
The rat as particle filter

Nathaniel D. Daw *
New York University
daw@cns.nyu.edu

Aaron C. Courville
Université de Montréal
aaron.courville@gmail.com

Abstract

The core tenet of Bayesian modeling is that subjects represent beliefs as distributions over possible hypotheses. Such models have fruitfully been applied to the study of learning in the context of animal conditioning experiments (and analogously designed human learning tasks), where they explain phenomena such as retrospective revaluation that seem to demonstrate that subjects entertain multiple hypotheses simultaneously. However, a recent quantitative analysis of individual subject records by Gallistel and colleagues cast doubt on a very broad family of conditioning models by showing that all of the key features the models capture about even simple learning curves are artifacts of averaging over subjects. Rather than smooth learning curves (which Bayesian models interpret as revealing the gradual tradeoff from prior to posterior as data accumulate), subjects acquire suddenly, and their predictions continue to fluctuate abruptly. These data demand revisiting the model of the individual versus the ensemble, and also raise the worry that more sophisticated behaviors thought to support Bayesian models might also emerge artifactually from averaging over the simpler behavior of individuals.

We suggest that the suddenness of changes in subjects' beliefs (as expressed in conditioned behavior) can be modeled by assuming they are conducting inference using sequential Monte Carlo sampling with a small number of samples — one, in our simulations. Ensemble behavior resembles exact Bayesian models since, as in particle filters, it averages over many samples. Further, the model is capable of exhibiting sophisticated behaviors like retrospective revaluation at the ensemble level, even given minimally sophisticated individuals that do not track uncertainty from trial to trial. These results point to the need for more sophisticated experimental analysis to test Bayesian models, and refocus theorizing on the individual, while at the same time clarifying why the ensemble may be of interest.

1 Introduction

A central tenet of the Bayesian program is the representation of beliefs by distributions, which assign probability to each of a set or continuum of different hypotheses. The prominent theoretical status accorded to such ambiguity seems rather puzzlingly at odds with the all-or-nothing nature of our everyday perceptual lives. For instance, subjects observing ambiguous or rivalrous visual displays famously report experiencing either percept alternately and exclusively; for even the most fervent Bayesian, it seems impossible simultaneously to interpret the Necker cube as potentially facing either direction.

One laboratory model for the formation of beliefs and their update in light of experience is Pavlovian conditioning in animals (and analogously structured prediction tasks in humans). There is a rich program of reinterpreting data from such experiments (which go back a century) in terms of statistical inference [1, 2, 3, 4, 5, 6]. The data do appear in a number of respects to reflect key features of the Bayesian ideal — specifically, that subjects represent beliefs as distributions with uncertainty and

*DRAFT version: NIPS 2007 preproceedings

appropriately employ it in updating them in light of new evidence. Most notable in this respect are *retrospective revaluation* phenomena (e.g., [7, 8]), which demonstrate that subjects are able to revise previously favored beliefs in a way suggesting they had entertained alternative hypotheses all along [6].

However, the data addressed by such models are, in almost all cases, averages over large numbers of subjects. This raises the question whether individuals really exhibit the sophistication attributed to them by the models, or if it instead somehow emerges from the ensemble. Recent work by Gallistel and colleagues [9] frames the problem particularly sharply. Whereas subject-averaged responses exhibit smooth learning curves approaching asymptote (interpreted by Bayesian modelers as reflecting the gradual tradeoff from prior to posterior as data accumulate), individual records exhibit neither smooth acquisition nor asymptote. Rather, most unlike the models, responding emerges abruptly and fluctuates perpetually.

Here we suggest that individuals’ behavior in conditioning might be understood in terms of Monte Carlo methods for sequentially *sampling* different hypotheses (e.g., [10]). Such a model preserves the insights of a statistical framing while accounting for the characteristics of individual records. Also, through the metaphor of particle filtering, it also explains why exact Bayesian reasoning is a good account of the ensemble. Finally, it addresses another common criticism of Bayesian models: that they attribute wildly intractable computations to the individual.

To make our point in the most extreme way, and to explore the most different corner of the model space, we here develop the idea that (as with percepts in the Necker cube) subjects sample only a single hypothesis at a time. That is, we treat them as particle filters employing only one particle. We show that even given individuals of such minimal capacity, effects like retrospective revaluation can emerge in the ensemble. Clearly intermediate models are possible, either employing more samples or mixtures of sampling and exact methods within the individual, and the insights developed here will extend to those cases. We therefore do not mean to defend the extreme claim that subjects never track or employ uncertainty — we think this would be highly maladaptive, though sampling likely has a role — but instead intend to explore the role of sampling and also point out how poor is the evidentiary record supporting more sophisticated theories, and how great is the need for more sophisticated experimental and analytical methods to test them.

2 Model

2.1 Conditioning as exact filtering

In conditioning experiments, a subject (say, a rat) experiences outcomes (“reinforcers,” say, food) paired with stimuli (say, a bell). That subjects learn thereby to predict the outcomes on the basis of the stimuli is demonstrated by the finding that they emit anticipatory behaviors (such as salivation to the bell) which are taken directly to reflect the expectation of the outcome. Human experiments are analogously structured, but using various cover stories (such as disease diagnosis or weather prediction) and with subjects typically simply asked to state their beliefs about how much they expect the outcome.

A standard statistical framing for such a problem [5], which we will adopt here, is to assume that subjects are trying to learn the conditional probability $P(r \mid \mathbf{x})$ of (real-valued) outcomes r given (vector-valued) stimuli \mathbf{x} . One simple generative model is to assume that each stimulus x_i (bells, lights, tones) produces reinforcement according to some unknown parameter w_i ; that the contributions of multiple stimuli sum; and that the actual reward is Gaussian in the the aggregate. That is, $P(r \mid \mathbf{x}) = \mathcal{N}(\mathbf{x} \cdot \mathbf{w}, \sigma_o^2)$, where we take the variance parameter as known. The goal of the subject is then to infer the unknown weights in order to predict reinforcement. If we further assume the weights \mathbf{w} can change with time, and take that change as Gaussian diffusion,

$$P(\mathbf{w}_{t+1} \mid \mathbf{w}_t) = \mathcal{N}(\mathbf{w}_t, \sigma_d^2 \mathbf{I}) \tag{1}$$

then we complete the well known generative model for which Bayesian inference about the weights can be accomplished using the Kalman filter algorithm [11]. Given a Gaussian prior on \mathbf{w}_0 , the posterior distribution $P(\mathbf{w}_t \mid \mathbf{x}_{1..t-1}, r_{1..t-1})$ also takes a Gaussian form, $\mathcal{N}(\hat{\mathbf{w}}_t, \Sigma_t)$, with the mean and covariance given by the recursive Kalman filter update equations.

Returning to conditioning, a subject’s anticipatory responding to test stimulus \mathbf{x}_t is taken to be proportional to her expectation about r_t conditional on \mathbf{x}_t , marginalizing out uncertainty over the weights. $E(r_t | \mathbf{x}_t, \hat{\mathbf{w}}_t, \Sigma_t) = \mathbf{x}_t \cdot \hat{\mathbf{w}}_t$.

2.2 Conditioning as particle filtering

Here we assume instead that subjects do not maintain uncertainty in their posterior beliefs, via covariance Σ_t , but instead that subject L treats her point estimate $\hat{\mathbf{w}}_t^L$ as true with certainty. Even given such certainty, because of diffusion intervening between t and $t + 1$, $\hat{\mathbf{w}}_{t+1}^L$ will be uncertain; let us assume that she recursively samples her new point estimate $\hat{\mathbf{w}}_{t+1}^L$ from the posterior given this diffusion and the new observation $\mathbf{x}_{t+1}, r_{t+1}$:

$$\hat{\mathbf{w}}_{t+1}^L \sim P(\mathbf{w}_{t+1}^L | \mathbf{w}_t = \hat{\mathbf{w}}_t^L, \mathbf{x}_{t+1}, r_{t+1}) \quad (2)$$

This is simply a Gaussian given by the standard Kalman filter equations. In particular, the mean of the sampling distribution is $\hat{\mathbf{w}}_t^L + \mathbf{x}_t \kappa (r_t - \mathbf{x}_t \cdot \hat{\mathbf{w}}_t^L)$. Here the Kalman gain $\kappa = \sigma_d^2 / (\sigma_d^2 + \sigma_o^2)$ is constant; the *expected* update in $\hat{\mathbf{w}}$, then, is just that given by the Rescorla-Wagner [12] model.

Such seemingly peculiar behavior may be motivated by the observation that, assuming that the initial $\hat{\mathbf{w}}_0^L$ is sampled according to the prior, this process also describes the evolution of a single sample in particle filtering by sequential importance sampling, with Equation 2 as the optimal proposal distribution [10]. (In this algorithm, particles evolve independently by sequential sampling, and do not interact except for resampling.)

Of course, the idea of such sampling algorithms is that one can estimate the true posterior $\hat{\mathbf{w}}_t$ by averaging over particles. In importance sampling, the average must be weighted according to *importance weights*. These (here, the product of $P(r_{t+1} | \mathbf{x}_{t+1}, \mathbf{w}_t = \hat{\mathbf{w}}_t^L)$ over each t) serve to squelch the contribution of particles whose trajectories turn out to be conditionally more unlikely given subsequent observations. If subjects were to behave in accord with this model, then this would give us some insight into the ensemble average behavior, though if computed without importance reweighting, the ensemble average will appear to learn more slowly than the true posterior.

2.3 Resampling and jumps

One reason why subjects might employ sampling is that, in generative models more interesting than the toy linear, Gaussian one used here, Bayesian reasoning is notoriously intractable. However, the approximation from a small number of samples (or in the extreme case considered here, one sample) would be noisy and poor. As we can see by comparing the particle filter update rule of Equation 2 to the Kalman filter, because the subject-as-single-sample does not carry uncertainty from trial to trial, she is systematically overconfident in her beliefs and therefore tends to be more reluctant than optimal in updating them in light of new evidence (that is, the Kalman gain is low). This is the individual counterpart to the slowness at the ensemble level, and at the ensemble level, it can be compensated for by importance reweighting and also by *resampling* (here we consider standard SIR importance resampling for sequential importance samplers; [13, 10]). Resampling kills off conditionally unlikely particles and keeps most samples in conditionally likely parts of the space, with similar and high importance weights. Since optimal reweighting and resampling both involve normalizing importance weights over the ensemble, they are not available to our subject-as-sample.

However, there are some generative models that are more forgiving of these problems. In particular, consider Yu and Dayan’s [14] diffusion-jump model, which replaces Equation 1 with

$$P(\mathbf{w}_{t+1} | \mathbf{w}_t) = (1 - \pi) \mathcal{N}(\mathbf{w}_t, \sigma_d^2 \mathbf{I}) + \pi \mathcal{N}(0, \sigma_j^2 \mathbf{I}) \quad (3)$$

with $\sigma_j \gg \sigma_d$. Here, the weights usually diffuse as before, but occasionally (with probability π) are regenerated anew. (We refer to these events as “jumps” and the previous model of Equation 1 as a “no-jump” model, even though, strictly speaking, diffusion is accomplished by smaller jumps.) Since optimal inference in this model is intractable (the number of modes in the posterior grows exponentially) Yu and Dayan [14] propose to make a sort of maximum likelihood determination whether a jump occurred or not; conditional on this the posterior is again Gaussian and inference proceeds as in the Kalman filter.

If we use Equation 3 together with the one-sample particle filtering scheme of Equation 2, then we simplify inference a bit further than this by not carrying over uncertainty from trial to trial even

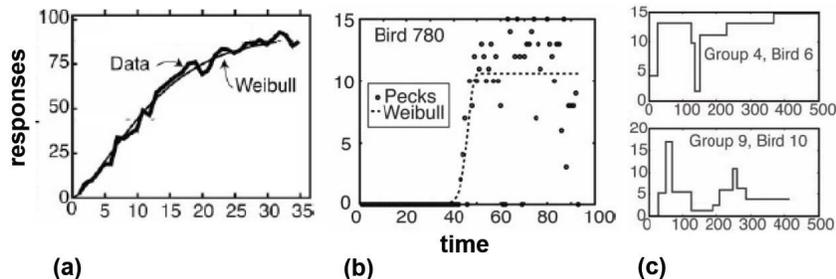


Figure 1: Aggregate versus individual behavior in conditioning, figures adapted from Gallistel et al (2004) (a) Mean over subjects reveals smooth, slow acquisition curve (timebase is in sessions). (b) Individual records are noisier and with more abrupt changes (timebase is in trials). (c) Examples of fits to individual records assuming the behavior is piecewise Poisson with abrupt rate shifts.

in a simplified posterior. Instead, as before, at each step, we sample from the posterior $P(\mathbf{w}_{t+1}^L \mid \mathbf{w}_t = \hat{\mathbf{w}}_t^L, \mathbf{x}_{t+1}, r_{t+1})$ given total confidence in our previous estimate. This distribution now has two modes, one representing the posterior given that a jump occurred, the other representing the posterior given no jump.

Importantly, we are more likely to infer a jump, and resample from scratch, if the observation r_{t+1} is far from that expected under the hypothesis of no jump, $\mathbf{x}_{t+1} \cdot \hat{\mathbf{w}}_t^L$. Specifically, the probability that no jump occurred (and that we therefore resample according to the posterior distribution given drift — effectively, the chance that the sample “survives” as it would have in the no-jump Kalman filter), by Bayes’ rule, is proportional to $P(r_{t+1} \mid \mathbf{x}_{t+1}, \mathbf{w}_t = \hat{\mathbf{w}}_t^L, \text{no jump})$, which is also the factor that the trial would contribute to the importance weight in the no-jump Kalman filter model of the previous section. The importance weight, in turn, is also the factor that would determine the chance that a particle would be selected during an exact resampling step [13, 10].

There is therefore a limited correspondence between sampling in this model and sampling with resampling in the simpler generative model of Equation 1. Of course, this cannot quantitatively accomplish optimal resampling, both because the chance that a particle survives should be normalized with respect to the population, and because the distribution from which a non-surviving particle resamples should also depend on the ensemble distribution. However, it has a similar qualitative effect of suppressing conditionally unlikely samples and replacing them ultimately with conditionally more likely ones.

We can therefore view the jumps of Equation 3 in two ways. First, they could correctly model a jumpy world; by periodically resetting itself, such a world would be relatively forgiving of the tendency for particles in sequential importance sampling to turn out conditionally unlikely. Alternatively, the jumps can be viewed as a fiction effectively encouraging a sort of resampling to improve the performance of low-sample particle filtering in the *non-jumpy* world of Equation 1. Whatever their interpretation, as we will show, they are critical to explaining subject behavior in conditioning.

3 Acquisition

In this and the following section, we illustrate the behavior of individuals and the ensemble in some simple conditioning tasks, comparing particle filter models with and without jumps (Equations 1 and 3).

Figure 1 reproduces some data from Gallistel and colleagues [9], who presented a number of analyses quantifying what had long been anecdotally known about conditioning: that individual records look nothing like the averages over subjects that have been the focus of much theorizing. Consider the simplest possible experiment, in which a stimulus A is paired repeatedly with food. (We write this as $A+$.) Averaged learning curves slowly and smoothly climb toward asymptote (Figure 1a, here the anticipatory behavior measured is pigeons pecking), just as does the average estimate \hat{w}_A in the Kalman filter models.

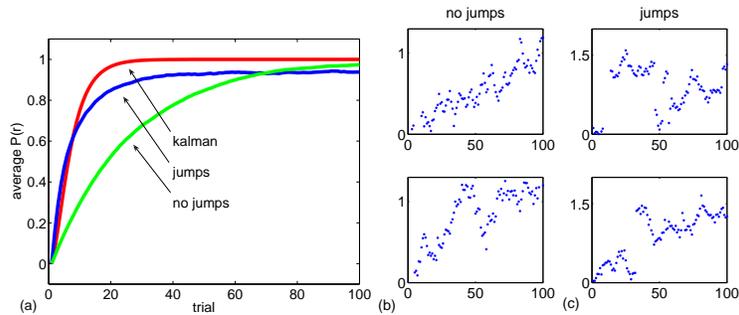


Figure 2: Simple acquisition in conditioning, simulations using particle filter model. (a) Mean behavior over samples for jump ($\pi = 0.075$) and no-jump ($\pi = 0$) particle filter models of conditioning, plotted against exact Kalman filter for same parameters ($\sigma_d = 0.1$; $\sigma_o = 0.5$). (b) Two examples of individual subject traces for no-jump particle filter model. (c) Two examples of individual subject traces for particle filter model incorporating jumps.

Viewed in individual records (Figure 1b), acquisition is much more abrupt (often it occurred in a single trial), and the subsequent behavior much more variable. The slow acquisition results from the average over abrupt transitions occurring at a range of latencies. Gallistel et al. [9] characterized the behavior as piecewise Poisson with instantaneous rate changes (Figure 1c). These results present a challenge to the bulk of models of conditioning — not just Bayesian ones, but also associative learning theories like the seminal model of Rescorla & Wagner [12] ubiquitously produce smooth, asymptoting learning curves of a sort that these data reveal to be essentially an artifact of averaging.

One further anomaly with Bayesian models even as accounts for the average curves is that acquisition is absurdly slow from a normative perspective — it emerges long after subjects using reasonable priors would be highly certain to expect reward. This was pointed out by Kakade and Dayan [5], who also suggested an account for why the slow acquisition might actually be normative due to unaccounted priors caused by pretraining procedures known as hopper training. However, Balsam and colleagues later found that manipulating the hopper pretraining did not speed learning [15].

Figure 2 illustrates individual and group behavior for the two particle filter models, with the group behavior plotted against the optimal Kalman filter for the same parameters. As expected, at the ensemble level (Figure 2a), particle filtering without jumps (and averaged without importance weighting or resampling) acquires slowly; the addition of jumps can speed this up. In individual traces (Figure 2c), the jumps both at and after acquisition capture the key qualitative features of the individual records. Notably, when a jump is sampled, the posterior distribution conditional on having jumped is centered near the observed r_t , meaning that the sampled weight will likely produce abrupt acquisition. (The sampling model without jumps exhibits smoother, albeit fluctuating, behavior; 2b.)

These simulations demonstrate, first, how sequential sampling using a very low number of samples is a good model of the puzzling features of individual behavior in acquisition, and at the same time clarify why subject-averaged records resemble the results of exact inference. Depending on the frequency of jumps (which help to compensate for this problem) the fact that these averages are of course computed without importance weighting may also help to explain the slowness of acquisition. This could be true regardless of whether other factors, such as those posited by Kakade and Dayan [5], could contribute.

4 Retrospective revaluation

So far, we have shown that sequential sampling provides a good qualitative characterization of individual behavior in the simplest conditioning experiments. But the best support for sophisticated Bayesian models of learning comes from more demanding tasks such as retrospective revaluation. These tasks give the best indication that subjects maintain something more than a point estimate of the weights, and instead strongly suggest that they maintain a full joint distribution over them. However, as we will show here, this effect can actually emerge due to covariance information be-

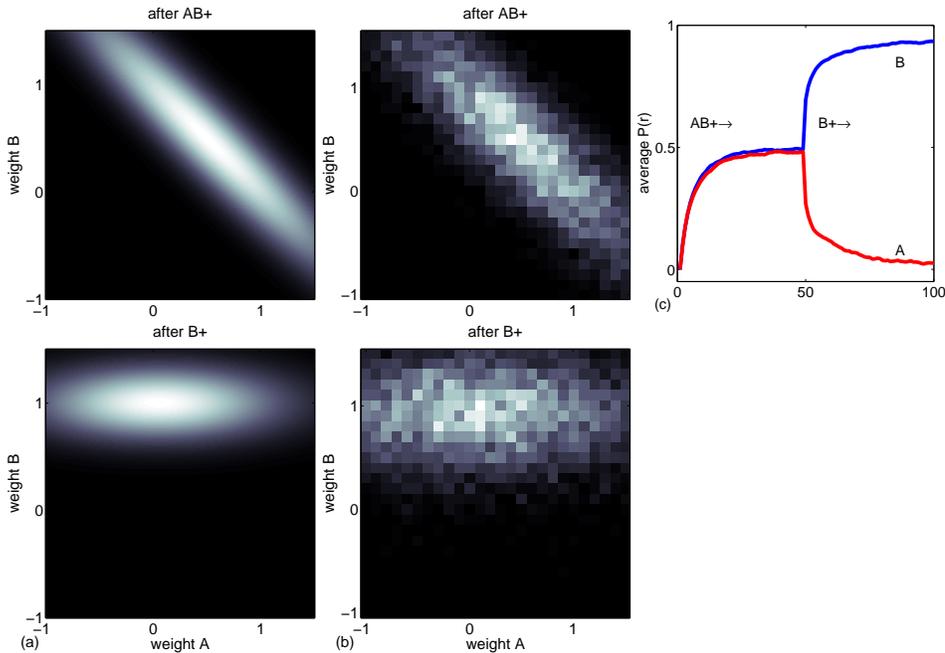


Figure 3: Simulations of backward blocking effect. (a) Exact Kalman filter simulations of backward blocking (after Kakade & Dayan, 2001); shown is joint distribution over \hat{w}_A and \hat{w}_B following first-phase $AB+$ training (top) and second phase $B+$ training (bottom). (b) Same as (a), but distribution is derived from histogram of individual particles’ joint point beliefs about the weights in the particle filter model with jumps. (c) Average over ensemble of \hat{w}_A and \hat{w}_B in particle filter model, showing development of backward blocking. Parameters as in Figure 2.

ing implicitly represented in the ensemble of beliefs over subjects, even if all the individuals are one-particle samplers.

Retrospective revaluation refers to how the interpretation of previous experience can be changed by subsequent experience. A typical task, called *backward blocking* [7, 8], has two phases. First, two stimuli, A and B , are paired with each other and reward ($AB+$), so that both develop a moderate level of responding. In the second phase, B alone is paired with reward ($B+$), and then the prediction to A alone is probed. The typical finding is that responding to A is attenuated; the intuition is that the $B+$ trials suggested that B alone was responsible for the reward received in the $AB+$ trials, so the association of A with reward is thereby retrospectively discounted. Backward blocking, like other retrospective revaluation phenomena, is hard to demonstrate in animals (though see [16]) but robust in humans [7, 8].

Kakade and Dayan [6] gave a more formal analysis of the task in terms of the Kalman filter model. In particular they point out that on the initial $AB+$ trials, the model will learn that \hat{w}_A and \hat{w}_B are *anticorrelated* — i.e., that they together add up to about one. This is represented in the covariance Σ ; the joint distribution is illustrated in Figure 3a. Subsequent $B+$ -training demonstrates that \hat{w}_B is high, which means, given its anticorrelation with \hat{w}_A , that the latter must be low. Note that this explanation turns crucially on the representation of the full joint distribution over the weights, rather than just a point estimate.

Figure 3b demonstrates the same thing in the particle filter model with jumps. At the end of $AB+$ training, the subjects as an ensemble represent the anti-correlated joint distribution over the weights, even though each individual maintains only a particular point belief. Moreover, $B+$ training causes a backward blocking effect. This is because individuals who believe that \hat{w}_A is high tend also to believe that \hat{w}_B is low, which makes them most likely to sample that a jump has occurred during subsequent $B+$ training. The samples most likely to stay in place have \hat{w}_A low and \hat{w}_B high; beliefs about \hat{w}_A are, on average, thereby reduced, producing the backward blocking effect in the ensemble.

Note that this effect depends on the subjects sampling using a generative model that admits of jumps (Equation 3). Although the population implicitly represents the covariance between \hat{w}_A and \hat{w}_B even using the diffusion model with no jumps (Equation 1; simulations not illustrated), subsequent $B+$ training has no tendency to suppress the relevant part of the posterior, and no backward blocking effect is seen. Again, this traces to the lack of a mechanism for downweighting samples that turn out to be conditionally unlikely.

5 Discussion

We have suggested that individual subjects in conditioning experiments behave as though they are sequentially sampling hypotheses about the underlying weights: like particle filters using a single sample. This model reproduces key and hitherto theoretically troubling features of individual records, and also, rather more surprisingly, has the ability to reproduce more sophisticated behaviors that had previously been thought to demonstrate that subjects represented distributions in a fully Bayesian fashion. One practical problem with particle filtering using a single sample is the lack of distributional information to allow resampling or reweighting; we have shown that use of a particular generative model previously proposed by Yu and Dayan [14] (involving sudden shocks that effectively accomplish resampling) helps to compensate qualitatively if not quantitatively for this failing. This mechanism is key to all of our results.

Gallistel and colleagues' [9] demonstration that individual learning curves exhibit none of the features of the ensemble average curves that had previously been modeled poses rather a serious challenge for theorists: After all, what does it mean to model only the ensemble? Surely the individual subject is the appropriate focus of theory — particularly given the evolutionary rationale often advanced for Bayesian modeling, that individuals who behave rationally will have higher fitness. The present work aims to refocus theorizing on the individual, while at the same time clarifying why the ensemble may be of interest. (We also speculate that, at the group level, there may be a fitness advantage to be gained by spreading different beliefs — say, about productive foraging locations — across subjects rather than having the entire population gravitate toward a single “best” belief. This may be a different motivation for sampling-based methods.)

In addition to addressing the empirical problem of fit to the individual, sampling is of course also a possible answer to a more abstract problem with Bayesian models: that they attribute to subjects the capacity for radically intractable calculations. While the simple Kalman filter used here is tractable to begin with, there has been a trend in modeling human and animal learning toward inference about model structure (e.g., recovering not just weights characterizing conditional rewards, but structural variables describing how different latent causes interact to produce observations; [4, 2, 3, 1]). Such inference cannot be accomplished exactly using simple recursive filtering like the Kalman filter. Indeed, it seems hard to imagine how it could be accomplished other than by sequentially sampling one or a small number of hypothetical model structures, since even with the structure known, there remains a difficult parametric inference problem. The present modeling is therefore motivated, in part, toward this setting.

While in our model, subjects don't explicitly track uncertainty about their beliefs from trial to trial, they do maintain hyperparameters (e.g., those controlling the speed of diffusion, the noise of observations, and the probability of jumps) that serve as a sort of constant proxy for uncertainty. We might expect them to adjust these so as to achieve the best performance. (Note that, because the inference is anyway approximate, it is not necessarily the case that veridical, generative settings of these parameters will perform the best.)

Of course, the present model is, by design, only the simplest possible sketch, and there is much work to do in developing it. In particular, it would be useful to develop less extreme models in which subjects either rely on sampling with more particles, or on some combination of sampling and exact inference. We posit that many of the insights developed here will extend to such models, which seem more realistic since *exclusive* use of low-sample particle filtering would be extremely brittle and unreliable. (The example of the Necker cube also invites consideration of Markov Chain Monte Carlo sampling to model nonsequential inference.) However, there is very little information available about individual-level behavior to constrain such a model. The present results on backward blocking stress again the perils of averaging and suggest that data must be analyzed much more delicately if they are ever to bear on issues of distributions and uncertainty. In the case of backward

blocking, if our account is correct, there should be a correlation, over individuals, between the degree to which they initially exhibited a low \hat{w}_B and the degree to which they subsequently exhibited a backward blocking effect. This would be straightforward to test. More generally, there has been a recent trend [17] toward comparing models against raw trial-by-trial data sets according to the cumulative log-likelihood of the data. Although this measure aggregates over trials and subjects, it measures the average goodness of fit, not the goodness of fit to the average, making it much more sensitive for purposes of studying the issues discussed in this article.

References

- [1] T.L. Griffiths and J.B. Tenenbaum. Structure and strength in causal induction. *Cognitive Psychology*, 51:334–384, 2005.
- [2] A. C. Courville, N. D. Daw, and D. S. Touretzky. Similarity and discrimination in classical conditioning: A latent variable account. In *Advances in Neural Information Processing Systems 17*, Cambridge, MA, 2004. MIT Press.
- [3] C. Kemp, A. Perfors, and J.B. Tenenbaum. Learning domain structures. In *Proceedings of the 26th Annual Conference of the Cognitive Science Society*, pages 720–725, 2004.
- [4] A. C. Courville, N. D. Daw, G. J. Gordon, and D. S. Touretzky. Model uncertainty in classical conditioning. In *Advances in Neural Information Processing Systems 16*, Cambridge, MA, 2003. MIT Press.
- [5] S. Kakade and P. Dayan. Acquisition and extinction in autoshaping. *Psychological Review*, 109:533–544, 2002.
- [6] S. Kakade and P. Dayan. Explaining away in weight space. In *Advances in Neural Information Processing Systems 13*, 2001.
- [7] D. R. Shanks. Forward and backward blocking in human contingency judgement. *Quarterly Journal of Experimental Psychology: Comparative & Physiological Psychology*, 37:1–21, 1985.
- [8] P. F. Lovibond, S.-L. Been, C. J. Mitchell, M. E. Bouton, and R. Frohardt. Forward and backward blocking of causal judgment is enhanced by additivity of effect magnitude. *Memory and Cognition*, 31:133–142, 2003.
- [9] C. R. Gallistel, S. Fairhurst, and P. Balsam. The learning curve: Implications of a quantitative analysis. *Proceedings of the National Academy of Sciences of the USA*, 101:13124–13131, 2004.
- [10] A. Doucet, S. Godsill, and C. Andrieu. On sequential Monte Carlo sampling methods for Bayesian filtering. *Statistics and Computing*, 10:197–208, 2000.
- [11] R. E. Kalman. A new approach to linear filtering and prediction problems. *Transactions of the ASME—Journal of Basic Engineering*, 82:35–45, 1960.
- [12] R. A. Rescorla and A. R. Wagner. A theory of Pavlovian conditioning: The effectiveness of reinforcement and non-reinforcement. In A. H. Black and W. F. Prokasy, editors, *Classical Conditioning, 2: Current Research and Theory*, pages 64–69. Appleton Century-Crofts, New York, 1972.
- [13] D. B. Rubin. Using the SIR algorithm to simulate posterior distributions. In J. M. Bernardo, M. H. DeGroot, D. V. Lindley, and A. F. M. Smith, editors, *Bayesian Statistics, Vol. 3*, pages 395–402. Oxford University Press, 1988.
- [14] A. J. Yu and P. Dayan. Expected and unexpected uncertainty: ACh and NE in the neocortex. In *Advances in Neural Information Processing Systems 15*. MIT Press, 2003.
- [15] P. D. Balsam, S. Fairhurst, and C. R. Gallistel. Pavlovian Contingencies and Temporal Information. *Journal of Experimental Psychology: Animal Behavior Processes*, 32:284–295, 2006.
- [16] R. R. Miller and H. Matute. Biological significance in forward and backward blocking: Resolution of a discrepancy between animal conditioning and human causal judgment. *Journal of Experimental Psychology: General*, 125:370–386, 1996.
- [17] N. D. Daw and K. Doya. The computational neurobiology of learning and reward. *Current Opinion in Neurobiology*, 16:199–204, 2006.