Extrapolative Beliefs in Perceptual and Economic Decisions:
Evidence of a Common Mechanism

Cary Frydman and Gideon Nave[1]

April 2015

ABSTRACT

A critical component of both economic and perceptual decision-making under uncertainty is the belief formation process. However, most research has studied belief formation in economic and perceptual decision-making in isolation. One reason for this separate treatment may be the assumption that there are distinct psychological mechanisms that underlie belief formation in economic and perceptual decisions. An alternative theory is that there exists a common mechanism that governs belief formation in both domains. Here, we test this alternative theory by combining a novel computational modeling technique with two well-known experimental paradigms. We estimate a drift-diffusion model (DDM) and provide an analytical method to decode prior beliefs from DDM parameters. Subjects in our experiment exhibit strong extrapolative beliefs in both paradigms. We find that a common computational model explains belief formation in both tasks, and that individual differences in belief formation are correlated across tasks.

---

[1] Authors contributed equally to this work and names are listed in alphabetical order. Frydman is in the Finance and Business Economics Department at the Marshall School of Business, University of Southern California, email: cfrydman@marshall.usc.edu. Nave is in the Department of Computational and Neural Systems, Caltech, email: gnave@caltech.edu.

1. Introduction

A common goal of economics and psychology is to understand the mechanisms that govern decision-making in uncertain environments (Platt and Huettel 2008). This is an ambitious goal that seeks to explain the computational processes underlying both (i) perceptual decisions – those that are determined by judgments about objective states of the world (Gold and Heekeren 2013), and (ii) value-based (economic) decisions – those that are determined by subjective beliefs and preferences (Rangel et al. 2008). Over the past several decades a vast amount of research has been devoted to studying these two types of decision-making. Most of this research has proceeded along two parallel tracks, perhaps because of an implicit assumption that the psychological mechanisms governing perceptual and economic decisions are distinct.

However, a small but growing body of literature has begun to investigate the links between economic and perceptual decision-making (Summerfield and Tsetsos 2012, 2015). For example, the evidence accumulation process that is used to describe perceptual decision-making can also explain the dynamics in simple economic decisions (Krajbich and Rangel 2011, Tsetsos et al. 2012, Polanía et al. 2014). Conversely, robust decision biases from economic decision-making – such as context effects – have recently been uncovered in perceptual decision-making (Trueblood et al. 2013). There is also likely to be an important interaction between the two domains as visual saliency can systematically bias economic decisions (Towal et al. 2013). Despite this growing evidence, it is unknown whether there is a similar link between *dynamic* (across-trial) aspects of perceptual and economic decision-making. In other words, does there exist a common belief formation mechanism across these two decision-making domains?

The answer to this question is important because it can help shed light on the origins and mechanisms of belief-based biases in judgment and decision-making. Moreover, these biases are key ingredients in many behavioral models in finance (Barberis et al. 1998, 2015, Hong and Stein 1999), and thus understanding the microfoundation of such biases may be useful in building more psychologically realistic models of financial markets. For example, recent survey evidence shows that many investors hold extrapolative beliefs, where they expect stock prices to continue rising after they have previously risen, and to fall have they have previously fallen (Greenwood and Shleifer 2014). If these extrapolative beliefs are driven by the same mechanism that governs belief formation in lower level perceptual processes, characterizing this common mechanism can be useful in modeling higher-level economic expectation formation.

In this paper, we develop an experimental paradigm that allows us to test the hypothesis that belief formation is governed by the same psychological mechanism in both economic and

perceptual decision-making. In order to provide a strong test of this hypothesis, it is necessary to precisely measure beliefs in both types of decision-making environments. Fortunately, economists have long since provided incentive-compatible experimental methods for measuring beliefs in economic decisions (Becker et al. 1964). However, directly measuring beliefs in a perceptual decision-making task remains challenging because subjects themselves may not have access to a conscious representation of their belief (Curran and Keele 1993). To overcome this challenge, we provide a novel computational technique that decodes beliefs from response-time (RT) data. This technique builds on the large literature of sequential sampling models (SSMs) in perceptual decision-making (Townsend and Ashby 1983, Usher and McClelland 2001, Ratcliff and Smith 2004, Bogacz et al. 2006) and more recently, in value based decision-making (Fehr and Rangel 2011, Webb 2013, Woodford 2014).

Most studies that use SSMs proceed by manipulating an attribute of the environment (e.g., stimulus motion coherence or subjective value) and testing which computational parameters encode the change in environment (Klauer et al. 2007, Mulder et al. 2012, White and Poldrack 2014). In contrast, the computational technique we employ to measure beliefs focuses on decoding *perceived* changes in the environment (i.e., changes in subjective beliefs) from an SSM. In particular, we show theoretically how the estimated initial point and boundary parameters of a drift diffusion model can be used to infer a subject's prior. Our study therefore provides a novel example of how neuro-computational models can be used to measure a subject's belief formation process.

To demonstrate this, we recruited subjects to participate in two separate tasks: an economic decision-making task (EDT) and a perceptual decision-making task (PDT). While each task has been used several times in its own literature separately (Bloomfield and Hales 2002, Cho et al. 2002, Huettel et al. 2002, Asparouhova et al. 2009), in the current study we have subjects participate in both the EDT and PDT in the same experimental setting. This within-subject design allowed us to measure and compare the computations governing belief formation across different decision domains. We hypothesized that if a single psychological mechanism governs belief formation across economic and perceptual decision domains, then (i) a common computational model should explain belief-updating (across trials) in both tasks and (ii) individual differences in the degree to which subjects rely on recent stimulus history to update beliefs should be correlated across tasks.

2. Methods

2.1 Subjects

Thirty-eight subjects (*17* females) aged *17-29* (mean: *20.24* SD: *3.11*) participated in the study. Subjects were students at Caltech or at a nearby community college and the sample size was chosen to match the exact sample size employed in previous work with the same task (Bloomfield and Hales, 2002). The California Institute of Technology and University of Southern California Institutional Review Boards approved this study, and the subjects gave informed consent.

2.2 Perceptual decision task

The PDT consisted of four blocks of *300* trials each, preceded by *5* training trials. Each trial began with the appearance of a white fixation cross in the center of a black screen; after 800 milliseconds, either a white circle (diameter: *10.5 cm*) or a blue square (width: *10.5 cm*) appeared in the place of the cross. Subjects were instructed to respond by either pressing the "right arrow" key when a circle had appeared or the "left arrow" key when a square had appeared (Figure 1A). Subjects were told that they would receive one cent for each correct response, and their earnings would be reduced by 0.05 cents for every 100 milliseconds of delay in their response. If the response was slower than 2 seconds they would receive 0 cents. A new trial started immediately following a response, with the appearance of a new fixation cross on the screen. Subjects were told that there were only two possible stimuli (a circle or a square), that the chance of seeing either shape was ½, and that the stimuli of previous trials had no influence on future trials. The actual sequence of stimuli was an independent and identically distributed pseudo-random binary process, such that the stimuli sequence was identical for all subjects. Subjects had a break of 20 seconds between blocks, during which they received feedback about the number of correct responses they made, but no feedback about their mean RT. The task instructions are available in the E-Companion.
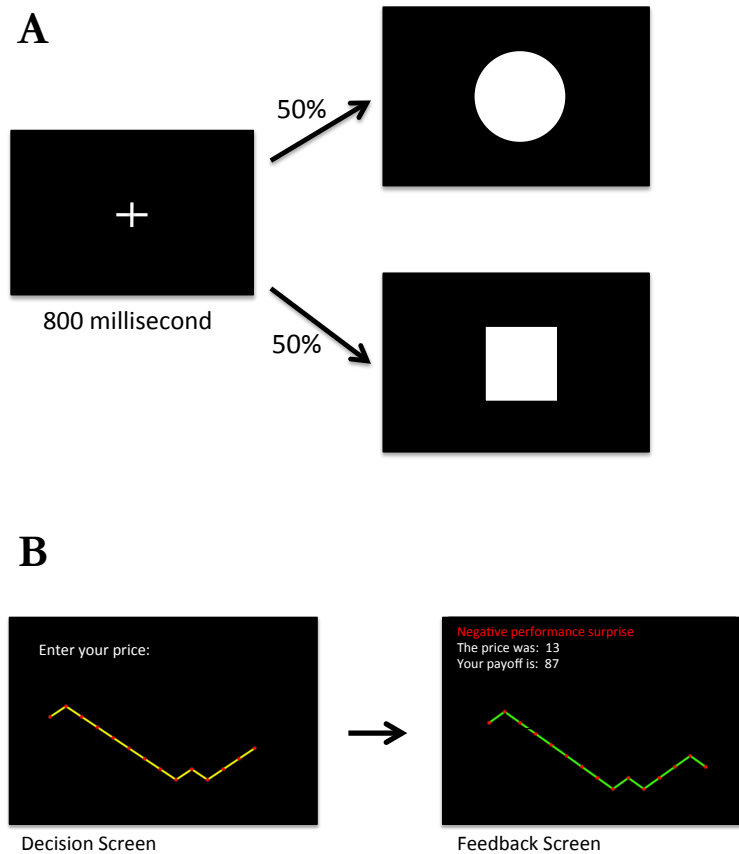
2.3 Economic decision task

In the EDT, subjects were told that data from publicly traded firms was used to create a model that generated sequences of "performance surprises" – a time series of actual performance minus predicted performance (Bloomfield and Hales 2002, Asparouhova et al. 2009). Subjects were presented with a sequence performance surprises from a typical firm, and were instructed to predict whether the next performance surprise would be positive or negative. In each period, subjects saw a history chart of performance surprises from the last 14 periods plotted in yellow on a black screen (see Figure 1B). Each period, subjects were endowed with 100 units of

4

experimental cash and were asked to state the maximal price at which they would be willing to buy a stock of the company. To elicit an incentive-compatible measure of a subject's willingness to pay (WTP) for the stock, we used a Becker–DeGroot–Marshak (BDM) auction (Becker et al. 1964). Thus, after the subject stated his WTP, the actual price was drawn from a uniform distribution between 0 and 100, and the stock would be purchased at the price drawn if and only if its price was less than the WTP. If the subject purchased the stock, it would return 100 units of experimental cash in the case of a positive surprise and 0 units of experimental cash otherwise. The WTP measures the expected value of the stock, and therefore, for a risk neutral subject, it provides a measure of the subjective belief of a positive performance surprise.

Subjects were explicitly told that in order to make the most money in the task, they should set the price equal to the probability of seeing a positive surprise. After stating their WTP, subjects received feedback about the price drawn, the stock performance and their financial outcome. The task consisted of a single block of 400 trials preceded by 10 training trials, where each period started at the same point in the sequence where the previous trial had ended. The actual sequence was a pseudo-random independent and identically distributed binary process (the stimuli were identical for all subjects). Every 50 trials subjects saw their accumulated payoff and were allowed to take a short break. At the end of the experiment, subjects were paid for all of their trials, using a conversion rate of 5,000 units of experimental cash = 1 USD. The task instructions are available in the E-Companion.

2.4 Experimental Procedures

The data was collected over three experimental sessions, all conducted at the Caltech Social Science Experimental Laboratory. At the beginning of each session, participants were randomly allocated to cubicles in the lab, where they could not see or interact with each other. Before each task, subjects received printed instructions that were also read out loud by the experimenter, and subjects subsequently had an opportunity to ask questions. Each session started with the PDT, followed by the EDT (this design choice is discussed in the E-Companion). Both tasks were programmed using Matlab Psychtoolbox (Brainard 1997).

Figure 1. Experimental design of the economic decision-making task and perceptual decision-making task. (A) PDT: Following a display of a fixation cross at the center of the screen (800 milliseconds), either a circle or a square was presented in random order over the course of 1200 trials. Subjects were incentivized to respond to each shape with a different key press as quickly and as accurately as possible. A new trial started immediately following the response, with the appearance of a new fixation cross. (B) EDT: On each of 400 trials, subjects entered the price, *p*, at which they would be willing to buy a stock. A price *x* was then randomly drawn and if *x<p,* the subject purchased the stock at a price of *x* on that trial. The stock then paid $100 if there was a positive performance surprise, and $0 otherwise.

3. Results

3.1 Basic Results from the Perceptual Decision-Making Task

We found that RTs and error rates systematically varied as a function of recent stimulus history, despite the fact that subjects were explicitly told the probability of displaying a circle was 0.5 on all trials, in accordance with previous studies (Cho et al. 2002, Huettel et al. 2002). On trials

where the stimulus continued the recent streak ("continuation" trials), RTs decreased with streak length. In contrast, on trials where the stimulus violated the recent streak (e.g. 3 circles followed by a square) we found that RTs moderately increase with streak length (Figure 2A). Figure 2B displays error rates as a function of streak length, where an error is defined as misclassifying the stimulus. We found that as the streak length continues, error rates decrease for continuation trials, but they increase for violation trials. For each streak length, the difference between the continuation RT and the violation RT, as well as the difference between the continuation error rate and the violation error rate are significantly different from zero at the 5% level (see E-Companion for statistics on each pairwise hypothesis test).

The perceptual decision-making literature has highlighted two main types of sequential effects. First, *automatic facilitation* (AF) effects occur when reaction times following a certain streak length are faster regardless of whether the current tryeial extends the streak or not. For example, after the sequence, {square, square, square}, an AF effect predicts that the response on the subsequent trial to either a circle of square will be faster than after the sequence {circle, square, square}. This type of effect could be driven by post-response residual activity (Roitman and Shadlen 2002) depending on the decay rate of neuronal activity in the motor cortex. In contrast, *strategic expectancy* (SE) effects occur when reaction times following a certain streak length are faster only when the current trial extends the streak. This means that after the sequence, [square, square, square], an SE effect predicts that the response on a square trial will be faster than on a circle trial.

The results in Figure 2 are inconsistent with AF effects as the mean RT for each streak length depends on the stimulus identity of the current trial. However, they are consistent with SE effects, which can be interpreted as effects driven by expectations about future stimuli. This finding is consistent with an extensive body of literature showing that AF effects do not occur when the response-stimulus interval (RSI) is greater than 200 milliseconds (Soetens et al. 1985). As the objective of this study was to investigate expectation formation, we intentionally designed the experiment with an RSI of 800 milliseconds in order to minimize AF sequential effects (Cho et al. 2002, Gao et al. 2009)

3.2 Basic Results from the Economic Decision-Making Task
Consistent with previous research, we found substantial evidence of extrapolation based on the previous history of stimuli (Bloomfield and Hales 2002, Asparouhova et al. 2009). Specifically, we found that the longer the current streak of positive (negative) performance surprises, the

higher the reported probability of a subsequent positive (negative) performance surprise (Figure 2C).
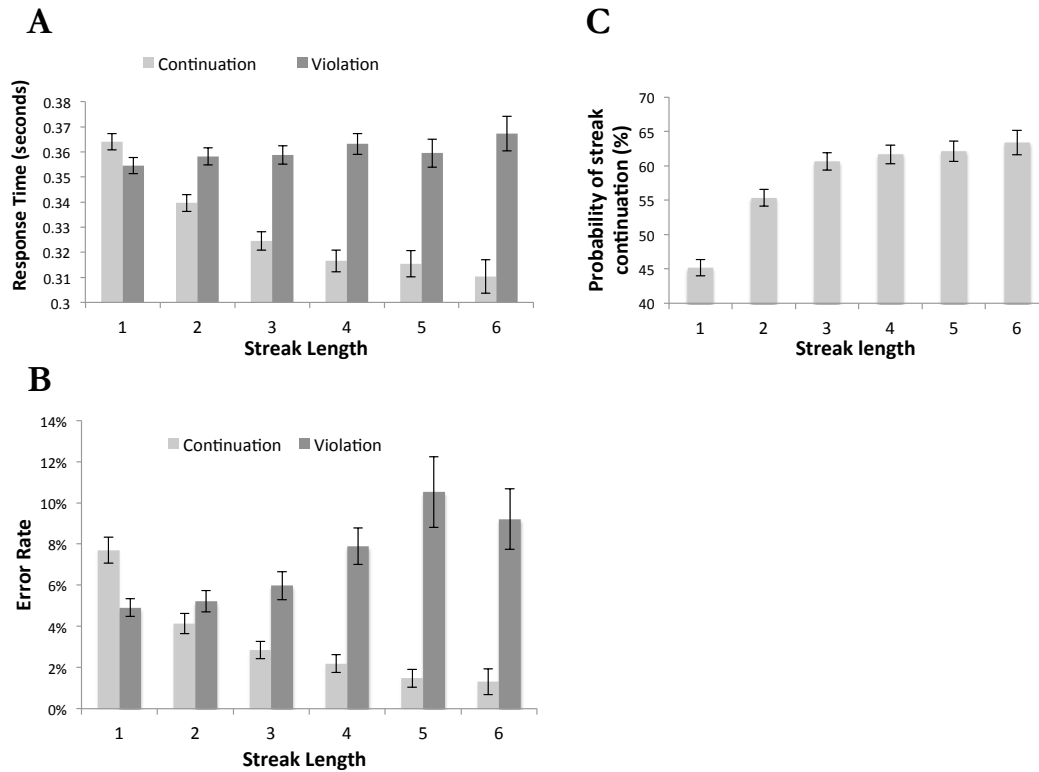
Figure 2. (A) Average RT as a function of current streak length (PDT), where streak length is defined as the number of consecutive identical stimuli. Continuation are those trials where the streak continues, violations are those trials where the steak is violated. (B) Average error rate as a function of current streak length (PDT). (C) Average reported beliefs (EDT) that the current streak would continue, as a function of streak length.

3.3 A Structural Model of Decision-Making in the PDT

3.3.1 DDM Overview

In order to investigate whether there is a common link between belief formation in the EDT and the PDT, we needed to first transform the RT and accuracy data from the PDT into probabilities (prior beliefs), allowing us to compare them with the beliefs elicited in the EDT. This is a non-trivial exercise, but we were able to overcome this obstacle by estimating a structural model of the decision-making process in the PDT. While there is a variety of structural models of perceptual decision-making in the psychology and neuroscience literature, we chose to model our data using a DDM for two main reasons. First, in addition to providing an accurate description of

perceptual decision-making data, recent work has shown that this model can also provide a good fit for behavior and reaction times in economic decision-making tasks (Fehr and Rangel 2011, Krajbich et al. 2014). As our goal is to provide links between belief formation in perceptual and economic decision-making, it is valuable to use a structural model that can explain the data well in both domains. Second, there are specific parameters of the DDM that have been shown to map on to changes in prior beliefs, which is the key object in our study (Mulder et al. 2012, Teodescuru and Usher 2013, White and Poldrack 2014).[2]

The DDM was originally developed to explain the response times and accuracy of perceptual decisions in binary choice tasks (Ratcliff 1978, Ratcliff and McKoon 2008). The basic assumptions of this model are that incoming sensory evidence about the identity of a stimulus (e.g., a circle or square) is noisy and decision time is costly. The DDM implements a choice algorithm that minimizes the decision time for a given level of accuracy[3] (Bogacz et al. 2006).

A brief description of the DDM as applied in our experimental setting is useful to understand the analytical technique we use to decode prior beliefs from the PDT. The DDM assumes that the brain computes a relative decision value (RDV) that measures the accumulated relative "evidence" in favor of the correct option, and this RDV evolves over time until a choice is made (Figure 3). The RDV follows a diffusion process:

$$(1) \quad dRDV(t) = Mdt + sdW ,$$

with initial point $RDV(0) = c_i$, for condition i. A choice is made once the RDV reaches one of two thresholds, where we normalize the lower threshold to zero and the upper threshold to a constant, a. In this model, M is the drift rate and represents the strength of incoming sensory information that a subject uses to infer the identity of the current shape. When the discriminability between the two possible stimuli is high, the drift rate is large; if instead, the two shapes are difficult to discriminate, the incoming sensory evidence in favor of one option versus the other is low, and the drift rate is small. The variable c_i represents the initial point in condition i and can parameterize the prior bias towards selecting the correct alternative (we use the convention that the upper boundary is associated with the correct alternative). Finally, s represents the standard

---

[2] In the next subsection, we test which specific parameters from the DDM encode changes in prior beliefs.
[3] Conversely, it can be seen as solving the dual optimization problem of maximizing decision accuracy for a given amount of decision time.

deviation of mean-zero Gaussian distributed noise, which we set to *s=0.1* without loss of generality, and dW is a Weiner process.[4]
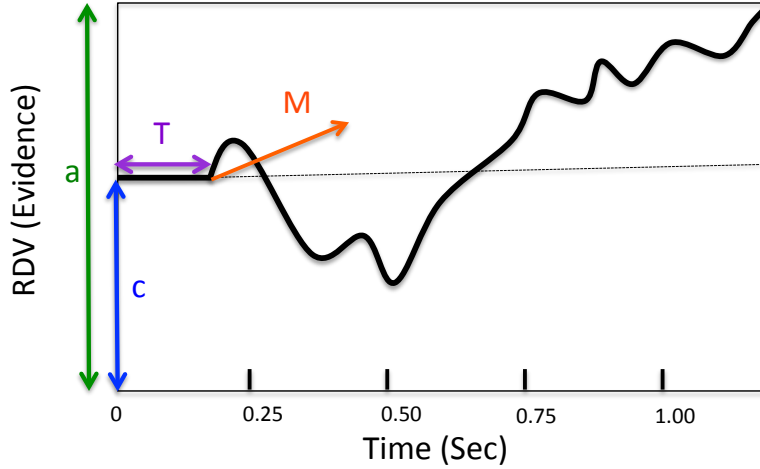
Figure 3. A graphical representation of the drift diffusion process. The bold path indicates the evolution of the relative decision value (RDV) that tracks the relative evidence in favor of the alternative associated with the upper boundary. *T* denotes the non-decision time, *c* denotes the initial point that captures prior bias, *a* denotes the upper boundary that captures the speed-accuracy trade-off and *M* denotes the drift rate that represents the quality of sensory input. When the RDV reaches a boundary, the process terminates and a decision is made. Without loss of generality, the lower boundary is set to zero.

The non-decision time, denoted by *T,* represents the time required to encode the stimulus and implement the motor action. In addition to the four parameters (*M, a, c_i,* and *T)* we include three additional across trial variability parameters that have been shown to improve the accuracy of describing observed RT distributions (Ratcliff and McKoon, 2008). First, a variance parameter M_{σ^2}, characterizes the distribution from which the drift rate is sampled in each trial, such that *M* is normally distributed with mean M_μ and variance M_{σ^2}. Second, a range parameter T_{σ^2} characterizes the distribution from which the non-decision time is sampled in each trial, such that $T \sim U[\frac{T_\mu - T_{\sigma^2}}{2}, \frac{T_\mu + T_{\sigma^2}}{2}]$. Finally, a range parameter $c_{i_{\sigma^2}}$ characterizes the distribution of the initial point in condition *i*, so that $c_i \sim U[\frac{c_{i_\mu} - c_{i_{\sigma^2}}}{2}, \frac{c_{i_\mu} + c_{i_{\sigma^2}}}{2}]$.

---

[4] Because the noise parameter *s*, drift rate *M*, and boundary separation *a,* are only defined up to positive affine transformations, one could fix any of these three parameters and estimate the remaining two. We choose to follow the convention of the studies that fix the noise parameter and estimate the boundary and drift rate (Ratcliff and Tuerlinckx 2002). The boundary and drift rate are then interpreted in units of standard deviation of noise.

The DDM was estimated at the individual subject level using the DMA-Toolbox (Tuerlinckx and Vandekerckhove 2008). This toolbox estimates the DDM parameters by minimizing the multinomial log-likelihood function:

$$(2)\ LL(x_t, b_t | \theta_i) = -2 \sum_{c=i}^{C} \sum_{t=1}^{T_i} \log\left(L(x_t, b_t | \theta_i)\right),$$

where $L(x_t, b_t | \theta_i)$ is the likelihood of observing, on trial t, a response (correct or incorrect) x_t and RT in bin b_t, conditional on parameters θ_i, where i denotes the experimental condition for trial t. Following the literature on sequential effects in perceptual decision-making we define an experimental condition by the 4-back history of repetitions (R) and alternations (A), thus setting $C=16$ total conditions (Soetens et al. 1985, Cho et al., 2002, Yu and Cohen 2009, Jones et al. 2013, Wilder et al. 2013).

3.3.2 Priors can be encoded by other DDM parameters

Our baseline model sets the parameter vector $\theta_i = [M, a, c_i, T, M_{\sigma^2}, T_{\sigma^2}, c_{\sigma^2}]$ which implies that only the initial point is allowed to vary across the C experimental conditions. This parameter restriction assumes that if prior probabilities vary across conditions, they must be encoded in the initial point. This baseline model is motivated by recent work showing that variation in prior beliefs is explained by variation in initial points, instead of by drift rates (Mulder et al. 2012, White and Poldrack 2014).

However, the debate on whether priors are encoded in the initial point or the drift rate is still ongoing (Rorie et al. 2010, Gao et al. 2009, Van Ravenzwaaij et al. 2012) and recent work has developed experimental paradigms to examine this particular issue (Mulder et al. 2012, Teodecuru and Usher 2013). Because this remains an open question in the literature, we estimated several additional versions of the DDM to investigate whether allowing drift rates to vary across experimental conditions yielded a better fit to our data. When comparing the baseline model to a model where both drift rates and initial points are allowed to vary, $\theta_i = [M_i, a, c_i, T, M_{\sigma^2}, T_{\sigma^2}, c_{\sigma^2}]$, we find that the baseline model performs better in 37 of 38 subjects according to the Akaike Information Criterion (AIC). Furthermore, when comparing the baseline model to a model where only the drift rate is allowed to vary, we find that the baseline model performs better in 31 of 38 subjects.

Another analysis that can help distinguish whether prior bias is encoded in the drift rate or initial point investigates the difference in average RTs on correct and incorrect trials. In particular, we followed Mulder et al. (2012) and defined (i) a valid prior trial as a trial where

either a stimulus repetition had followed two or more repetitions, or an alternation had followed two or more alternations; (ii) an invalid prior trial as a trial where either stimulus alternation had followed two or more repetitions, or repetition had followed two or more alternations; and (iii) a neutral prior trial as all other trials. The key prediction is that on valid trials where subject respond incorrectly, RTs should be fast if prior bias is encoded in the drift rate and slow if prior bias is encoded in the initial point. In line with this hypothesis, we found that when making correct responses, subjects were fastest in trials with valid priors compared to neutral priors, and slowest in trials with invalid priors; critically, the opposite pattern was found for incorrect trials (see Figure S4). These results support the notion that for the PDT in our setting, changes in the prior are better captured by changes in the initial point of evidence accumulation rather than changes in the drift rate.

Finally, although previous research on sequential effects in binary perceptual decision-making tasks often sets the number of conditions to C=16 (Cho et al. 2002, Yu and Cohen 2009), this itself is a parameter that is set by the researcher. To investigate this parameter choice, we re-estimated the baseline DDM allowing *C* to vary over the set [2, 4, 8, 16] by collapsing the number of repetition and alternation histories accordingly. We find that the baseline DDM that uses 16 conditions performs better in 35 of 38 subjects (according to AIC) than the baseline DDM that use 2, 4 or 8 conditions. Similarly, when comparing a model where only the drift rate is allowed to vary across conditions, the model with 16 conditions performs best according to the AIC in 36 of 38 subjects. Taken together, the baseline model where only initial points are allowed to vary across 16 experimental conditions is the best fit, and we therefore use its estimated parameters in all subsequent analyses.

3.4 Decoding Prior Beliefs Using the DDM

Using the estimated DDM parameters, we decoded the prior probability for each subject and condition. To understand how the decoding process works, recall that the drift rate of the DDM encodes the informativeness of the sensory signals in discriminating between the two shapes. At every instant within a trial, a new noisy signal is sampled where the noise is governed by the volatility of the process, s^2. When volatility increases, the sensory signals become noisier and a subject relies more heavily on his prior. In the limit, when the volatility goes to infinity, the subject relies exclusively on his prior. In the E-Companion, we analytically solve for the probability of hitting the upper boundary when the volatility goes to infinity, and find that, in condition *i,* this prior equals $\frac{c_i}{a}$.

Using this analytical result, we then collapsed the data from 16 to 8 conditions, because the identity of the current trial does not vary with the prior (e.g., trials in condition AAAR and trials in condition AAAA provide information only about $Pr(A \mid AAA) = 1 - Pr(R \mid AAA)$). Figure 4 plots the average priors for each of the eight conditions (all possible three element histories) for both the PDT and EDT. The figure shows that prior probabilities are indeed a function of the recent stimulus history, and that these priors are highly correlated across the PDT and EDT (r =0.90, p<0.005).

3.5 Individual differences in belief formation across tasks

Using our decoding strategy, we found significant individual differences in the extent to which priors in the PDT deviated from the rational prior of 0.5. To quantify this deviation, for each subject u, we computed the sum of squared deviations from 0.5 and define this as the *irrationality index (II)*:

$$(3) \qquad II_u = \frac{1}{8}\sum_{i=1}^{8}(p_{i,u} - 0.5)^2.$$

Because subjects were explicitly told that the probability of seeing either shape was 0.5, independent of the stimulus history, a rational subject would exhibit an irrationality index of 0 in the PDT. Instead, we found that the average *II* across subjects was *1.93* (SD: 0.391), which is significantly greater than the optimal level of *0* (p<0.001, Tobit regression left-censored at 0).

If the extrapolative beliefs from the EDT are driven by the same psychological mechanism that generated the irrationality in the PDT, then the *II* should explain a portion of the cross-subject heterogeneity in the extrapolation of beliefs in the EDT (Appelt et al. 2011). To test this, we defined an *Extrapolation Index* (*EI*) for each subject u, as

$$(4) \qquad EI_u = \frac{1}{400}\sum_{t=1}^{400}(b_{t,u} - 0.5)^2,$$

where $b_{t,u}$ is the belief reported by subject u in trial t that a repetition would occur on trial $t+1$. As illustrated in Figure 5A, we found a significant positive correlation of 0.57 between the *II* and the *EI* (p<0.001). To verify that the *II* was not simply capturing noise in the RTs, we regressed the *EI* on the *II* and an individual subject measure of the standard deviation of RTs. We found

that after controlling for the standard deviation of RTs, the *II* was still a significant predictor of the *EI* (p<0.001).

Finally, we tested whether the irrationality index of the PDT was predictive of the *payoffs* in the EDT. We found a significant negative correlation between the *II* and the payoff in the EDT, demonstrating that irrational behavior in the perceptual domain predicts performance in the economic domain (r= -0.47, p<0.01, Figure 5B).
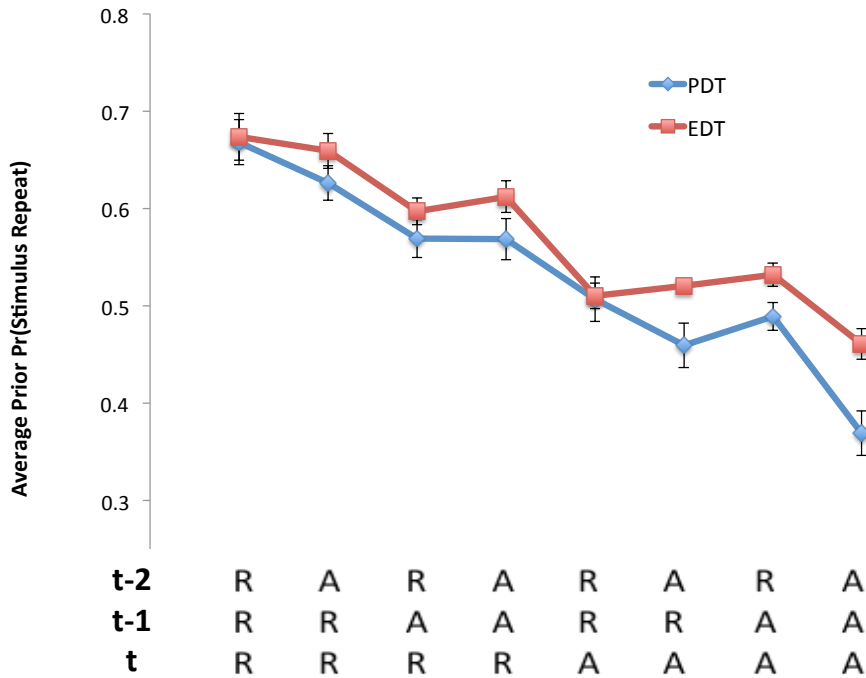
Figure 4. Average priors for the EDT and PDT as a function of the four most recent stimuli. For each of the eight different conditions, the red line shows the average belief that a repetition will occur on the subsequent trial in the EDT, elicited using the BDM procedure. The blue line shows the average prior that a repetition will occur on the subsequent trial in the PDT, decoded from the initial point and boundary parameters of the DDM.

A

B
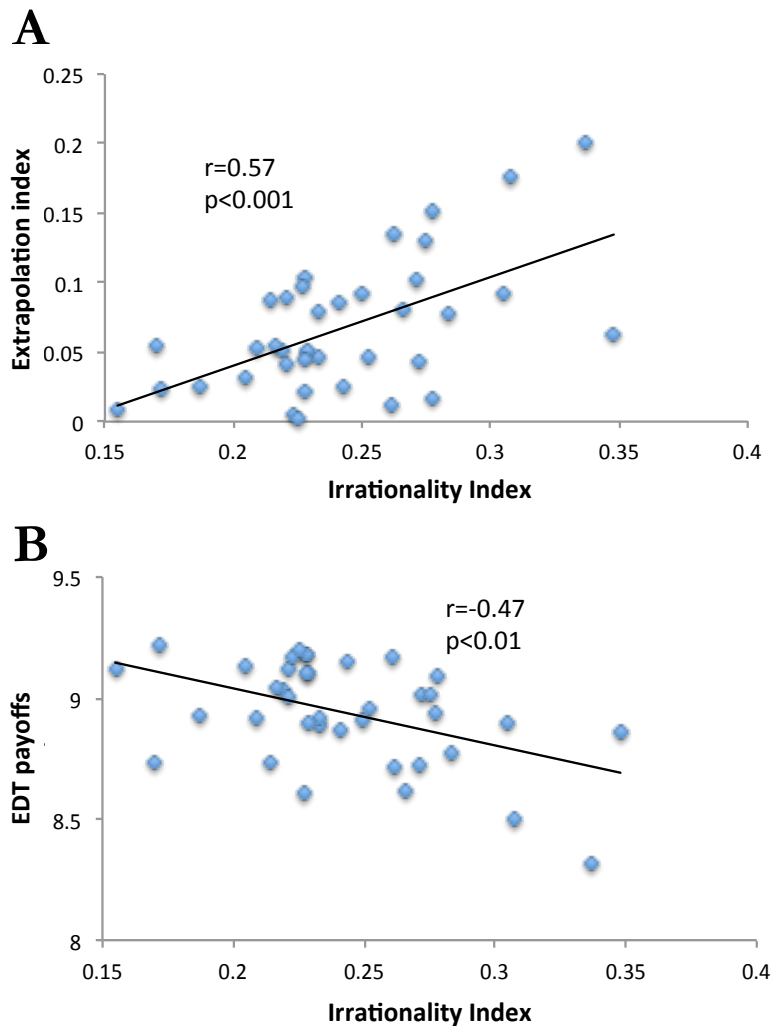
Figure 5. Individual differences. (A) Correlation across subjects between the extrapolation index and the irrationality index. (B) Correlation across subjects between the irrationality index and the EDT payoffs. Each point represents a single subject.

3.6 A Common Dynamic Belief Model

So far we have shown that (i) conditional on recent stimulus history, the average beliefs in the EDT and PDT are correlated (Figure 4) and (ii) individual differences in the degree to which beliefs deviate from the 0.5 benchmark in the EDT and PDT are correlated across subjects (Figure 5A). These two pieces of evidence suggest that a common belief formation mechanism governs subjects' responses in both tasks. To further investigate this conjecture, we use a model

from the computational neuroscience literature, the *Dynamic Belief Model* (DBM) (Yu and Cohen 2009, Wilder et al. 2010, Zhang et al. 2014). The DBM relies on two key assumptions. First, the agent believes that the probability of observing a repetition on trial t (i.e., the stimulus on trial t matches the stimulus on trial t-1) is governed by a parameter γ_t. Second, the parameter γ_t is time-varying and changes on each trial with a constant probability α. These assumptions are motivated by empirical evidence in cognitive neuroscience, showing that the brain is well adapted to learning in non-stationary environments (Behrens et al. 2007, Nassar et al. 2012). Furthermore, the DBM can be well approximated by an exponential filter (Yu and Cohen 2009), which has received empirical support both at the behavioral and neurophysiological levels (Sugrue et al. 2004).

To fix ideas, we describe the model in the context of the PDT. Let $X_t \in$ $\{Square, Circle\}$, and suppose the agent believes that the state of the world is captured by γ_t, which represents the time-varying repetition rate, $Pr_t\ (X_t = X_{t-1})$. On each trial, with probability α, γ_t is resampled from a reset prior $P_0(.)$ that is uniform over $[0, 1]$ (Fox and Rottenstreich, 2003). Formally, let $z_t = 1$ if X_t=*square*, and let $z_t = 0$ if X_t=*circle*, so that the probability that the next instance, X_t, is a square is as follows:

$$(5) \quad Pr(X_t = square | X_{t-1}, X_{t-2}, \dots, X_1) =$$
$$Pr(X_t = square | \gamma_t, X_{t-1}) = z_{t-1}\gamma_t + (1 - z_{t-1})(1 - \gamma_t).$$

The model's prediction depends explicitly only on the most recent observation, X_{t-1} and on the current estimate of γ_t. The DBM algorithm operates iteratively by maintaining a prior distribution over γ_t, $Pr\ (\gamma_t | X_{t-1}, \gamma_{t-1})$. After observing a new stimulus, the posterior, $\hat{P}\ (\gamma_t | X_t, \gamma_{t-1})$ is computed using Bayesian updating:

$$(6) \qquad \hat{P}\ (\gamma_t | X_t) \propto P\ (X_t | \gamma_t, X_{t-1}) P(\gamma_t | X_{t-1}).$$

The posterior of the current trial is then used to compute the prior for the next trial, as a sum of the posterior weighted by $(1 - \alpha)$ and the reset prior weighted by α:

$$(7) \qquad P(\gamma_{t+1} | X_t) = (1 - \alpha)\ \hat{P}\ (\gamma_t | X_t) + \alpha P_0\ (\gamma_{t+1}).$$

The model generates predictions, $P\ (X_t | X_{t-1})$, by integrating Eq. 6 over the prior on γ_t. In our simulations, we maintain a discrete approximation to the continuous prior by dividing the interval

[0,1] into 100 equally spaced bins, where expectations are computed by summing over the discrete probability mass function.

To test whether the DBM provides a common computational model of belief formation across both the EDT and PDT, we estimated α at the subject level for the PDT, and found that the mean level was $\alpha* = 0.44$. We then used $\alpha*$ to calibrate a DBM for the EDT, and computed the DBM predicted time series of beliefs for the EDT. Figure 6 shows the DBM theoretical predictions plotted against the actual average beliefs from the EDT. The two time series exhibit a strong positive correlation (r=0.66 p<0.001)[5]. Moreover, we also find that the mean level of α from the EDT is 0.46, which is nearly identical to the mean level of α from the PDT of 0.44. Figure S5 shows the sum of squared errors from the EDT as a function of α, and indicates that the global minimum is indeed at 0.44.
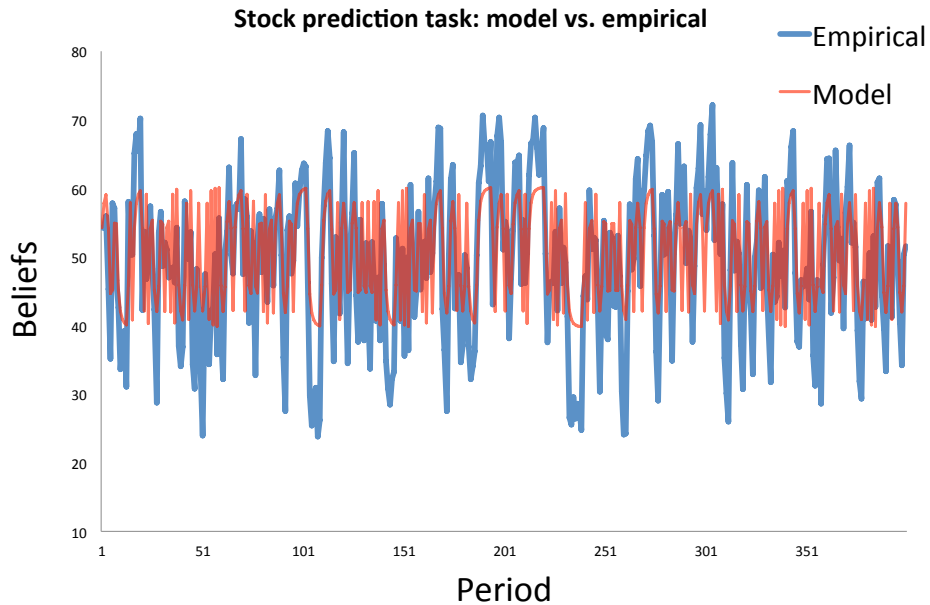
Figure 6. Out of sample DBM predictions of the average EDT beliefs. The predictions are calibrated from the PDT priors. A DBM model is estimated for each subject from the PDT. We then input the stimuli from the EDT into the DBM using the average α across subjects from the PDT, and generate a time series of theoretical predictions (red). The empirical average beliefs from the EDT are plotted in blue. The correlation between the two time series is 0.66 (p<0.001).

---

[5] We ran an additional analysis where we first computed the thirty-eight predicted time series for the EDT using the thirty-eight individually estimated alphas, and then averaged these time series. This average time series exhibits a 0.686 correlation with the actual average time series from the EDT, thus serving as a robustness check.

4. Discussion

Our experimental results provide support for the hypothesis that belief formation in perceptual and economic decision-making is governed, at least in part, by a common psychological mechanism that is rooted in Bayesian models of decision-making (Chater et al. 2006, Oaksford and Chater 2009). Interestingly, the use of the DBM in the PDT is not optimal, because subjects are explicitly told the data generating process, and hence there is no reason to learn the underlying model parameters. When viewed separately, the results from the PDT and EDT are successful replications of several studies that have employed similar tasks, as RTs in the PDT and predictions in the EDT are heavily dependent on recent stimulus history (Bloomfield and Hales 2002, Cho et al. 2002, Huettel et al. 2002, Asparouhova et al. 2009). However, by employing a within-subject design, we were also able to compare the computational processes underlying belief formation across the tasks.

Because subjects were explicitly told that the probability of seeing a square on each trial in the PDT was 0.5, an optimal Bayesian agent using the DBM would set α equal to zero 0 (i.e., the autocorrelation parameter would change with 0 probability on each trial). Instead, we found the average level of α in our subject pool was 0.44. This suggests that subjects may have a strong prior on a non-stationary model of the world, and they have difficulty adjusting their belief formation mechanism in the face of explicit information that would increase earnings (Yu and Cohen 2009). Moreover, the average value of α in the EDT was 0.46, nearly identical to the parameter's average value in the PDT (see Figure S5). This similarity in α rules out the possibility that working memory constraints are responsible for subjects' strong reliance on recent stimulus history; because we display the previous fourteen stimuli in the EDT and we display no previous stimuli in the PDT, we would expect differences in average α if behavior was driven by working memory constraints, but we do not find this.

It is important to highlight that the combination of the DDM with the DBM is just one of many models that can be used to explain sequential effects in perceptual decision-making at long RSIs. For example, another candidate model is proposed by Gao et al. (2009) and allows for a unified explanation of both within and across trial dynamics over a broad range of RSIs. While this model surely provides a generalized and detailed account of the sequential effects in our data, we choose to model our data with a joint DBM-DDM because the DBM is a portable model that can be applied directly to many other settings, including the EDT. As the core question in our paper is about the relationship between belief formation across domains, the ability to use the

same single parameter model across separate tasks is especially useful. Moreover, the DDM is increasingly being used to model economic decisions, and we therefore believe it is valuable to use a model that has been shown to explain data well in both perceptual and economic decision-making tasks (Fehr and Rangel 2011).

Furthermore, the key feature of the DBM that enables it to flexibly explain sequential effects is the assumption that the autocorrelation parameter is perceived to be time varying, and that subjects use Bayesian inference to update their estimate of this parameter. While it may be difficult for subjects to implement these precise Bayesian computations, it has been shown theoretically that the DBM is well approximated by an exponential filter (Yu and Cohen 2009), which has received empirical support both at the behavioral and neurophysiological levels (Sugrue et al. 2004). This property is especially relevant to the current study because of recent theories of belief formation in economics and finance that explicitly use this exponential decay property (Malmendier and Nagel 2009, 2015, Barberis et al. 2015).

At a methodological level, one contribution of our study is to provide the analytical framework that enabled us to measure the trial-by-trial subjective beliefs in the PDT. In particular, we develop a new methodology that can be used in future work on perceptual decisions to elicit prior probabilities. This is valuable because of the possibility that subjects themselves may not have a conscious representation of these priors, thus making it difficult to directly measure them (Curran and Keele 1993). While most previous DDM research has focused on estimating computational parameters as a function of experimental conditions, we instead use the estimated computational parameters to reverse-engineer the belief formation process. While reaction time and accuracy data have been used for decades to infer mental representations of the environment (Luce 1986, Achtziger and Alós-Ferrer 2013, Jones et al. 2013), our study provides the extra structure necessary to map RTs into exact probabilities, thus providing a common framework for studying beliefs in perceptual and economic decisions.

One potential concern about our experimental design is that the EDT and PDT are not substantially different from each other, and therefore it not surprising that we find evidence for a common belief formation process across the two tasks. Indeed, the stimuli in both the EDT and PDT are binary processes and both tasks require making judgments about the underlying state of the world. However, the tasks do differ on three fundamental dimensions. First, in the EDT subjects are asked to make judgments about the likelihood of a future event (a firm's performance) whereas in the PDT subjects are asked to make judgments about the identity of the

currently displayed stimulus. In other words, the EDT relies on cognitive resources to predict the future whereas the PDT relies on perceptual resources to classify the present state of the world. Another major difference between the two tasks is the information that is given to subjects about the underlying data generating process. In the PDT subjects are explicitly informed that the probability of observing each stimuli is 0.5, whereas subjects are not told anything about the process governing the stimuli in the EDT. Finally, the tasks differ based on the information that is available to subjects about the stimulus history. In the EDT, subjects are given access on-screen to a history of the previous fourteen stimuli; in the PDT, only the current stimulus is displayed.

One could also argue that because of these different aspects across the two tasks, the individual differences in behavior that we document may be driven by two distinct mechanisms. Perhaps the strongest alternative hypothesis is that the sequential effects we observe in the PDT are not driven by expectation formation, but are instead driven by post-response residual activity of the motor cortex. This type of effect has indeed been documented in the sequential effects literature, but it is most prevalent in tasks that feature short RSIs on the order of 250 milliseconds or less (Soetens et al. 1985). Because the goal of this paper is to examine expectation formation explicitly, we designed the PDT with a long RSI of 800 milliseconds, for which expectation formation effects have been shown to dominate lower level automatic facilitation effects that are driven by post-response residual activity (Gao et al. 2009).

Our results also relate to the literature on belief-based biases in judgment and decision-making and behavioral economics. Numerous studies have documented that after seeing a sequence of identical stimuli (e.g. successful basketball shots or positive stock returns), humans have a tendency to extrapolate the past and believe that the streak will continue (Gilovich et al. 1985, Greenwood and Shleifer 2014). The origin of these extrapolative beliefs – also known as the "hot hand fallacy" -- is still not completely understood, and multiple models have been proposed to explain it and its implications for financial markets (Rabin 2002, Massey and Wu 2005, Asparouhova et al. 2009, Oskarsson et al. 2009, Barberis et al. 1998, 2015). Because of the common computational processes and the similar belief in continuing streaks across the two tasks, our results suggest that the extrapolative beliefs in economic decision-making may stem from low-level perceptual processes instead of deliberative top-down judgments. One possible interpretation is that the inability to maintain a constant initial point across trials in the PDT is driven by the fact that it is optimal to flexibly change the initial point from decision to decision in many other environments. This inability to maintain a constant initial point may then be inherited

by the belief-formation mechanism deployed in economic decision-making, though we emphasize this is just a conjecture.

Because subjects are explicitly told that the data generating process in the PDT does not contain any predictability, these perceptual pattern recognition processes may be hard-wired and difficult to suppress. Understanding whether extrapolative beliefs, and judgment biases in general, arise from deliberative top-down decisions or from bottom-up perceptual processes is important because the distinction has implications for policy-makers. If some biases are driven by low-level processes, as our data suggests, then a policy that mandates additional information disclosure may not be effective at impacting decision-making (Brav and Heaton 2002). While we examine only one specific decision bias in this study, future work may benefit from linking perceptual and economic decision-making via computational models to better understand the origins of other well-known behavioral biases (Tsetsos et al. 2012).

References:

Achtziger, Anja, and Carlos Alós-Ferrer, 2013, Fast or rational? A response-times study of bayesian updating, *Management Science* 60, 923-938.

Appelt, Kirstin C., Kerry F. Milch, Michel J. J. Handgraaf, and Elke U. Weber, 2011, The decision making individual differences inventory and guidelines for the study of individual differences in judgment and decision-making research, *Judgment and Decision Making* 6, 252-262.

Asparouhova, Elena, Michael Hertzel, and Michael Lemmon, 2009, Inference from streaks in random outcomes: Experimental evidence on beliefs in regime shifting and the law of small numbers, *Management Science* 55, 1766-1782.

Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer, 2015, X-CAPM: An extrapolative capital asset pricing model, *Journal of Financial Economics*, 115, 1-24.

Barberis, Nicholas, Andrei Shleifer, Robert Vishy, 1998, A model of investor sentiment, *Journal of Financial Economics* 49, 307-343.

Becker, Gordon M, Morris H DeGroot, and Jacob Marschak, 1964, Measuring utility by a single response sequential method, *Behavioral Science* 9, 226-232.

Behrens, Timothy EJ, Mark W Woolrich, Mark E Walton, and Matthew FS Rushworth, 2007, Learning the value of information in an uncertain world, *Nature Neuroscience* 10, 1214-1221.

Bloomfield, Robert, and Jeffrey Hales, 2002, Predicting the next step of a random walk: Experimental evidence of regime-shifting beliefs, *Journal of Financial Economics* 65, 397-414.

Bogacz, Rafal, Eric Brown, Jeff Moehlis, Philip Holmes, and Jonathan D Cohen, 2006, The physics of optimal decision making: A formal analysis of models of

performance in two-alternative forced-choice tasks, *Psychological Review* 113, 700-765.

Brainard, David H, 1997, The psychophysics toolbox, *Spatial Vision* 10, 433-436.

Brav, Alon, and John B Heaton, 2002, Competing theories of financial anomalies, *Review of Financial Studies* 15, 575-606.

Chater, Nick, Joshua B. Tenenbaum, and Alan Yuille, 2006, Probabilistic models of cognition: Conceptual foundations, *Trends in Cognitive Sciences* 10, 287-291.

Cho, Raymond Y, Leigh E Nystrom, Eric T Brown, Andrew D Jones, Todd S Braver, Philip J Holmes, and Jonathan D Cohen, 2002, Mechanisms underlying dependencies of performance on stimulus history in a two-alternative forced-choice task, *Cognitive, Affective, & Behavioral Neuroscience* 2, 283-299.

Curran, Tim, and Steven W Keele, 1993, Attentional and nonattentional forms of sequence learning, *Journal of Experimental Psychology: Learning, Memory, and Cognition* 19, 189-202.

Fehr, Ernst, and Antonio Rangel, 2011, Neuroeconomic foundations of economic choice—recent advances, *The Journal of Economic Perspectives* 25, 3-30.

Fox, Craig R., and Yuval Rottenstreich, 2003, Partition priming in judgment under uncertainty, *Psychological Science* 14, 195-200.

Gao, Juan, KongFatt Wong-Lin, Philip Holmes, Patrick Simen, and Jonathan D Cohen, 2009, Sequential effects in two-choice reaction time tasks: Decomposition and synthesis of mechanisms, *Neural Computation* 21, 2407-2436.

Gilovich, Thomas, Robert Vallone, and Amos Tversky, 1985, The hot hand in basketball: On the misperception of random sequences, *Cognitive Psychology* 17, 295-314.

Gold, Joshua I. , and Hauke R. Heekeren, 2013, Neural mechanismilliseconds for perceptual decision making, *Neuroeconomics: Decision Making and the Brain*.

Greenwood, Robin and Andrei Shleifer, 2014, Expectations of returns and expected returns, *Review of Financial Studies* 27, 714-746.

Huettel, Scott A, Peter B Mack, and Gregory McCarthy, 2002, Perceiving patterns in random series: Dynamic processing of sequence in prefrontal cortex, *Nature Neuroscience* 5, 485-490.

Hong, Harrison, and Jeremy Stein, 1999, A unified theory of underreaction, momentum trading and overreaction in asset markets, *Journal of Finance* 54, 2143-2184.

Jones, Matt, Tim Curran, Michael C Mozer, and Matthew H Wilder, 2013, Sequential effects in response time reveal learning mechanisms and event representations, *Psychological Review* 120, 628-666.

Klauer, Karl Christoph, Andreas Voss, Florian Schmitz, and Sarah Teige-Mocigemba, 2007, Process components of the implicit association test: A diffusion-model analysis, *Journal of Personality and Social Psychology* 93, 353-368.

Krajbich, Ian, and Antonio Rangel, 2011, Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions, *Proceedings of the National Academy of Sciences* 108, 13852-13857.

Luce, Robert Duncan, 1986. *Response times: Their role in inferring elementary mental organization3* (Oxford University Press).

Malmendier, Ulrike, and Stefan Nagel, 2011, Depression babies: Do macroeconomic experiences affect risk-taking?, *Quarterly Journal of Economics* 126, 373-416.

Malmendier, Ulrike, and Stefan Nagel, 2015, Learning from inflation experiences, *Quarterly Journal of Economics*, forthcoming.

Massey, Cade, and George Wu, 2005, Detecting regime shifts: The causes of under-and overreaction, *Management Science* 51, 932-947.

Mulder, Martijn J, Eric-Jan Wagenmakers, Roger Ratcliff, Wouter Boekel, and Birte U Forstmann, 2012, Bias in the brain: A diffusion model analysis of prior probability and potential payoff, *The Journal of Neuroscience* 32, 2335-2343.

Nassar, Matthew R, Katherine M Rumillisecondsey, Robert C Wilson, Kinjan Parikh, Benjamin Heasly, and Joshua I Gold, 2012, Rational regulation of learning dynamics by pupil-linked arousal systemilliseconds, *Nature Neuroscience* 15, 1040-1046.

Oaksford, Mike, and Nick Chater, 2009, Précis of bayesian rationality: The probabilistic approach to human reasoning, *Behavioral and Brain Sciences* 32, 69-84.

Oskarsson, An T, Leaf Van Boven, Gary H McClelland, and Reid Hastie, 2009, What's next? Judging sequences of binary events, *Psychological Bulletin* 135, 262-285.

Platt, Michael L, and Scott A Huettel, 2008, Risky business: The neuroeconomics of decision making under uncertainty, *Nature Neuroscience* 11, 398-403.

Polanía, Rafael, Ian Krajbich, Marcus Grueschow, and Christian C Ruff, 2014, Neural oscillations and synchronization differentially support evidence accumulation in perceptual and value-based decision making, *Neuron* 82, 709-720.

Rabin, Matthew, 2002, Inference by believers in the law of small numbers, *Quarterly Journal of Economics* 117, 775-816.

Rangel, Antonio , Colin Camerer, and P. Read Montague, 2008, A framework for studying the neurobiology of value-based decision-making, *Nature Reviews Neuroscience* 9, 545-556.

Ratcliff, Roger, and Gail McKoon, 2008, The diffusion decision model: Theory and data for two-choice decision tasks, *Neural Computation* 20, 873-922.

Ratcliff, Roger, and Jeffrey N Rouder, 1998, Modeling response times for two-choice decisions, *Psychological Science* 9, 347-356.

Ratcliff, Roger, and Philip L Smith, 2004, A comparison of sequential sampling models for two-choice reaction time, *Psychological Review* 111, 333-367.

Roitman, Jamie D, and Michael N Shadlen, 2002, Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task, *Journal of Neuroscience* 22, 9475-9489.

Rorie, Alan E, Juan Gao, James L McClelland, and William T Newsome, 2010, Integration of sensory and reward information during perceptual decision-making in lateral intraparietal cortex (LIP) of the macaque monkey, *PloS One* 5, e9308.

Soetens, Eric, LC Boer, and JE Hueting, 1985, Expectancy or automatic facilitation? Separating sequential effects in two-choice reaction time, *Journal of Experimental Psychology: Human Perception and Performance* 11, 598-616.

Sugrue, Leo P, Greg S Corrado, and William T Newsome, 2004, Matching behavior and the representation of value in the parietal cortex, *Science* 304, 1782-1787.

Summerfield, Christopher, and Konstantinos Tsetsos, 2012, Building bridges between perceptual and economic decision-making: Neural and computational mechanisms, *Frontiers in Neuroscience* 6, 70.

Summerfield, Christopher, and Konstantinos Tsetsos, 2015, Do humans make good decisions?, *Trends in Cognitive Sciences* 19, 27-34.

Teodorescu, Andrei R, and Marius Usher, 2013, Disentangling decision models: From independence to competition, *Psychological Review* 120, 1-38.

Towal, R. Blythe, Milica Mormann, and Christof Koch, 2013, Simultaneous modeling of visual saliency and value computation improves predictions of economic choice, *Proceedings of the National Academy of Sciences* 110, 3858-3867.

Townsend, James T, and F Gregory Ashby, 1983. *The stochastic modeling of elementary psychological processes* (CUP Archive).

Trueblood, Jennifer S, Scott D Brown, Andrew Heathcote, and Jerome R Busemeyer, 2013, Not just for consumers context effects are fundamental to decision making, *Psychological Ccience* 24, 901-908.

Tsetsos, Konstantinos, Nick Chater, and Marius Usher, 2012, Salience driven value integration explains decision biases and preference reversal, *Proceedings of the National Academy of Sciences* 109, 9659-9664.

Tuerlinckx, Francis, and Joachim Vandekerckhove, 2008, Diffusion model analysis with matlab: A DMAT primer, *Behavior Research Methods* 40, 61-72.

Usher, Marius, and James L McClelland, 2001, The time course of perceptual choice: The leaky, competing accumulator model, *Psychological Review* 108, 550-592.

Van Ravenzwaaij, Don, Martijn J Mulder, Francis Tuerlinckx, and Eric-Jan Wagenmakers, 2012, Do the dynamics of prior information depend on task context? An analysis of optimal performance and an empirical test, *Frontiers in Psychology* 3.

Webb, Ryan, 2013, Dynamic constraints on the distribution of stochastic choice: Drift diffusion implies random utility, *Working paper.*

White, Corey N, and Russell A Poldrack, 2014, Decomposing bias in different types of simple decisions, *Journal of Experimental Psychology: Learning, Memory, and Cognition,* 40, 385-398

Wilder, Matthew H, Matt Jones, and Michael Mozer, 2010, Sequential effects reflect parallel learning of multiple environmental regularities, NIPS.

Wilder, Matthew, Matt Jones, Alaa Ahmed, Tim Curran and Michael Mozer, 2013, The persistent impact of incidental experience*, Psychonomic Bulletin & Review,* 20, 1221-1231.

Woodford, Michael, 2014, Stochastic choice: An optimizing neuroeconomic model, *American Economic Review* 104, 495-500.

Yu, J Angela, and Jonathan D Cohen, 2009, Sequential effects: Superstition or rational behavior?, NIPS.

Zhang, Shunan, Crane He Huang, and J Angela Yu, 2014, Sequential effects: A bayesian analysis of prior bias on reaction time and behavioral choice, Proceedings of the Cognitive Science Society Conference.

E-Companion

**Extrapolative Beliefs in Perceptual and Economic Decisions:
Evidence of a Common Mechanism**

This E-companion contains supplementary material. Section 1 includes supplementary information on preprocessing the reaction time data and on the ordering of tasks in the experimental sessions. Section 2 provides robustness checks and extended statistical tests of the basic results from the PDT and EDT. Section 3 provides details on derivation of the prior probability from the DDM parameters. Section 4 contains the full set of instructions given to our experimental subjects.

1. Data preprocessing and order of experimental tasks

1.1 Preprocessing of reaction time data. The PDT consisted of 4 blocks of 300 trials each. RTs and error rates systematically increased over the course of each block, likely due to subjects' fatigue (See Fig. S1 and S2). To control for this, we removed a linear time trend, within each block, for each subject. All results and analyses in the text use this de-trended RT data and are robust to exclusion of the de-trending step.

1.2 On the ordering of the tasks in the experimental session. In our within-subjects design, the ordering of the tasks was not randomized and the PDT always took place first. We began the experiment with the PDT for all subjects because during pilot testing, we observed a sharp fatigue effect in subjects' response times in this task (see Fig. S1-S3). We were therefore concerned that this fatigue effect would vary between subjects if the PDT was administered later in the session for a subset of subjects. Furthermore, because subjects in the PDT were explicitly informed that the probability of seeing either shape was 0.5, we believe that any possible spillover effects between the two tasks should bias us against finding the extrapolation effect observed in the EDT (where subjects were not explicitly informed about the underlying random process).

2. Robustness checks and extended statistical tests

The following matrices summarize the pairwise t-statistics that result from comparing mean RTs presented in Figures 2A and 2B (Each subject mean is used as single observation).

2.1 Difference in mean continuation response times at different streak lengths

Continuation	1	2	3	4	5	6
1		t(37)=9.48 p<0.001	t(37)=10.53 p<0.001	t(37)=10.27 p<0.001	t(37)=6.99 p<0.001	t(37)=8.81 p<0.001
2			t(37)=7.99 p<0.001	t(37)=5.99 p<0.001	t(37)=4.15 p<0.001	t(37)=5.95 p<0.001
3				t(37)=2.09 p=0.04	t(37)=1.63 p=0.11	t(37)=3.05 p=0.004
4					t(37)=0.27 p=0.79	t(37)=1.34 p=0.19
5						t(37)=1.08 p=0.29
6						

2.2 Difference in mean error rates at different streak lengths

Continuation	1	2	3	4	5	6
1		t(37)=7.20 p<0.001	t(37)=12.91 p<0.001	t(37)=9.99 p<0.001	t(37)=10.42 p<0.001	t(37)=-7.22 p<0.001
2			t(37)=-3.25 p<0.002	t(37)=-3.95 p<0.001	t(37)=-4.98 p<0.001	t(37)=-3.84 p<0.001
3				t(37)=-1.49 p=0.14	t(37)=-3.09 p=0.004	t(37)=-1.98 p=0.05
4					t(37)=-1.15 p=0.26	t(37)=-1.31 p=0.19
5						t(37)=-0.22 p=0.83
6						

2.3 Difference in mean response times between continuation and violation trials

Streak	1	2	3	4	5	6
	t(37)=0.45 p=0.65	t(37)=3.95 p<0.001	t(37)=5.86 p<0.001	t(37)=6.41 p<0.001	t(37)=6.56 p<0.001	t(37)=7.19 p<0.001

2.4 Difference in mean error rates between continuation and violation trials

Streak	1	2	3	4	5	6
	t(37)=1.02 p=0.31	t(37)=3.65 p<0.001	t(37)=4.61 p<0.001	t(37)=4.94 p<0.001	t(37)=5.36 p<0.001	t(37)=4.95 p<0.001

3. Derivation of Prior Decoding Technique

The equation that we use to derive the prior probability from the DDM parameters is based on the basic DDM, but it can be extended to the full DDM in a straightforward manner by integrating the original equation against the across trial variability parameters. In particular, we start from equation A3 from Ratcliff and Smith (2004), which provides the probability of responding to the upper boundary for the basic DDM. In our setting, this corresponds to the probability of responding to the "correct" alternative:

$$q^{basic}(c, a, b) = \frac{e^{\frac{-2Mc}{s^2}} - e^{\frac{-2Mb}{s^2}}}{e^{\frac{-2Ma}{s^2}} - e^{\frac{-2Mb}{s^2}}}$$

In this expression, c represents the initial point which always lies in between the lower boundary, b, and the upper boundary, a. The prior probability of choosing the correct response can be computed by assuming that the within trial noise is arbitrarily large, in which case any information that the stimulus provides is dominated by the within trial noise. Using the fact that $e^x \approx (1 + x)$ as $x \to 0$, it can be shown that:

$$\lim_{s^2 \to \infty} q^{basic}(c, a, b) = \frac{c - b}{a - b}$$

Using this result, we can then compute the prior probability of crossing the upper boundary under the extended DDM by integrating $q^{basic}(c, a, b)$ against the across trial variability parameters, and then allowing s^2 to become arbitrarily large. In particular, the probability of crossing the upper boundary under the extended DDM is:

$$q^{advanced}(c, a, b) = \iiint q^{basic}(c, a, b) f(c) g(M) h(T) \, dc dM dT,$$

where f, g, and h represent the probability density functions of the initial point, drift rate, and non-decision time, respectively. To compute the prior probability, we take the limit as the within trial noise goes to infinity:

$$\lim_{s^2 \to \infty} q^{advanced}(c, a, b) = \lim_{s^2 \to \infty} \iiint q^{basic}(c, a, b) f(c) g(M) h(T) \, dc dM dT$$

$$= \iiint \lim_{s^2 \to \infty} q^{basic}(c, a, b) f(c) g(M) h(T) \, dc dM dT$$

$$= \iiint \frac{c - b}{a - b} f(c) g(M) h(T) \, dc dM dT$$

$$= \int \frac{c - b}{a - b} f(c) dc$$

$$= \frac{c - b}{a - b}$$

The second equality is justified by the Dominated Convergence Theorem, which allows us interchange the order of the limit and the integration. The fifth equality is based on the facts that c is uniformly distributed and the integrand $\frac{c - b}{a - b}$, is linear in c. Finally, for the case at hand, since we set the lower boundary b to zero, the prior probability of reaching the upper boundary under the extended DDM is given by $\frac{c}{a}$.

4. Instructions

Thank you for participating in this experiment.
For your participation you have already made $5. During the rest of the experiment you have the chance to make more money. Your final payoff for participating depends on your decisions in Parts I, II.

Part I:

In this part of the task, you will see a sequence of shapes; each element of the sequence will be displayed one at a time. There are only two possible shapes: a white circle and a white square. Your task is to accurately classify which shape is currently being presented, as quickly as possible. If you see the **circle**, press the **right** arrow button, and if you see the **square**, press the **left** arrow button. The trials will be broken up into 4 separate blocks of 300 trials; after each block, you will have a 20 seconds break. Please use only one hand to enter both buttons.
For every shape you correctly classify, you will be paid 1 cent. If you classify all the shapes correctly, you will make 1200*0.01=$12.00. However, for every 0.05 second it takes you to respond, you will lose 0.1 cents. (you will have a maximum of 2 seconds/trial to respond). **Therefore, to make the most money possible, you should answer as quickly and as accurately as you can.**

In each trial, the chance that you will see a circle is ½, and the chance that you will see a square is ½. Shapes on previous trials have no influence on the shape in the current trials; in other words the shape you see on the current trial is completely independent of all other shapes you've already seen.
Before the real task starts, you will start with 5 practice trials.

Part II (instructions were given only after part I was completed)

We have studied large numbers of publicly traded companies, and constructed models of their performance patterns. Using these models, we created sequences to represent patterns of "surprises" (actual performance minus predicted performance). An upward movement indicates a "positive surprise," which results when the firm performs better than expected, and a downward movement indicates a "negative surprise" when the firm performs worse than expected.

In this task, you will see a sequence of 400 performance surprises from a typical company, and your job is to estimate whether the next performance surprise will be positive or negative. For each of the 400 periods, you will see the performance surprises of the last 14 periods on the screen.

In each period, you will be asked to give a price at which you would be willing to buy a share of stock in this company. If you buy the stock and see a positive surprise, the stock will pay you $100. If you buy the stock and see a negative surprise, the stock will pay you $0. The important thing to understand is the following: the price you are willing to pay will, in general, not be the price you actually pay for the stock. Instead, the actual price of the stock will be drawn randomly between $0 and $100. If your willingness to pay is above this random price, you will pay the random price and receive a share of the company. If your willingness to pay is below the random price, you do not buy the share of the company. **In order to make the most money under this rule, the best thing for you to do is set the price equal to probability you think there will be a positive surprise.**

Examples

1. Suppose you believe that there will be a positive earnings surprise with 75% chance. You should then be willing to pay exactly $75 for this share; if the actual price is $50, then you will pay $50 for something that has a 75% chance of winning $100 which on average, will make you money. If instead the random price drawn was $90, the rule says that you will not buy this stock since $75<$90. This is good because you avoid paying $90 for something that has only a 75% chance of paying you $100.

2. Suppose you are certain (a 100% chance) that there will be positive performance surprise. Then you would be willing to pay any price between $0-$100 to buy this stock. The only way to guarantee that you buy this stock is to set your price exactly equal to $100. If you made a mistake and set the price of the stock to $90, then if the random price drawn is $92, you would not be able to buy the $92 stock, which has a 100% chance to pay $100.

3. Suppose you are certain that there will be negative performance surprise (0% chance of a positive surprise). Then you are not willing to pay any price to buy this stock. The only way to guarantee that you don't end up paying something for this stock, is to set your price exactly equal to $0. If instead, you made a mistake and entered $10, then if the actual price drawn was $8, you would end up paying $8 for a stock that has 0% chance of paying you.

In each period, you will be given $100 in experimental currency to buy a share of the stock. Since the maximum price you would ever pay for a share is $100, you will always have enough cash to buy a share of this stock, since you receive a new $100 endowment each period. Your payoff in each period will depend on the three things: your willingness to pay, the actual price, and whether there was a positive or negative surprise. To illustrate your payoffs consider the two scenarios.

If you believe there will be a positive surprise for sure, and your willingness to pay is $100, and the actual price drawn is $0, and there is actually a positive surprise, then you will end the period with $100 – $0 + $100= $200. That is, you will end the period with the $100 you started with, you don't pay any cost since the price was $0, and you earn $100 for buying the stock and having a positive earning surprise.

If you believe there will be a positive surprise with 60% chance, the actual price drawn is $30, and there is a positive surprise, then your total earnings this period will be $100-$30+$100=$170.

Your final earnings will be the sum of each of your individual period earnings, divided by 5,000. It is important to emphasize once more: **the only way to maximize your final earnings is to enter your willingness to pay equal to the probability you think there will be a positive surprise.**

After every 50 trials, you will see your accumulated payoff, and will be allowed to take a short break. Before the real task starts, you will start with 5 practice trials.
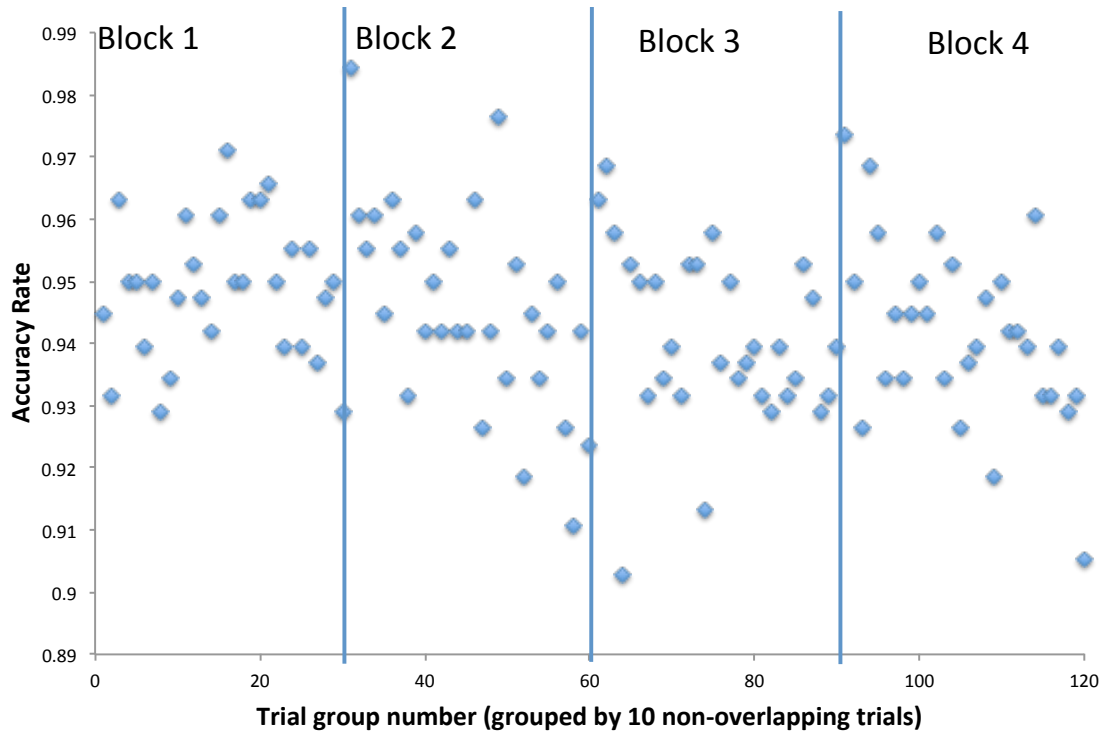
Figure S1. Average reaction times across subjects and across four blocks of trials.
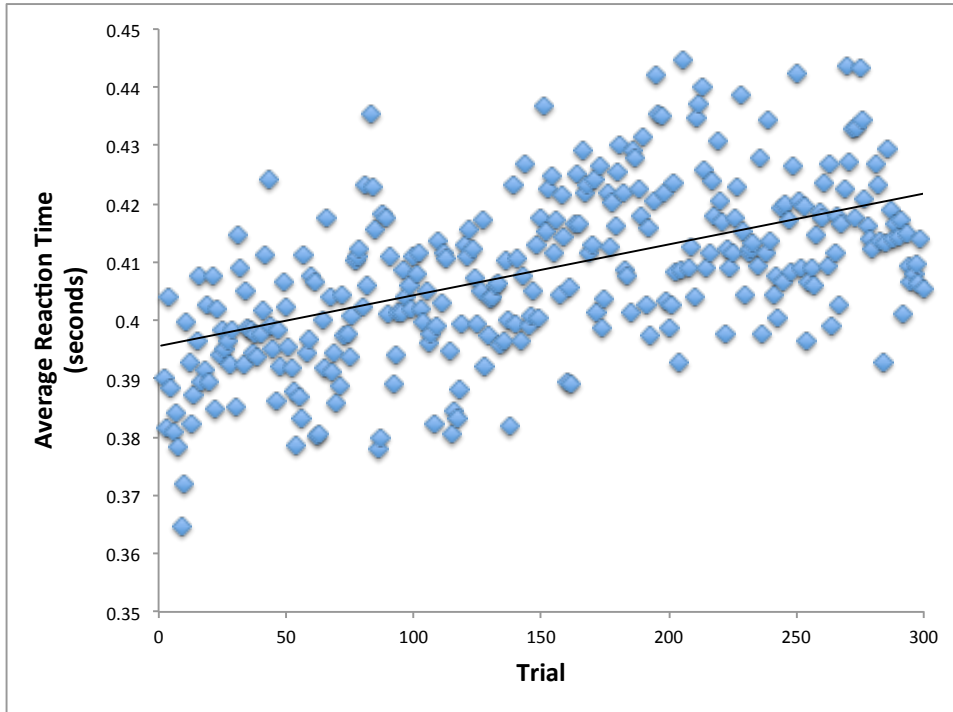

Figure S2. Average reaction times across subjects for each of the four blocks of trials (grouped by trials of 10).
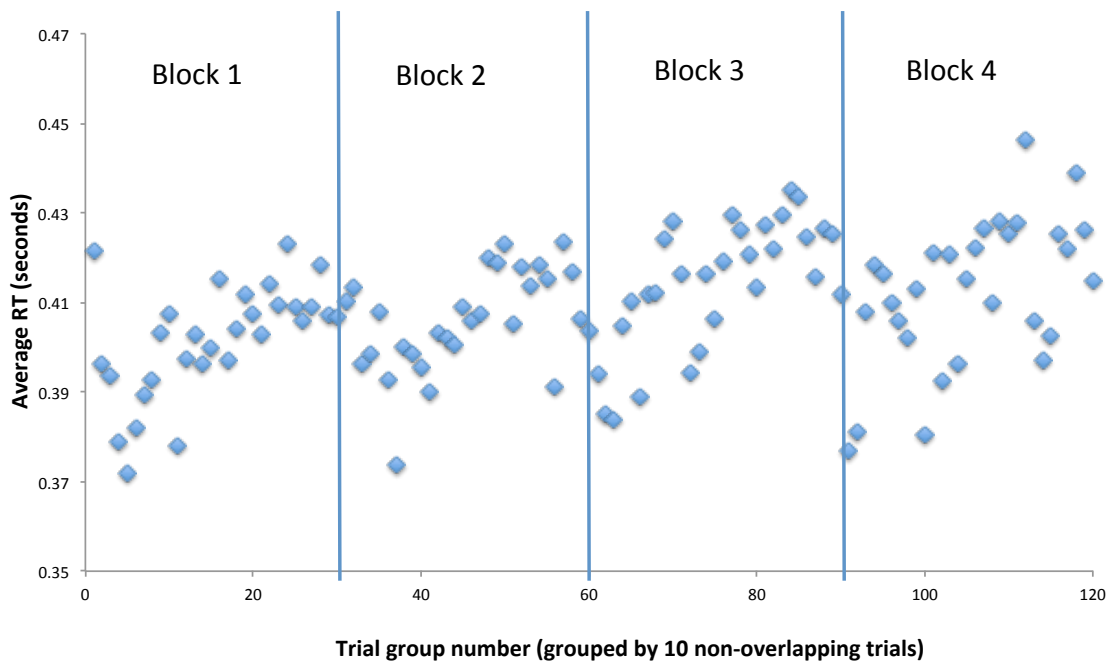
**Figure S3.** Average rate of correct responses across subjects for each of the four blocks of trials.

Figure S4. Average response times of correct and incorrect responses following "valid" cues (i.e., repetition following two or more repetitions or alternation following two or more alternations), 'invalid' cues (i.e., alternation following two or more repetitions or repetition following two or more alternations) and 'neutral' cues (all other trials)
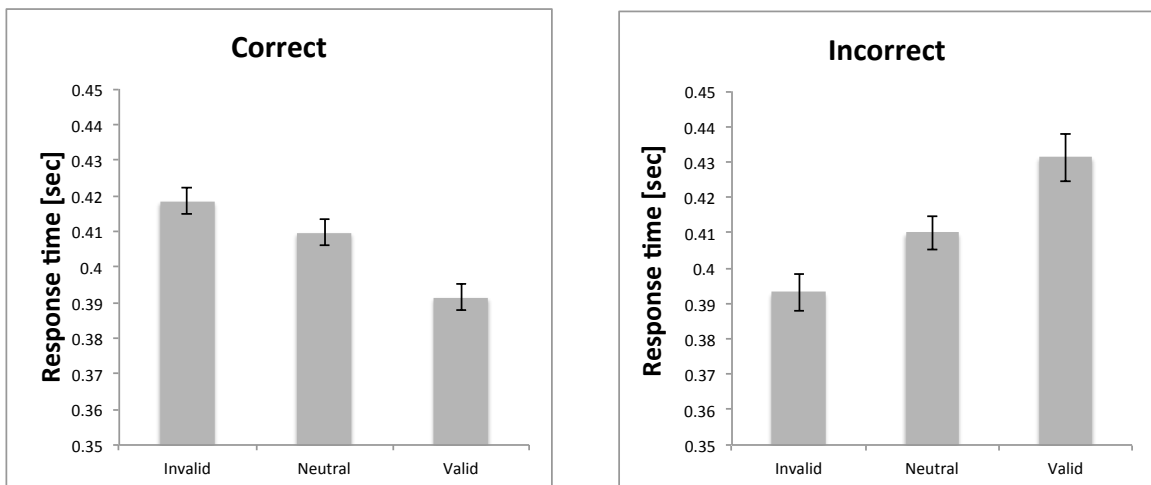
Figure S5. The sum of square errors of the DBM prediction, for the average reported beliefs in the EDT, as a function of parameter α. The mean parameter found in the PDT, α*, is marked by the dashed line.
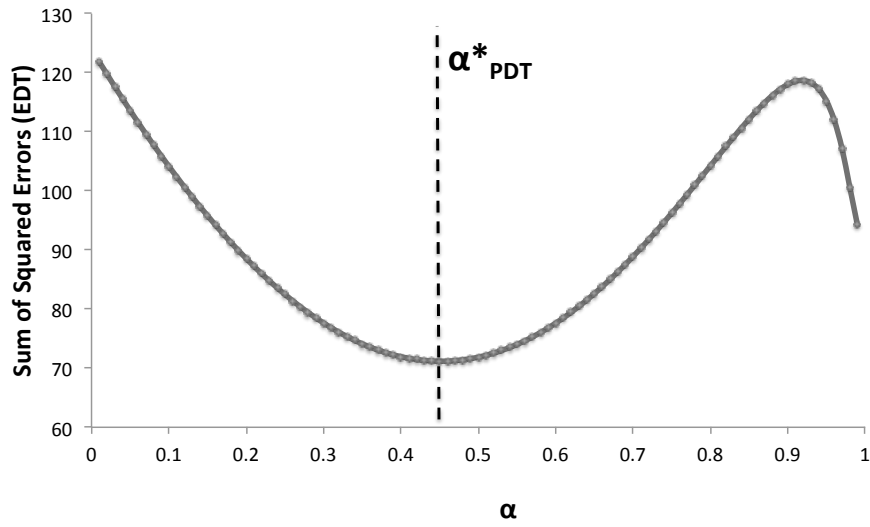
REFERENCES

1.   F. Tuerlinckx, J. Vandekerckhove, Diffusion model analysis with MATLAB: a DMAT primer. *Behavior Research Methods* **40**, 61 (2008).
2.   R. Ratcliff, F. Tuerlinckx, Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychonomic Bulletin & Review* **9,** 3 (2002).
3.   R. Ratcliff, R. Smith, A comparison of sequential sampling models for two-choice reaction time. *Psychological Review* **111** (2004).