003$), but despite the appearance of a better fit for the power model, the interaction between horizon and model was not reliable ($F(8,72)=1.033$, $p=.42$).

The Jentzsch and Sommer (2002) study was limited because higher-order sequence statistics were not controlled—introducing an additional source of variability—and because distinguishing predictions of the two models is difficult when sequences have no structure. The latter point is due to the fact that, when the two trial types—repetition and alternation—occur with equal probability, their influence tends to cancel out, regardless of how strongly individual trials are weighted.

Experiment 1: Autocorrelation in the Sequence Structure

We therefore conducted a 2AFC study with a biased sequence structure in two opposing conditions, one in which 2/3 of the trials were repetitions of the preceding trial and one in which 2/3 of the trials were alternations of the preceding trial—*positive* and *negative* autocorrelation, respectively.

Method

Participants

Twenty-eight young adults (age 21.5 ± 2.9 yrs, 9 female, 19 male) participated for monetary compensation. 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. One subject was removed from the analysis because of an error recording responses during one block. Participants gave informed consent in accordance with the University of Colorado's Institutional Review Board.

Stimuli and Apparatus

The subject's task 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 visible throughout the task. Responses were made using a button box, oriented vertically so as to be spatially compatible with target locations. The left and right index fingers were assigned to the two buttons, with the assignment counterbalanced across subjects and fixed across sessions for each subject. Stimulus duration was 100 ms. A 700 ms response-to-stimulus interval followed each response. Reaction time was recorded at 1000 HZ.

Procedure

Each session consisted of 3402 experimental trials divided into 14 blocks of 243 trials. Within each block, local stimulus histories were controlled to a depth of six trials, and the frequency of each of the 64 (2^6) different trial sequences was exactly as dictated by the repetition rate for the condition (1/3 and 2/3 for the Negative and Positive conditions, respectively). The actual stimulus identities (above or below the fixation line) were equally probable. Subjects were given rest breaks roughly every 116 trials, and additional practice and post-rest contextual lead-in trials were inserted into the sequence for a total of 3744 trials.

Results and Discussion

As expected, RTs were modulated by the short-term context (Figure 2a). However, behavior also depended on the autocorrelation structure: RTs for repetition trials (left side of Figure 2a) were faster in the positive condition than the negative condition and vice versa for alternation trials. The difference due to autocorrelation structure when conditioned on the immediate context indicates that the influence of the past extends beyond four trials back. Although one cannot determine how far back from Figure 2a, a preference for the power model emerges when fits to the per-subject trial-by-trial data are aggregated according to the four-back sequence history. R^2 between model and data across the 32 histories (16 in each condition) was greater for the power model for 25 of the 27 subjects (mean R^2 across subjects: .798 vs. .730; paired t-test, $t(26)=7.45$, $p<.001$). R^2 values reported for the means of the four-back sequence histories are higher than those for the individual trial data—e.g., Figure 1a—because some sources of variability are averaged out.

Support for a long-range sequential effect is obtained by examining the accumulative prediction error values for the two models with varied context horizons (Figure 2b). There is a significant main effect of horizon ($F(8,208)=85.1$, $p<.001$) and an interaction between model type and horizon ($F(7,182)=62.3$, $p<.001$). The exponential model fit improves reliably as more trials are included out to 32 trials (comparing 32 vs. 16, $t(26)=4.12$, $p=.0003$) but no further (1024 vs. 32, $t(26)=0.36$, $p=.72$). In contrast, the power model fit improves out to 1024 trials (1024 vs. 512, $t(26)=2.84$, $p=.0086$). Behavior in this task is clearly affected by a long history of prior experience.

Further support for power over exponential decay is obtained by studying a *lag profile* derived from the data, plotted on a log-log scale in Figure 2c. The lag profile isolates the effect of the trial l trials in the past by computing the difference between mean RT when the current trial does not match the lag-l trial and mean RT when those trials do match. Because the exponential and power models both predict a lag profile that matches the decay function, this analysis offers another means of differentiating the models. The empirical lag profile appears linear in log-log coordinates suggesting power decay. We fit individual subject lag profiles to both power and exponential functions and obtained a better fit for the power function (mean R^2 across subjects: .878 vs. .855; $t(26)=2.17$, $p=.039$).

Even though both the power and exponential functions have a single free parameter, one could argue that the power function fits better because it has more flexibility. To rule out this possibility, we compare out-of-sample fits using leave-one-out cross validation. The power fit is consistently better than the exponential fit across lags and across subjects: mean absolute deviation between the empirical and predicted lag values is smaller for the power function ($F(1,26)=10.47$, $p=.003$). For 9 of the 10 lags, the mean absolute deviation is smaller for the power function. Furthermore, we compared fits for individual subjects using an extension of the likelihood ratio test that is appropriate for non-nested models (Vuong, 1989). Figure 2d histograms the log-likelihood ratios across subjects. A preference for the power model is evidenced by both the larger number of significantly negative ratios according to the Vuong test (11 blue vs. 3 red boxes) and the larger total number of negative ratios (18 vs. 9).

If incidental experience has a long-lasting influence, a cumulative effect of trial statistics across the entire course of the experiment might be observable. Figure 2e reveals a preference for repetitions in the positive condition that increases as the experiment progresses, and a preference for alternation in the negative condition. When superimposed over Figure 2e, predictions derived from power model fits capture the long-range effect of condition. In contrast, the trajectory from the exponential model fits is roughly flat because the model cannot benefit from integrating beyond about 64 past trials.

The power model is appealing because it is capable of explaining effects across a range of timescales, from the variation due to the immediate four-back context to the bias that grows over the hour-long duration of the sequence in each autocorrelation condition.

Experiment 2: Sequential Dependencies in Motor Control

Although we have argued for a unified explanation of short- and long-term adaptation via the power model, there is an alternative, though somewhat less parsimonious, possibility that the two timescales reflect distinct mechanisms. For instance, in Experiment 1, sequence structure might have been detected by subjects, leading to explicit learning and deliberate biasing of behavior. We thus aimed to strengthen our account by demonstrating the persistence of incidental experience in the absence of sequential structural regularity.

However, as our reanalysis of the Jentzsch and Sommer (2002) data revealed, it was difficult to uncover long-range effects when the sequence history was balanced and response latency was the dependent variable. We conjectured that response latency may not be a terribly sensitive measure because speedy responses are a secondary consideration in the performance of 2AFC; responding correctly is the subjects' primary goal. Consequently, RTs may be more susceptible to perturbation by task-unrelated factors. A task whose behavioral measures are better aligned with the subjects' primary goals might be more effective in exposing a persistent influence of incidental experience, despite the previously described cancellation of far-back effects that results from balanced sequences.

One domain of study that seems suitable is motor control because movement trajectories reflect planning processes. Long-term motor adaptation has been observed when systematic and consistent perturbations were applied to the control system (e.g., Hoppand & Fuchs, 2004; Robinson, Soetedjo, & Noto, 2006; Shadmehr & Mussa-Ivaldi, 1994). Some support for the persistent influence of incidental experience is found in an eye movement task in which error-based adaptation was observed extending back nearly one hundred trials and decaying according to a power function (Wong & Shelhamer, 2011). However, in this task, correlations could be attributed to endogenous variation rather than exogenous effects of the target sequence because target timing and position were completely predictable on every trial. Though ignored in many motor control studies, short-term sequential dependencies have been demonstrated in reaching tasks where straight-line arm movements were disrupted by variable perpendicular perturbation forces (Fine & Thoroughman, 2006; Scheidt, Dingwell, & Mussa-Ivaldi, 2001).

To bridge the gap between traditional 2AFC experiments and motor studies that exhibit sequential effects, we explored a reaching task with a sequential structure akin to that of 2AFC. Rather than imposing an autocorrelation structure as in Experiment 1, the two trial types were controlled to be equally probable.

Method

Participants

Twenty right-handed young adults (age 18.3 ± 0.7 yrs, 14 female, 6 male) participated for monetary compensation. Participants gave informed consent in accordance with the University of Colorado's Institutional Review Board.

Stimuli and Apparatus

Subjects sat in a chair with full back support and made horizontal planar reaching movements while grasping the handle of a robotic arm (Interactive Motion Technologies, Shoulder-Elbow Robot 2). The handle position, handle velocity, and robot-generated force were recorded at 20 Hz.

The task involved making rapid 15cm out-and-back movements along the midline of the transverse plane. Visual feedback of a cursor representing hand position and the home and target circles was presented on an LCD screen in front of the subjects (Figure 3a). Once subjects had centered the cursor within the home circle, the target appeared, and an audio cue signaled the trial onset. On each trial, a perturbing force was applied perpendicular to the desired direction of movement. The force increased linearly as a function of distance from the home circle over the first 5cm (1 N/cm) and remained fixed at 5N for the remaining 10cm. No forces were applied on the return. Subjects received warning messages if trial duration exceeded 1.4 seconds.

Procedure

Two versions of the task were run, identical except for the control of the stimulus sequences, with 10 subjects each. In version 1, 10 introductory null trials with no forces were followed by 490 force trials with the force direction randomly selected with equal probability. Subjects were given a 30 second break after every 100 trials. In version 2, subjects completed a total of 1106 trials with 10 introductory null trials and 30 second rests every 137 trials. The 9 trials following each rest were excluded from analyses. Local stimulus histories of right and left trials were controlled to a depth of 9 trials so that each of the 512 (2^9) trial sequences occurred exactly twice. For model fitting, the deflection measures for right and left trials were normalized—for each subject—to have the same mean and standard deviation, thus eliminating imbalances due to structural constraints of the arm. All statistical analyses focus on model fits to individual subjects and collapsed across data from the two versions of the experiment.

Results and Discussion

Individual trial movement trajectories were affected by the recent trial sequence: subjects compensated for the current perturbation more accurately when it was consistent with the recency-weighted sequence of prior perturbations (Figure 3b). For the purpose of modeling, accuracy of the trajectory on a given trial was quantified as the absolute value of the maximum horizontal deviation of the trajectory. However, other deflection measures—e.g., initial angle, mean deviation, area under deflection curve—gave similar results. The persistence of past experience is revealed by analyzing accumulative prediction error as a function of context horizon (Figure 3c). We find support for the hypothesis that sequential effects extend back more than 32 trials (one-tailed t-test for 64 vs. 32, exponential: $t(19)=1.93$, $p=.035$; power: $t(19)=1.86$, $p=.040$). Because the exponential and power models differ primarily in the weights they assign to far-back trials, we expected that the balanced sequences in this experiment would make it difficult to compare the two models directly. Despite this limitation, evidence for power decay over exponential decay is found in the near linear trend of the lag profile in log-log coordinates (Figure 3d). Per-subject fits to the lag profile values are reliably better for a power function than an exponential function (mean R^2 across subjects: .891 vs. .835; $t(19)=4.98$, $p<.001$).

Using leave-one-out cross validation, the power fit is significantly better than the exponential fit across lags and across subjects: mean absolute deviation between the empirical and predicted lag values is smaller for the power function ($F(1,19)=15.26$, $p=.001$). Additionally, for 9 of the 10 lags, the mean absolute deviation is smaller for the power function. Figure 3e shows a strong preference for the power model according to Vuong's test (Vuong, 1989) with more significantly negative log-likelihood ratios (12 blue vs. 0 red) and a larger total number of negative ratios (17 vs. 3).

A Normative Account of Long-Range Effects

Many theoretical accounts characterize sequential dependencies as a by-product of adaptation to the statistical structure of a dynamic environment (e.g., Jones & Sieck, 2003; Mozer, Kinoshita, & Shettel, 2007; Wilder, Jones, & Mozer, 2010; Yu & Cohen, 2009). These accounts suppose that statistics of the environment are tracked over time—statistics such as relative stimulus frequency or the magnitude and direction of perturbing forces. The statistics represent not only a summary of the past, but an expectation for the future, facilitating tuning of perceptuo-motor control to perform optimally in the anticipated environment.

If environments have temporal nonstationarity, more recent experience is most indicative of what an individual will experience next. Specific theoretical formulations lead to specific characterizations of how past experiences should optimally be combined to predict future events. Yu and Cohen's (2009) Dynamic Belief Model (DBM) explains sequential effects as a consequence of optimal Bayesian inference in an environment whose characteristics are stationary for an interval of time until they are redrawn from a reset distribution at abrupt changepoints distributed in time according to a Bernoulli process. The DBM assumptions lead to predictions about behavior that are consistent with an exponentially decaying lag profile. Consequently, the model fails to produce long-range effects of experience.

We propose an extension of the DBM, called the Hierarchical Dynamic Belief Model or HDBM (Figure 4a), that yields roughly a power function lag profile and consequently outperforms the DBM when fit to the entire experimental data in one pass (Figure 4b; Experiment 1: $t(26)=7.69$, $p<.0001$, Experiment 2: $t(19)=3.87$, $p=.0010$). The HDBM relaxes a seemingly unnatural assumption in the DBM: that environmental statistics have a time-invariant probability of change. For example, it would seem that the dynamics of change during a four-hour plane flight are not the same as those during the half hour it takes to deplane, walk through the terminal, collect bags, catch a taxi, and check into a hotel. The HDBM avoids this restrictive assumption by taking a hierarchical Bayesian approach in which the underlying generative model is a non-homogeneous Bernoulli process, i.e., a process with a fluctuating changepoint probability that is driven by a separate Markov process. Because the HDBM models a spectrum of environments—ranging from rapidly changing to stable—its expectations of the future reflect strong short-term recency as well as long-range dependencies.

The success of the HDBM in fitting the data suggests a normative explanation for the long-range influence of incidental experience on behavior. Under the assumptions of the HDBM, the mind optimally adapts to a complex dynamic environment in which even seemingly irrelevant experiences that occur far in the past offer predictive information about upcoming environmental states and task demands. Specifically, the expected relevance of a past experience to the current moment falls off according to an approximate power function.

As previously mentioned, human forgetting of explicit (declarative) knowledge in long-term memory is often characterized in terms of power decay. This decay function has been cast as rational via the observation that in diverse domains—newspaper articles, parental speech, and electronic mail—the empirical probability of needing access to a specific piece of information is well fit by a power function of time (Anderson & Schooler, 1991). The present analyses of the DBM and HDBM indicate that this observation is not well explained by nonstationarity with a fixed change probability, but that introducing variable change rates offers the basis for a normative explanation. Thus power decay serves as an informative connection between sequential effects, long-term memory, and the statistical structure of the environment.

Concluding Remarks

Contrary to the prevailing assumption that variations in experience produce only fleeting perturbations in behavior, we have argued that incidental priming yields enduring modulations of behavior. Modeling indicates that past experience is integrated to anticipate the future using a weighting that is strongly recency based but also has a heavy tail, consistent with power but not exponential discounting. Power discounting can be characterized as optimal adaptation to the statistics of an environment with second-order nonstationarity.

To perform optimal prediction in nonstationary environments with changepoint dynamics, the complete history of experience must be maintained (Adams & MacKay, 2006). Consequently, our results are consistent with the perspective that as individuals interact with their world, they continually log their experiences, forming a library of memory traces that is called on to adapt behavior to an environment that can change on timescales ranging from seconds to months. Alternatively, a good approximation to optimal prediction can be achieved by combining across several exponentially decaying sequence statistics that span a range of timescales (e.g., Kording & Tenenbaum, 2007; Mozer, Pashler, Cepeda, Lindsey, & Vul, 2009; Sikström, 1999, 2002; Staddon, Chelaru, & Higa, 2002; Wixted, 2004). Indeed, Mozer et al. (2009) and Murre and Chessa (2011) demonstrate mathematically that power functions emerge when an infinite collection of exponential functions are averaged together assuming certain constraints on the distribution of decay rates. Our work suggests the necessity of combining across multiple timescales ranging from just a few trials to hundreds of trials to the entire duration of the experiment. The presence of power decay, regardless of the precise mechanisms that produce it, suggests that sequential dependencies in rapid decision making are best understood as a memory phenomenon akin to human long-term declarative memory rather than as a byproduct of short-term incremental learning.

The perspective that sequential effects reflect memory storage and updating offers a novel interpretation of the continual and often long-range (Gilden, Thornton, & Mallon, 1995) fluctuations observed in human behavior and cognition. Far from being internal noise in the system, trial-to-trial variability in choice, response latency, and movement reflect an adaptive process in which individuals exploit their extensive experience to respond optimally to a dynamic world.

References

Adams, R. P. & MacKay, D. J. C. (2006). Bayesian online changepoint detection. Technical report, Cavendish Laboratory, University of Cambridge.

Anderson, J. R., et al. (2004). An integrated theory of the mind. *Psychological Review*, *111* (4), 1036-1060.

Anderson, J. R. & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science, 2*, 396-408.

Cho, R. et al. (2002). Mechanisms underlying dependencies of performance on stimulus history in a two-alternative forced-choice task. *Cognitive, Affective, & Behavioral Neuroscience, 2* (4), 283-299.

Doshi, A., Tran, C., Wilder, M., Mozer, M. C., & Trivedi, M. M. (2012). Sequential dependencies in driving. *Cognitive Science, 36*, 948-963.

Fine, M. S. & Thoroughman, K. A. (2006). Motor adaptation to single force pulses: sensitive to direction but insensitive to within-movement pulse placement and magnitude. *Journal of Neurophysiology, 96* (2), 710-720.

Furnham, A. (1986). The robustness of the recency effect: Studies using legal evidence. *Journal of General Psychology, 113*, 351-357.

Gilden, D.L., Thornton, T., & Mallon, M.W. (1995). 1/f noise in human cognition. *Science, 267*, 1837-1839.

Hebb, D. O. (1949). The Organization of Behavior. New York: Wiley & Sons.

Hogarth, R. M., & Einhorn, H. J. (1992). Order effects in belief updating: The belief adjustment model. *Cognitive Psychology, 24*, 1-55.

Hoppand, J. J. & Fuchs, A. F. (2004). The characteristics and neuronal substrate of saccadic eye movement plasticity. *Progress in Neurobiology, 72* (1), 27-53.

Jentzsch, I. & Sommer, W. (2002). Functional localization and mechanisms of sequential effects in serial reaction time tasks. *Perception and Psychophysics, 64* (7), 1169-1188.

Jones, M. & Sieck, W. (2003). Learning myopia: An adaptive recency effect in category learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 29*, 626-640.

Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. *Journal of Basic Engineering, 82* (1), 35-45.

Kasif, S., Salzberg, S., Waltz, D., Rachlin, J., & Aha, D. (1998). A probabilistic framework for memory-based reasoning. *Artificial Intelligence, 104* (1-2), 287-311.

Kording, K. P., Tenenbaum, J. B., & Shadmehr, R. (2007). The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nature Neuroscience, 10*, 779-786.

Link, B. V., Kos, B., Wager, T. D., & Mozer, M. C. (2011). Past experience influences judgment of pain: Prediction of sequential dependencies. In L. Carlson, C. Hoelscher, & T. F. Shipley (Eds.), *Proceedings of the 33d Annual Conference of the Cognitive Science Society*, 1248-1253. Austin, TX: Cognitive Science Society.

Loftus, G. R. & Masson, M. E. J. (1994). Using confidence intervals in within-subject designs. *Psychonomic Bulletin & Review, 1*, 476-490.

Mozer, M., Kinoshita, S., & Shettel, M. (2007). Sequential dependencies offer insight into cognitive control. In W. Gray (Ed.), Integrated Models of Cognitive Systems, 180-193. Oxford University Press.

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.

Mumma, G. H., & Wilson, S. B. (2006). Procedural debiasing of primacy/anchoring effects in clinical-like judgments. Journal of Clinical Psychology, 51, 841-853.

Murre, J. M. J., & Chessa, A. G. (2011). Power laws from individual differences in learning and forgetting: mathematical analyses. *Psychonomic Bulletin & Review, 18* (3), 592-597.

Neiman, T. & Loewenstein, Y. (2011). Reinforcement learning in professional basketball players. *Nature Communications, 2*, Article number: 569. doi:10.1038/ncomms1580

Remington, R. J. (1969). Analysis of sequential effects in choice reaction times. *Journal of Experimental Psychology, 82*, 250-257.

Rescorla, R. A. & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A.H. Black & W.F. Prokasy (Eds.), Classical Conditioning II. 64-99. Appleton-Century-Crofts.

Robinson, F. R., Soetedjo, R., & Noto, C. (2006). Distinct short-term and long-term adaptation to reduce saccade size in monkey. *Journal of Neurophysiology, 96* (3), 1030-1041.

Rubin, D. C. & Wenzel, A. E. (1996). One hundred years of forgetting: A quantitative description of retention. *Psychological Review, 103*, 734-760.

Scheidt, R. A., Dingwell, J. B., & Mussa-Ivaldi, F. A. (2001). Learning to move amid uncertainty. *Journal of Neurophysiology, 86* (2), 971-985.

Shadmehr, R. & Mussa-Ivaldi, F. A. (1994). Adaptive representation of dynamics during learning of a motor task. *Journal of Neuroscience, 14*, 3208-3224.

Sikström, S. (1999). Power-function forgetting curves as an emergent property of biologically plausible neural networks model. *International journal of psychology, 34*(5/6), 460-464.

Sikström, S. (2002). Forgetting curves: Implications for connectionist models. *Cognitive Psychology, 45*(1), 95-152.

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 & Performance, 11*, 598-616.

Staddon, J. E. R., Chelaru, I. M., & Higa, J. J. (2002). Habituation, memory and the brain: The dynamics of interval timing. *Behavioural Processes, 57*, 71-88.

Stanfill, C. & Waltz, D. (1986). Toward memory-based reasoning. *Communications of the ACM*, *29* (12), 1213-1228.

Sutton, R. S., Barto, A. G. (1998). Reinforcement Learning: An Introduction. MIT Press.

Vuong, Q. H. (1989). Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica, 57*(2), 307-333.

Wagenmakers, E.-J., Grünwald, P., & Steyvers, M. (2006). Accumulative prediction error and the selection of time series models. *Journal of Mathematical Psychology, 50*, 149-166.

Ward, L. & Lockheed, G. (1970). Sequential effects and memory in category judgements. *Journal of Experimental Psychology, 84* (1), 27-34.

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 22* (pp. 2053-2061). La Jolla, CA: NIPS Foundation.

Wixted, J. T. (2004). On common ground: Jost's (1897) law of forgetting and Ribot's (1881) law of retrograde amnesia. *Psychological Review, 111*, 864-879.

Wixted, J. T. & Ebbesen, E. B. (1997). Genuine power curves in forgetting: A quantitative analysis of individual subject forgetting functions. *Memory & Cognition, 25*, 731-739.

Wixted, J. T. & Carpenter, S. K. (2007). The Wickelgren power law and the Ebbinghaus savings function. *Psychological Science, 18*, 133-134.

Wixted, J. T. & Ebbesen, E. B. (1991). On the form of forgetting. *Psychological Science, 2,* 409-415.

Wong, A. L. & Shelhamer, M. (2011). Exploring the fundamental dynamics of error-based motor learning using a stationary predictive-saccade task. *PLoS ONE, 6* (9): e25225.

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.

Appendix

Modeling Details

The models form an expectation for trial t based on a weighting of past trials, $w(l)$ for trial t-l, and yield a quantity ϕ_t reflecting the match between expectation and actual outcome:

$$\phi_t = x_t \sum_{l=1}^{\min(t-1,T)} w(l) x_{t-l} \,,$$

where $w(l) = \lambda^l$ and $w(l) = (1+l)^\kappa$ for the exponential and power models, respectively, $x_t \in \{-1,1\}$ denotes the binary type of trial t (repetition versus alternation for Experiment 1, left versus right for Experiment 2), and T is the context horizon.

To fit data, ϕ_t is converted to an RT or movement error via an affine transformation. In both experiments, an additive offset was incorporated in the transformation of repetition trials to allow for a default bias towards repetitions or alternations commonly observed in 2AFC studies. Transformation and model parameters were fit to each subject separately to minimize the mean squared error across individual trial predictions for the entire sequence of trials and were constrained to be equal for the two conditions of Experiment 1.

HDBM Mathematical Specification

The Dynamic Belief Model (DBM) assumes that individuals maintain a distribution over a single environmental statistic, γ_t, that represents the probability of a repetition vs. alternation (Experiment 1) or left vs. right (Experiment 2). The value of γ_t is inferred from the sequence history, $\mathbf{x}_{t-1}$, subject to the constraints of a fixed change probability, α. The expectation match, ϕ_t, is defined to be $P(x_t \mid \mathbf{x}_{t-1}, \alpha)$ which is given by $E[\gamma_t \mid \mathbf{x}_{t-1}]$ when x_t is a repetition and $1 - E[\gamma_t \mid \mathbf{x}_{t-1}]$ when x_t is an alternation. The posterior distribution over γ_t is iteratively updated:

$$p(\gamma_t \mid \mathbf{x}_{t-1}, \alpha) = (1-\alpha) p(\gamma_{t-1} \mid \mathbf{x}_{t-1}, \alpha) + \alpha p_\gamma \text{, with}$$

$$p(\gamma_{t-1} \mid \mathbf{x}_{t-1}, \alpha) \propto P(x_{t-1} \mid \gamma_{t-1}) p(\gamma_{t-1} \mid \mathbf{x}_{t-2}, \alpha),$$

where p_γ is the standard uniform. (See Yu & Cohen, 2009, for more details).

In the Hierarchical Dynamic Belief Model (HDBM), instead of assuming a fixed change probability α, we define α_t as a time-varying change probability subject to the same dynamics that govern γ_t in the DBM. Specifically, with probability η, called the 'meta change probability', α_t will be redrawn from a Beta resampling distribution, p_α, and with probability $1-\eta$, α_t will remain unchanged. In the HDBM, ϕ_t is defined as

$$P(x_t \mid \mathbf{x}_{t-1}) = \int_0^1 P(x_t \mid \mathbf{x}_{t-1}, a) p(\alpha_t = a \mid \mathbf{x}_{t-1}) da,$$

where $P(x_t \mid \mathbf{x}_{t-1}, a)$ is the DBM probability for the fixed changepoint *a*. The posterior distribution over α_t is recomputed iteratively:

$$p(\alpha_t = a \mid \mathbf{x}_{t-1}) = (1-\eta) p(\alpha_{t-1} = a \mid \mathbf{x}_{t-1}) + \eta p_\alpha(\alpha_t = a), \text{ with}$$

$$p(\alpha_{t-1} = a \mid \mathbf{x}_{t-1}) \propto P(x_{t-1} \mid \mathbf{x}_{t-2}, a) p(\alpha_{t-1} = a \mid \mathbf{x}_{t-2}).$$

The HDBM has 3 free parameters: the meta-change probability and 2 parameters for the resampling distribution p_α.

Acknowledgements

We thank Ines Jentzsch for sharing her data. We thank Brent Young, Colin Argys, Chris Bird, Casey DeBuse, Alex Eichenbaum, Elizabeth Eustis, Megan Freeman, Robert Garcia, William Hall, Lindsey Johnson, and Emily Kleinfelder for helping with Experiment 1; and Bianca Bzdel, David Busick, and Stefan Garcia for Experiment 2.

Research supported by NSF grants SBE-0518699, BCS-0720375, SBE-0542013, SMA-1041755, and AFOSR grant FA9550-10-1-0177.

Figure Captions

Figure 1

Reanalysis of a representative sequential effects study (Jentzsch & Sommer, 2002, Expt. 1) a) Mean RTs for current trial type—repetition (R) or alternation (A)—as a function of sequence history (current trial at top of label). Error bars here and elsewhere for behavioral data indicate standard error. Exponential (red) and power (blue) models—with full context horizon—fit to per-subject trial-by-trial data and averaged across subjects. b) Accumulative prediction error (R^2) as a function of context horizon. Error bars indicate standard error of R^2 difference between models (Loftus & Masson, 1994), thus aiding in comparing models but not horizons.

Figure 2

a) Mean RTs for positive and negative autocorrelation conditions as a function of sequence history. Exponential and power models—with full context horizon—fit to per-subject trial-by-trial data and averaged across subjects. b) Accumulative prediction error (R^2) as a function of horizon. Error bars as in Figure 1. c) Lag profile averaged across conditions and subjects in log-log coordinates. Mean of exponential and power function fits to per-subject lag profiles. d) Difference in mean RT for repetition and alternation trials by block (234 trials) for each autocorrelation condition. e) Histogram of the log-likelihood ratios for individual subject fits (negative supports power model and positive supports exponential model). Significance determined by the Vuong's closeness test.

Figure 3

a) Experimental setup. b) Mean trajectories from blue dot to green dot for different sequences of right (R) and left (L) perturbations (current trial at right end of label). Sequential dependencies here result from the history of right and left forces rather than the repetition/alternation sequences (we anticipated this based on the theoretical division between perceptual and response sequential effects, see Wilder, Jones, & Mozer, 2010). c) Accumulative prediction error (R^2) as a function of context horizon. Error bars as in Figure 1. d) The lag profile in log-log coordinates with mean exponential and power function fits. e) Histogram of the log-likelihood ratios for individual subject fits (negative supports power model and positive supports exponential model). Significance determined by the Vuong's closeness test.

Figure 4

a) The graphical model for the Hierarchical Dynamic Belief Model (HDBM). x_t is the trial type at time t, γ_t is the parameter of the Bernoulli process generating x_t, and α_t is the change probability. The original DBM (Yu and Cohen, 2009) consists of only the black part of the graph, with α constant. b) Comparison of model performance for Experiment 1 and 2. Error bars for the power and exponential models—and similarly for the HDBM and DBM models—represent the standard error of the R^2 difference between the two models across subjects.

Figures

Figure 1a

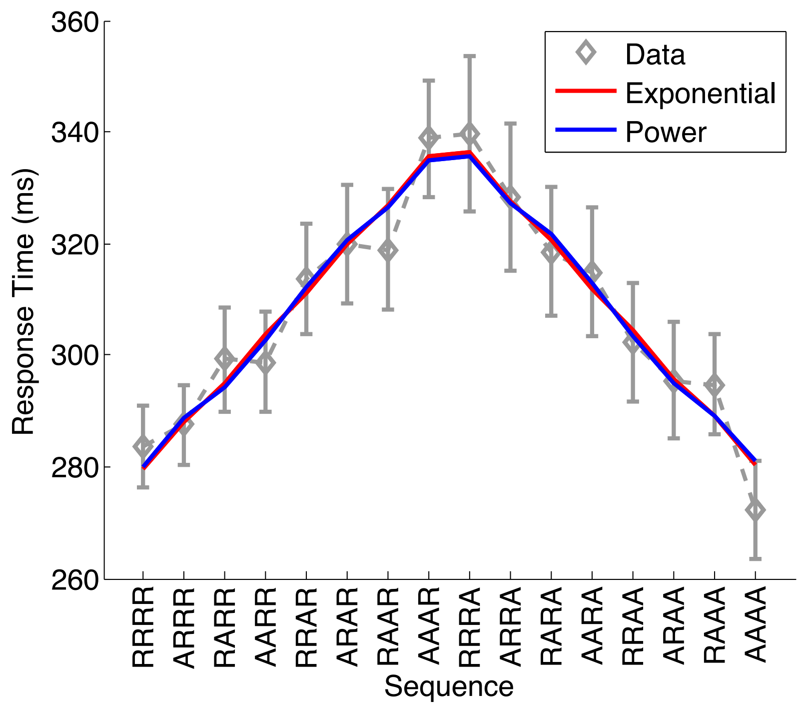

Figure 1b

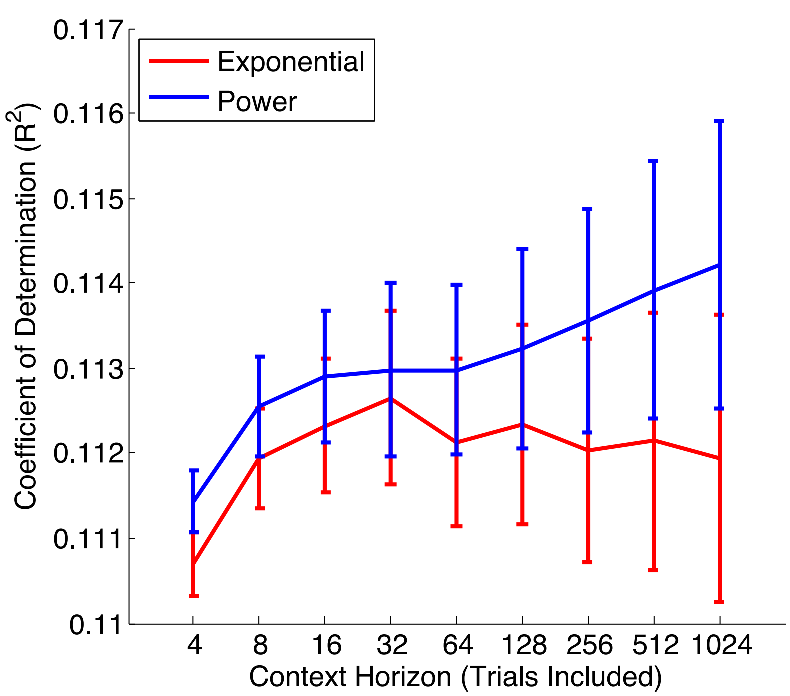

Figure 2a

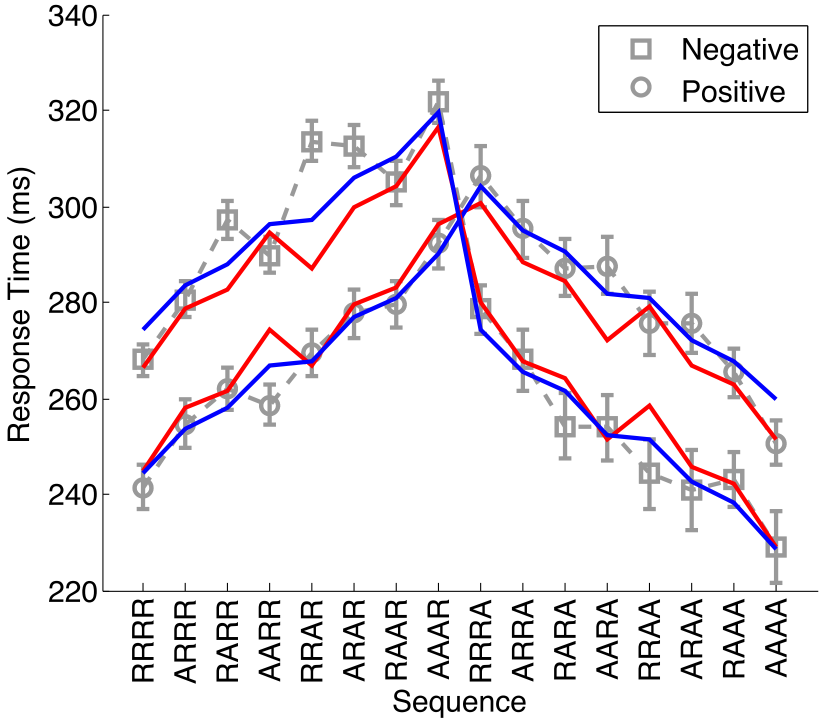

Figure 2b

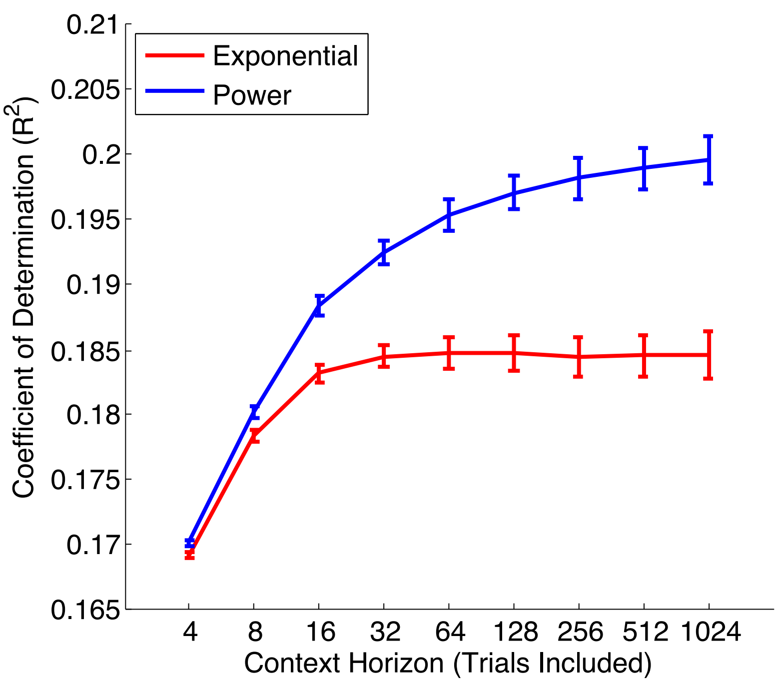

Figure 2c

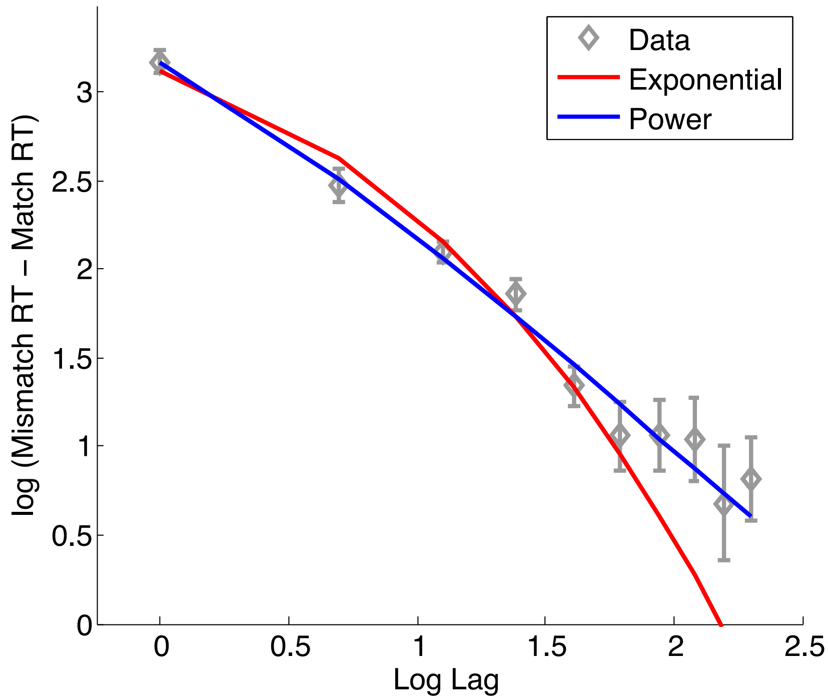

Figure 2d

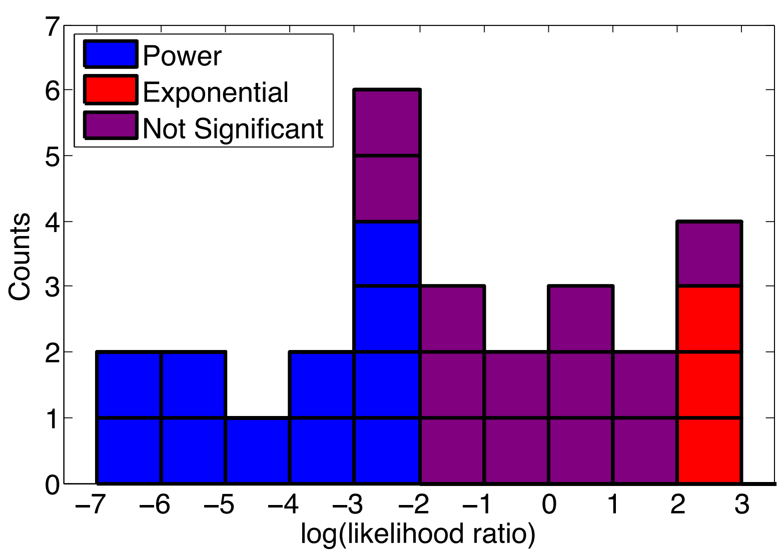

Figure 2e

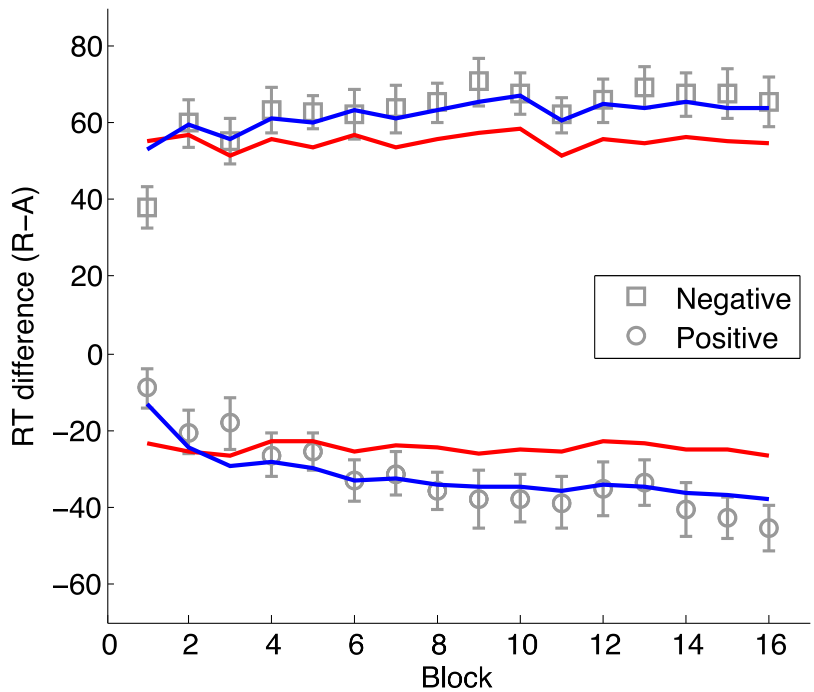

Figure 3a

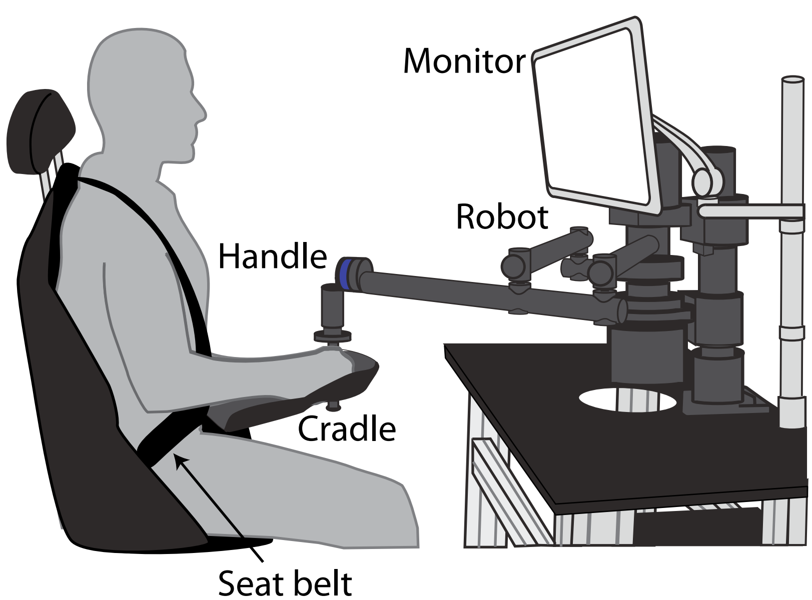

Figure 3b

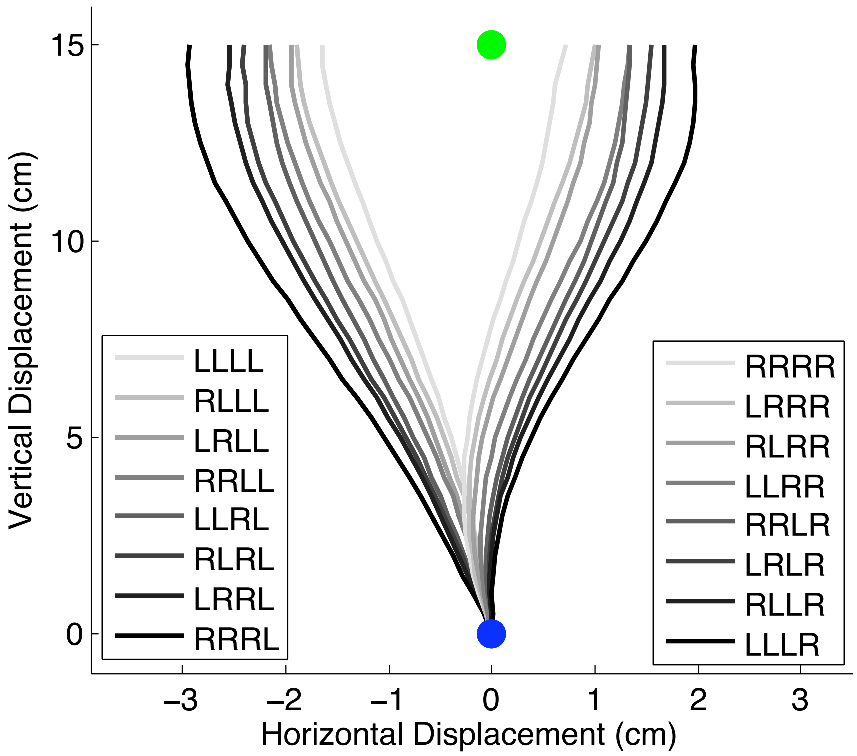

Figure 3c

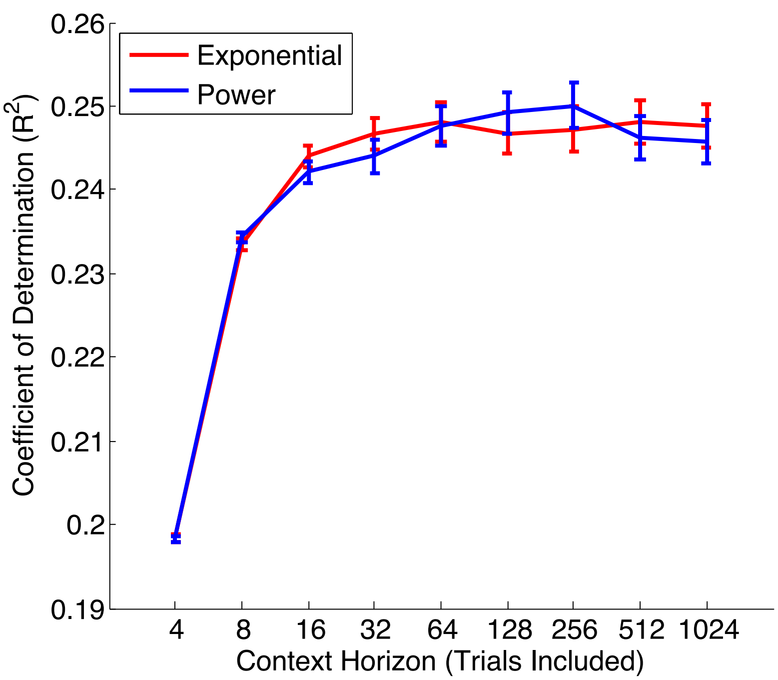

Figure 3d

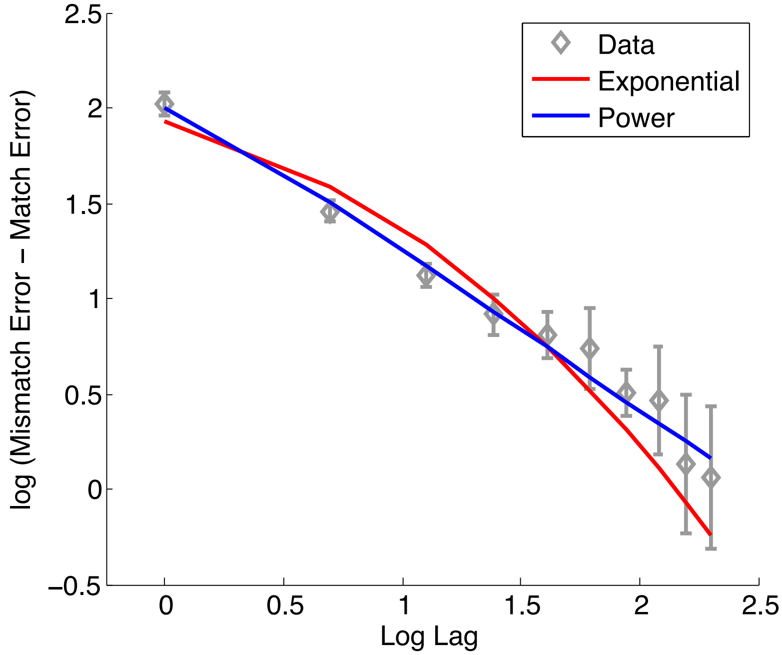

Figure 3e

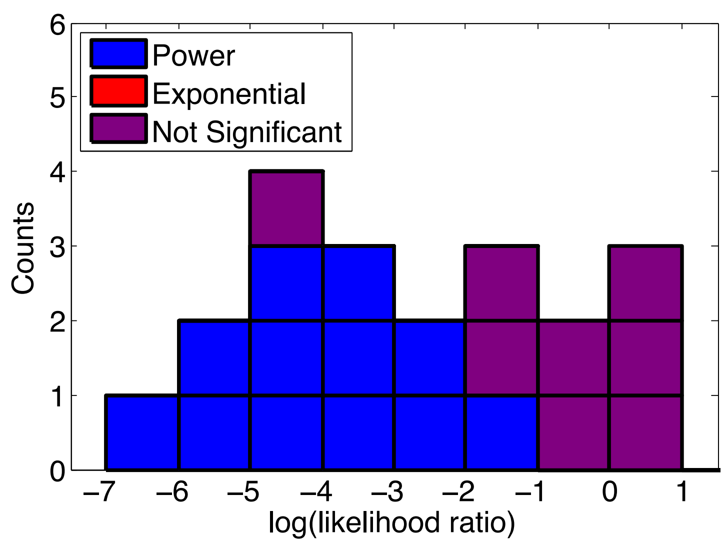

Figure 4a

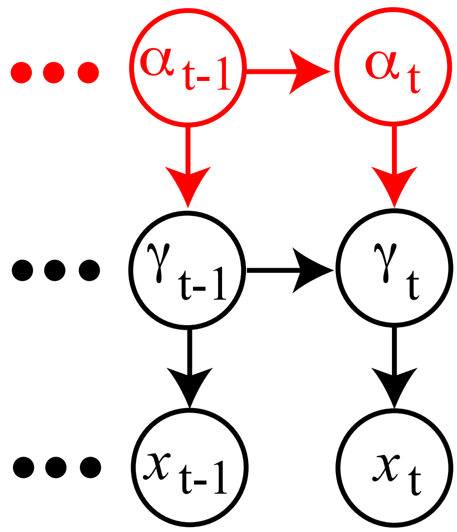

Figure 4b