Modeling sampling duration in decisions from experience

Nisheeth Srivastava, Johannes Müller-Trede
UC San Diego, 9500 Gilman Drive
La Jolla, CA 92093 USA

Paul Schrater
University of Minnesota
Minneapolis, MN 55455 USA

Edward Vul
UC San Diego, 9500 Gilman Drive
La Jolla, CA 92093 USA

Abstract

Cognitive models of choice almost universally implicate sequential evidence accumulation as a fundamental element of the mechanism by which preferences are formed. When to stop evidence accumulation is an important question that such models do not currently try to answer. We present the first cognitive model that accurately predicts stopping decisions in individual economic decisions-from-experience trials, using an online learning model. Analysis of stopping decisions across three different datasets reveals three useful predictors of sampling duration - relative evidence strength, how long it takes participants to see all rewards, and a novel indicator of convergence of an underlying learning process, which we call predictive *volatility*. We quantify the relative strengths of these factors in predicting observers' stopping points, finding that predictive volatility consistently dominates relative evidence strength in stopping decisions.

Keywords: response time; decision-making; evidence accumulation; sequential sampling; decisions from experience

Introduction

In moving decision theory from analyzing static economic descriptions towards the more dynamic decisions humans face in everyday life, the decisions-from-experience (DFE) paradigm presents an important step forward (Hertwig, Barron, Weber, & Erev, 2004). The DFE paradigm is a modification of certainty equivalence experiments in experimental economics. Certainty equivalence (CE) is a commonly used procedure for eliciting subjects' utility functions for money amounts (Hershey & Schoemaker, 1985). In a typical certainty equivalence task, participants are asked to choose between a risky option that pays H with a probability p and L with probability $1-p$ and a safe option that always pays M, where $H > M > L$. In standard CE experiments, participants are explicitly shown and are asked to choose between the two options given for each such tuple.

Decisions from experience modify this protocol: Participants are not shown the lottery payoffs and odds, and have to learn them via experience. Several interesting observations emerge from research on DFE. Subjects sample more variable options and options with higher stakes for longer, for example (Lejarraga, Hertwig, & Gonzalez, 2012). They also appear to underestimate the probability of rare events (Hertwig et al., 2004), though much less so with increasing experience (Zhang & Maloney, 2012).

In the particular variety of DFE we consider throughout this paper, typically known as the sampling paradigm, participants are permitted to sample each of two options without consequence as long as they like, before finally committing to a binding choice. This protocol is particularly interesting since it closely mirrors the flow of information in everyday human decisions - learn from the environment *ad libitum*, then make a choice. Importantly, such choices are actually composed of two decisions: a latent decision to terminate learning, and an overt decision to choose the risky or the safe option, based on the information acquired.

Efforts to model the overt decision about which option to choose have been relatively successful (Erev et al., 2010). Little attention has been paid, however, to modeling the earlier latent decision about how long to sample information (or when to stop learning). Research on DFE has used simple statistical approaches as place-holders, assuming an underlying probability distribution over sampling lengths and fitting this distribution to the empirical distribution of sampling lengths observed in the data (Gonzalez & Dutt, 2011). Markant et al. have recently proposed a model that jointly predicts choices and sampling length distribution (Markant, Pleskac, Diederich, Pachur, & Hertwig, 2015) However, since this model uses *known* lottery stakes, it cannot be applied to sampling-based DFE where the rewards and probabilities of the lottery must be learned from experience.

In this paper, we investigated variables that, on theoretical grounds, are expected to predict sampling lengths in DFE, without assuming *a priori* knowledge of lottery stakes and probabilities. Our analysis revealed that, across three different datasets, the predictive value of economic variables in DFE is exceeded by that of a novel psychological predictor, predictive *volatility*, which tracks abrupt changes in the magnitude of prediction error an observer experiences while learning about a DFE decision. We further developed a computational model of sequential sampling for DFE, incorporating these predictors, which makes accurate sampling length predictions for individual DFE trials. The ability of our model to predict sampling duration for individual DFE trials is categorically unprecedented, and is the principal contribution of this paper.

DFE: Psychology and economics

An intuitive economically motivated predictor for sampling duration in DFE is the difference between the imputed value of both options. Presumably, if two competing options are close in experienced value, observers could be expected to sample them for longer to differentiate them more precisely. Thus, if we measure the relative economic evidence contained in observers' sampling sequences, we'd expect to see smaller amounts of relative evidence associated with longer sequences. Such a conclusion would also mesh well with pre-existing theories of response time, that ground the evidence accumulation process formally in the sequential probability ratio test (Usher & McClelland, 2001; Busemeyer & Townsend, 1993).

Formally, we track the expected value difference (EVD) between the two gambles, measured at every sample for each DFE trial,

$$EVD = pH + (1-p)L - M, \quad (1)$$

where $p = \frac{|H|}{|H|+|L|}$, and $|x|$ is the number of times the outcome x has occurred in the sequence up to the time at which the measurement is taken. This quantity is called the EVD predictor in our following analyses.

We argue that such economic predictors, while useful, are not sufficient to model stopping criteria in DFE. Since information search is a universal activity in animals, functional parsimony suggests that observers use more general psychological variables to determine when to terminate information search, with the economic variables specific to DFE complementing them.

But which psychological variables? We suggest that the information-theoretic goal of DFE observers is to efficiently learn the reward rate of the lottery(ies) they are sampling. We treat the decision to terminate information gathering as a rational response to the agent realizing that the learning procedure has saturated. In practice, observers, while trying to learn useful models of their environment, have access to the prediction errors in such models. The central theoretical novelty of our current proposal is an intuition that unexpected rises in prediction error magnitude (as illustrated in Fig.1) signal the presence of unlearned environmental dynamics, stimulating rational observers to sample more.

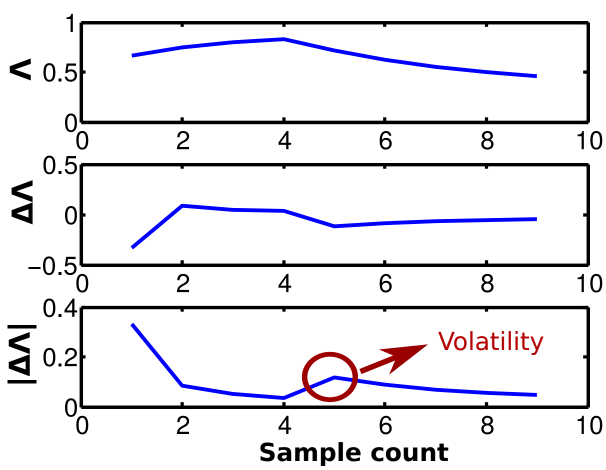

Figure 1: An intuitive view of predictive volatility. When the prediction error in a sequential learning process increases abruptly, it is reasonable to infer that the process has yet to converge. This indicator of the need to keep learning is what we call predictive volatility.

Assuming a simple parametric observer model, we can acquire insight into the learning process by tracking the evolution of the learned parameter over individual sampling sequences. If the learning procedure is efficient, the prediction error is expected to show an asymptotic gradual decrease, reflecting increasingly precise estimation through acquiring more data samples. Critically, we assume that human observers are intuitive statisticians in this particular sense - they are implicitly aware that the trajectory of prediction error corresponding to efficient learning takes this particular shape. However, in individual learning sequences, such declines are not always monotonic. We call breaks in the expected prediction error trend episodes of *predictive volatility*, and suggest that, since observers, through experience, become implicitly sensitive to the expected trajectory of prediction error, it is rational for them to treat observed volatility episodes as evidence that learning has not yet completed, and respond by sampling for longer.

Formally, we model this process, in the context of DFE, as a statistically efficient observer sequentially updating estimates of the mean parameter Λ of a Poisson distribution tracking the frequency with which the high outcome of the risky option occurs in the sampling sequence. Given every new sample, the parameter estimate shifts by some quantity $\Delta\Lambda$. The normative (and psychological) expectation is for the magnitude of this quantity to decay over time $|\Delta\Lambda_t| < |\Delta\Lambda_{t-1}|$. Deviations from this trend, therefore, suggest to an observer that some aspects of the task still remain potentially unlearned, and therefore justify continued sampling. These episodes, we assert, constitute predictive *volatility*, which we operationalize formally as

$$v(t) = 1 \text{ iff } |\Delta\Lambda_t| > \kappa \times |\Delta\Lambda_{t-1}|, \quad (2)$$

and 0 otherwise, with the proportionality constant taking values $\kappa > 1$. In all our experiments, we used a value of $\kappa = 2$; substantially larger magnitudes than this would degrade the information present in this signal, since such large fluctuations are statistically infeasible in DFE reward rate estimation, where the set of possible outcomes is very limited; substantially smaller values would add noise to the signal, in the form of volatility false positives. Across an entire sampling sequence, the cumulative effect of such episodes is measured by the trial volatility load

$$V = \sum_t v(t), \quad (3)$$

and is referred to, in our following analysis, as the *volatility* predictor

Finally, based on the task description of DFE, it is normatively expected that observers will try to see each of the three reward outcomes at least once before terminating sampling. Depending on how skewed the risky option's odds are, this can take a relatively long time. Thus, the number of samples it takes to see all three reward outcomes at least once contains valuable information about how long the sampling sequence can be, and enters our predictive model in the form of a *counting* predictor.

Results

We present two sets of results. We first demonstrate, using a proportional hazards regression analysis, that both the magnitude of difference in expected value and the amount of volatility seen in the sampling sequence influence sampling durations, in theoretically expected directions. Model selection reveals that volatility plays a more influential role in this process.

These results, however, are calculated using *post hoc* predictors that an actual observer would not have access to in real-time. Our second set of results uses sequential counterparts to these predictors to develop a sequential model that simulates the trajectory of the stopping probability of any DFE trial, sensitive to the influence of both differences in expected value and episodes of volatility experienced in real-time.

Data

We procured data from two sources: the decisions-by-sampling condition from the Technion Prediction Tournament, which involved two sets (one for estimation and one for competition) of 40 participants solving 30 such problems each, and a sample of 37 participants solving 19 different DFE problems we collected in the decisions-by-sampling condition of a different experiment (Experiment 2 in (Lejarraga & Muller-Trede, 2016)). Throughout this paper, we will refer to the Technion estimation dataset as **TE**, the Technion competition dataset as **TC**, and our own sample as **LM**.

Experimental protocols were largely[1] identical across the datasets. Participants could sample both options in each lottery pair as often as they liked, and subsequently committed to one final draw that would correspond to their actual payout. All participants were compensated via a random incentive scheme, and earned real money corresponding to their payout in one randomly selected choice problem. We note, however, that participants in LM revisited each choice problem in a group setting between individual trials (for details, see (Lejarraga & Muller-Trede, 2016)), whereas participants in TE and TC did not.

Volatility matters more

What can the expected value difference between options tell us about the decision to stop sampling? Theory suggests that large magnitudes, irrespective of sign, should be negatively correlated with sampling lengths. That is, if the evidence accumulated favors one of the gambles predominantly, a rational observer would terminate sampling and select it. Thus, if observers are using the weight of economic evidence to decide when to terminate information search, greater average magnitudes of the expected value difference, i.e. $\frac{1}{T}\sum_t^T |EVD_t|$ for sequences of duration T, should be correspond with earlier sampling termination and vice versa.

We tested this hypothesis by running a Cox proportional hazards regression, assessing the direction and magnitude of effect the average expected value difference measured during a sampling sequence has on the hazard rate across all trials per subject. As the top panels in Fig.2(A) show, EVD consistently increases hazard rates across participants in all three datasets, as expected.

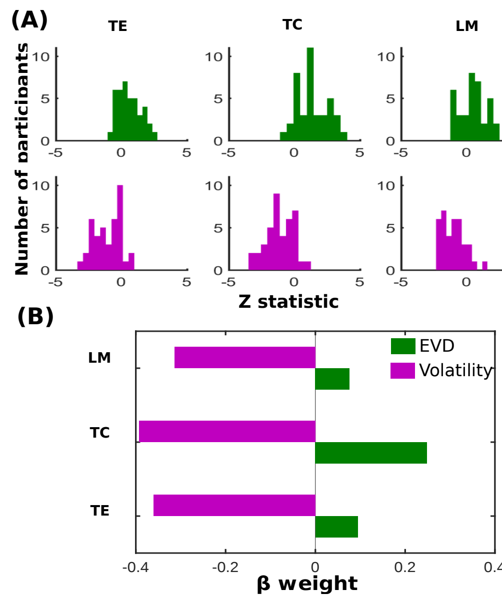

Figure 2: **(A)** Histograms of subject-wise Z-statistics obtained by Cox regression of predictors against sampling sequence lengths. Negative Z-statistics indicate that the predictor reduces the hazard rate, yielding longer sampling durations than baseline expectations. **(B)** Coefficients from Cox multiple regressions using normalized EVD and volatility predictors for all three datasets.

We ran a similar proportional hazard regression using volatility load as a predictor of sampling sequence lengths. As the bottom row in Fig.2(A) demonstrates, volatility load consistently retards hazard rates in participants across all three datasets tested. Panel B in Fig.2 compares the regression coefficients obtained when we use both predictors - EVD and volatility - normalized and combining all participants in each dataset. In all three cases, the normalized predictors are uncorrelated ($r < 0.05, p > 0.25$), which makes the regression coefficients (β-weights) informative about the relative importance of the two predictors (Nathans, Oswald, & Nimon, 2012). Volatility consistently dominates EVD as a predictor across all three datasets.

A similar conclusion can be drawn by computing the Bayesian Information Criterion (BIC) for regressions using the two predictors indvidually, and then together. Not only is the joint predictor preferable by BIC, corresponding ΔBIC values show that the volatility-alone model is considerably closer to the joint predictor than the EVD-alone model, suggesting that it is more powerful predictor. Together, these analyses demonstrate (i) that EVD and volatility are independent sources of information for predicting sampling duration

[1]Participants in the LM experiment reported choices by alloting fractions of an allocation budget to either option; TE and TC choice selections were binary. We do not believe this difference is salient for our purpose.

Table 1: Model selection using BIC. Lower values are better within individual datasets. ΔBIC from best model reported in brackets.

	TE	TC	LM
EVD	14312 (+165)	14643 (+177)	7090 (+63)
Volatility	14157(+10)	14529 (+63)	7031(+4)
Both	14147	14466	7027

and (ii) that volatility is a more informative predictor for sampling durations than EVD.

DFE sampling durations are *post hoc* predictable

But how much predictive ability do these predictors actually give us? As a simple indicator, a simple additive model combining these two predictors,

$$\text{duration} = \text{volatility} + \beta\ |\text{EVD}|$$

yields correlations $r = \{0.56, 0.53, 0.45\}$ with human sampling lengths for the TE, TC and LM datasets[2] respectively, suggesting that these predictors can explain between 20-30% of variance in sampling durations for human observers in DFE. Since no previous models have reported results for DFE sampling durations at the individual trial level, it is difficult to assess this performance in a comparative sense.

But we can do better than this - by incorporating the *counting predictor*. Recall that this predictor simply counts the number of samples it took an observer to see each of the three possible outcomes $\{H, M, L\}$ at least once during a trial. Incorporating it yields an augmented linear model

$$\text{duration} = \text{count} + \alpha\ \text{volatility} + \beta\ |\text{EVD}|$$

which shows correlations of $r = \{0.56, 0.69, 0.71\}$ with human data for the TE, TC and LM datasets respectively[3].

One might question why adding this new predictor did not improve the data-model correlation for the TE dataset in particular. This is because participants in this dataset strongly violated the normative expectation that participants would see all three outcomes at least once. Of the 1200 total trials in this dataset (40 participants × 30 problems), as many as 742 trials (62%) were terminated without having seen all three outcomes, including 192 (16%) that were terminated after drawing just two samples altogether. For comparison, participants in the TC and LM datasets terminated 41% and 14% of all trials before seeing all three possible outcomes, respectively. The higher sampling effort in the LM dataset may reflect additional intrinsic motivation their participants derived from the repeated social interactions between choice problems.

Given greater congruence with the assumption implicit in the counting predictor - that observers will not terminate sampling until they have seen all three outcomes at least once - its addition boosts the overall correlation of the model to ≈ 0.7 in the other two datasets, thereby showing that it is capable of explaining around 50% of the variance in the data. In light of the large variability in sampling lengths across both participants and problems, contingent on momentary fluctuations in valuation, attention, and motivation, it is likely that our model explains a much larger fraction of the total variance explicable.

Note though, that the large improvement in correlation by adding the counting predictor somewhat overstates its true explanatory value. By definition, the sample count at which all three options have been seen once cannot exceed the overall sampling sequence length. Thus, the counting predictor is upper-bounded by the independent variable it is trying to predict. The definitional absence of counting predictor values greater than the actual sampling length adds substantial linearity to the correlation, leading to the elevated number we see. Hence, while the counting predictor, *prima facie* adds substantial predictive value to our model, it does so for reasons that need not be theoretically insightful.

Real-time stopping point prediction

While we show that sampling lengths in DFE are substantially predictable *post hoc* using objectively observable predictors, not all these predictors are available to decision-makers at the time they're making their decisions. Neither the average EVD magnitude, nor the cumulative volatility load across the entire sequence is available to an observer while they are still in the process of sampling. In this section, we use elements of our predictors that are available to such observers in real-time, and test how well they predict observers' eventual stopping decisions.

To do so, we use computational models that perform the same sampling task as the observer, stepping through each trial sample by sample, predicting sequence lengths indirectly by estimating stopping probabilities λ_t at each sample. To make this analysis feasible, we make a simplifying assumption - that there are no individual differences across individuals within datasets. Doing so yields multiple data points for each DFE problem, instead of just one per problem-by-participant pair. From these, we construct a stopping point distribution for each problem.

We then use these stopping point distributions to fit a basic piece-wise constant hazard model that assumes the stopping probability increases linearly with the sampling count t, i.e.,

$$\lambda_{t+1} - \lambda_t = \delta, \tag{4}$$

and fit $\{\lambda_0, \delta\}$ for each unique problem such that the model's predicted sequence lengths,averaged across multiple runs (N=1000) become statistically indistinguishable (two-sided T-test $p > 0.9$) to the empirical stopping point distributions. We use these same $\{\lambda_0, \delta\}$ values in the subsequent models we test below.

This simple model is theoretically and empirically similar to previous sampling length models proposed in the literature, which assess model fit by testing whether it produces the

[2] Best fit $\beta = 0, -0.1, -0.4$

[3] For best fit values of $\alpha = \{3.8, 5.9, 2.6\}$ and $\beta = \{-0.2, -0.4, -0.9\}$ respectively.

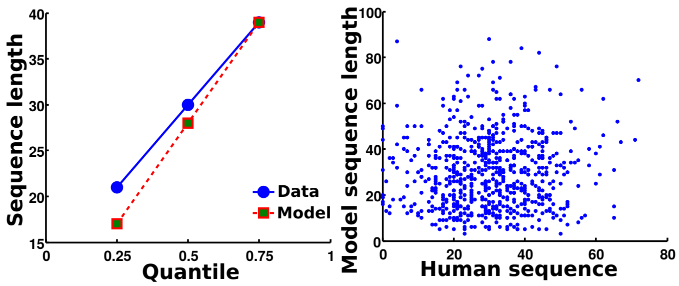

Figure 3: While simple statistical models of sampling lengths can reproduce population-level statistics *(left)*, they *(right)* fail to predict sampling sequence lengths for individual trials. Results shown for LM dataset.

same statistical distribution of sampling lengths as the underlying data.

Fig.3 illustrates that our baseline model, too, closely approximates the empirical distribution of sampling lengths (compare left panel with Figure 1 in (Markant et al., 2015)). However, it is a poor predictor of sampling lengths at the individual trial level ($r = 0.03$, right panel), illustrating a basic limitation of this modeling approach.

Table 2: Best fit correlations of sampling durations predicted by sequential models with human data in all three datasets.

Models	**TE**	**TC**	**LM**
Baseline	-0.04	0.02	0.03
Baseline + Vol	0.20	0.11	0.19
Baseline + EVD	0.15	0.11	0.06
Baseline + EVD + Vol	0.26	0.18	0.21
Baseline + EVD + Vol + Counting	**0.44**	**0.38**	**0.41**

Next, we added a modification: incrementing the baseline stopping probability by Δ every time the observer encounters volatility in the sampling sequence,

$$\lambda_{t+1} - \lambda_t = \delta + v(t)\Delta. \tag{5}$$

Here, volatility refers to single episodes of volatility within a sampling sequence, as defined in Eqn.2, not the cumulative quantity we measure across the entire trial. If, as we suspect, volatility retards the termination probability, we should expect negative values of Δ to correspond to greater correlations of the model's sample sequences with human data. As Fig.4 illustrates, across all three datasets, observer models that reduce stopping probability when encountering volatility do fit the data better. Hence, our hypothesis about the influence of volatility is substantiated.

We ran a similar analysis to measure the sample-by-sample impact of the EVD predictor. We assumed that incoming signals of greater relative evidence strength would affect the stopping probability following a logistic relationship, privileging the impact of larger values. Thus,

$$\log \frac{\lambda_t}{1-\lambda_t} = \log \frac{\lambda'_t}{1-\lambda'_t} + k \log |d_t + 1|, \tag{6}$$

where λ'_t is obtained from Eqn 4, k is fitted to the data to maximize the model-data correlation, and d_t is the EVD calculated using the sequence up to the t^{th} sample.

To combine the influence of volatility and EVD, we can reuse Eqn 6, with λ' calculated using Eqn 5 with $\{\Delta, k\}$ as free parameters. The best overall model fit yielded weakly positive correlations across the three datasets[4].

Finally, if we incorporate the counting predictor into our model in the form of a real-time decision threshold - if all options seen at sample t, terminate with probability λ_t, otherwise, terminate with probability λ_0 - the correlations improve substantially, to $\{0.44, 0.38, 0.41\}$, for the same parameter values as in the previous model. These final correlation values indicate the upper limit of our model's ability to predict observers' stopping points using only information they themselves would have available in real-time.

Discussion

This paper develops theory and algorithms to predict sampling durations in economic decisions-from-experience, wherein observers get to sample options before committing to a choice. We argue, and then empirically demonstrate, that a combination of economic evidence strength, psychological predictive volatility, and simply tracking how long it takes participants to observe the entire reward structure of the particular DFE problem they are solving goes a considerable way in explaining the individual sampling stopping points of participants. Whereas earlier approaches can reproduce the population's sampling duration distribution, they are effectively random at making trial-level predictions. Our account replicates the distribution-level performance, while making reasonably accurate trial level predictions. Finally, since it directly emits stopping point predictions, our model can easily be combined with choice models that respect the epistemic limitations of sampling-based DFE, e.g., primed sampling, natural means, etc. (Erev et al., 2010) to make joint predictions of choice and sampling duration.

While one of the predictors we use, time till all rewards are observed, is specific to DFE, the other two are substantially more general, and hence, could be used to predict sampling duration in other experimental modalities. For example, Juni et al. have demonstrated how observers sample for longer when encountering noisier stimuli in a visuomotor estimation task (Juni, Gureckis, & Maloney, 2011). In this modality, greater stimulus noise corresponds directly to evidence strength, with lower values, due to greater noise, associated with longer sampling, as in our case.

Our discovery of the significant influence of predictive *volatility* on sampling duration warrants further investigation.

[4]Best fit parameter values for all three datasets, $\Delta = \{-0.20, -0.15, -0.125\}, k = \{0.02, 0.02, 0.02\}$

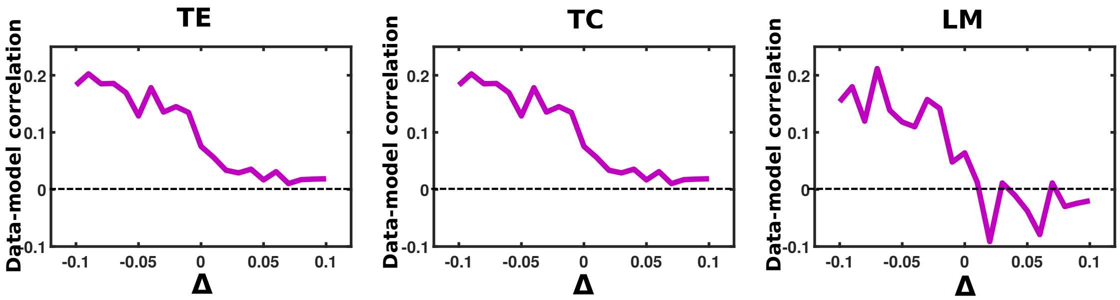

Figure 4: Model-data correlations for observer models fitted using different increments to stopping probability when encountering volatility.

Predictive volatility is a very accessible information signal that observers could be using in a variety of real-world decisions from both experience and memory. To date, researchers have modeled response durations as arising from either evidence accumulation rising to a fixed threshold (Usher & McClelland, 2001), or a time-sensitive threshold collapsing to meet accumulating evidence (Thura, Beauregard-Racine, Fradet, & Cisek, 2012). Our results show that thresholds need not stay fixed or fall over time; they could rise adaptively within trials also, sensitive to sequence-dependent predictive volatility. The representational generality of this measure, alongside our demonstration of its consistent and considerable impact on DFE stopping point decisions, invites further exploration in other experiment designs.

The role of predictive volatility in determining when to terminate sampling could also streamline the functional interpretation of cortico-striatal dopaminergic activity in decision-making (Schultz, Dayan, & Montague, 1997). Whereas dopamine has been experimentally associated with encoding both reward and prediction error, the latter association appears to be more robust, in the sense that it is congruent with a larger literature on the role of prediction error in multiple motor, cognitive and perceptual functions (Friston et al., 2012). Our account shows how dopaminergic activity could be temporally correlated with the choice process, as a critical participant in the decision to terminate information search, without actually encoding reward.

References

Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, *100*(3), 432.

Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S. M., Hau, R., … Lebiere, C. (2010). A choice prediction competition: Choices from experience and from description. *Journal of Behavioral Decision Making*, *23*(1), 15–47.

Friston, K., Shiner, T., FitzGerald, T., Galea, J., Adams, R., Sporns, O., et al. (2012). Dopamine, affordance and active inference. *PLoS Comput Biol*, *8*(1), e1002327.

Gonzalez, C., & Dutt, V. (2011). Instance-based learning: Integrating sampling and repeated decisions from experience. *Psychological Review*, *118*(4), 523.

Hershey, J. C., & Schoemaker, P. J. (1985). Probability vs certainty equivalence methods in utility measurement: Are they equivalent? *Management Science*, *31*(10), 1213–1231.

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, *15*(8), 534–539.

Juni, M. Z., Gureckis, T. M., & Maloney, L. T. (2011). Dont stop til you get enough: adaptive information sampling in a visuomotor estimation task. In *33rd Annual conference of the cognitive science society* (pp. 2854–2859).

Lejarraga, T., Hertwig, R., & Gonzalez, C. (2012). How choice ecology influences search in decisions from experience. *Cognition*, *124*(3), 334–342.

Lejarraga, T., & Muller-Trede, J. (2016). When experience meets description: How dyads integrate experiential and descriptive information in risky decisions. *Management Science (in press)*.

Markant, D., Pleskac, T. J., Diederich, A., Pachur, T., & Hertwig, R. (2015). Modeling choice and search in decisions from experience: A sequential sampling approach. In *37th annual conference of the cognitive science society* (pp. 1512–1517).

Nathans, L. L., Oswald, F. L., & Nimon, K. (2012). Interpreting multiple linear regression: A guidebook of variable importance. *Practical Assessment, Research & Evaluation*, *17*(9), 1–19.

Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*(5306), 1593–1599.

Thura, D., Beauregard-Racine, J., Fradet, C.-W., & Cisek, P. (2012). Decision making by urgency gating: theory and experimental support. *Journal of Neurophysiology*, *108*(11), 2912–2930.

Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychological Review*, *108*(3), 550.

Zhang, H., & Maloney, L. T. (2012). Ubiquitous log odds: a common representation of probability and frequency distortion in perception, action, and cognition. *Frontiers in Neuroscience*, *6*.