

Resource-Rational Account of Sequential Effects in Human Prediction

Arthur Prat-Carrabin^{1,2}, Florent Meyniel³, and Rava Azeredo da Silveira^{1,4,5}

¹Laboratoire de Physique de l'École Normale Supérieure, ENS, Université PSL, CNRS, Sorbonne Université, Université de Paris

²Department of Economics, Columbia University, New York, USA

³Cognitive Neuroimaging Unit, Institut National de la Santé et de la Recherche Médicale, Commissariat à l'Énergie Atomique et aux énergies alternatives, Université Paris-Saclay, NeuroSpin center, 91191 Gif/Yvette, France

⁴Institute of Molecular and Clinical Ophthalmology Basel

⁵Faculty of Science, University of Basel

June 20, 2022

Abstract

An abundant literature reports on ‘sequential effects’ observed when humans make predictions on the basis of stochastic sequences of stimuli. Such sequential effects represent departures from an optimal, Bayesian process. A prominent explanation posits that humans are adapted to changing environments, and erroneously assume non-stationarity of the environment, even if the latter is static. As a result, their predictions fluctuate over time. We propose a different explanation in which sub-optimal and fluctuating predictions result from cognitive constraints (or costs), under which humans however behave rationally. We devise a framework of costly inference, in which we develop two classes of models that differ by the nature of the constraints at play: in one case the precision of beliefs comes at a cost, while in the other beliefs with high predictive power are favored. To compare model predictions to human behavior, we carry out a prediction task that uses binary random stimuli, with probabilities ranging from .05 to .95. Although in this task the environment is static and the Bayesian belief converges, subjects’ predictions fluctuate and are biased toward the recent stimulus history. Both classes of models capture this ‘attractive effect’, but they depart in their characterization of higher-order effects. Only the precision-cost model reproduces a ‘repulsive effect’, observed in the data, in which predictions are biased away from stimuli presented in the more distant past. Our experimental results reveal systematic modulations in sequential effects, which our theoretical approach accounts for in terms of rationality under cognitive constraints.

In many situations of uncertainty, some outcomes are more probable than others. Knowing the probability distributions of the possible outcomes provides an edge that can be leveraged to improve and speed up decision making and perception [1]. In the case of choice reaction-time tasks, it was noted in the early 1950s that human reactions were faster when responding to a stimulus whose probability was higher [2, 3]. In addition, faster responses were obtained after a repetition of a stimulus (i.e., when the same stimulus was presented twice in a row), even in the case of serially-independent stimuli (i.e., when the preceding stimulus carried no information on subsequent ones; [3, 4]). The observation of this seemingly suboptimal behavior has motivated in the following decades a profuse literature on “sequential effects”, i.e., on the dependence of reaction times on the recent history of presented stimuli [5–12]. These studies consistently report a *recency effect* whereby the more often a simple pattern of stimuli (e.g., a repetition) is observed in recent stimulus history, the faster subjects respond to it. In tasks in which subjects are asked to make predictions about sequences of random binary events, sequential effects are also observed and they have given rise since the 1950s to a rich literature [13–25].

Sequential effects are intriguing: why do subjects change their behavior as a function of the recent past observations when those are in fact irrelevant to the current decision? A common theoretical account is that humans infer the statistics of the stimuli presented to them, but because they usually live in environments that change over time, they may believe that the process generating the stimuli is subject to random changes even when it is in fact constant [8, 9, 11, 12]. Consequently, they may rely excessively on the most recent stimuli to predict the next ones. Instead of positing that subjects hold an incorrect belief on the dynamics of the environment and do not learn that it is stationary, we propose a different account, whereby a cognitive constraint is hindering the inference process and prevent it from converging to the correct, constant belief about the unchanging statistics of the environment. This proposal calls for the investigation of the kinds of choice patterns and sequential effects that would result from different cognitive constraints at play during inference.

We derive a framework of constrained inference, in which a cost hinders the representation of belief distributions (posteriors). As for the nature of the cost, we consider two natural hypotheses: first, that it is difficult for subjects to carry computations with high precision, and thus that more precise posteriors come with higher costs; and second, that it is difficult for subjects to deal with uncertain outcomes, and thus that they are averse to unpredictable environments. The first cost results in an inference process in which remote observations are discarded, and beliefs do not converge but fluctuate instead with the recent stimulus history. By contrast, under the second cost, the inference process does converge, although not to the correct (Bayesian) posterior, but rather to a posterior that implies a biased belief on the temporal structure of the stimuli. In both cases, sequential effects emerge as the result of a constrained inference process.

We examine experimentally the degree to which the models derived from our framework account for human behavior, with a task in which we repeatedly ask subjects to predict the upcoming stimulus in sequences of Bernoulli-distributed stimuli. Most studies on sequential effects only consider the equiprobable case, in which the two stimuli have the same probability. However, the models we consider here are more general than this singular case and they

apply to the entire range of stimulus probability. We thus manipulate in separate blocks of trials the *stimulus generative probability* (i.e. the Bernoulli probability that parameterizes the stimulus) to span the range from .05 to .95 by increments of .05. This enables us to examine in detail the behavior of subjects in a large gamut of environments from the singular case of an equiprobable, maximally-uncertain environment (with a probability of .5 for both stimuli) to the strongly-biased, almost-certain environment in which one stimulus occurs with probability .95.

To anticipate on our results, the predictions of subjects depend on the stimulus generative probability, but also on the history of stimuli. We examine whether the occurrence of a stimulus, in past trials, increase the proportion of predictions identical to this stimulus (‘attractive effect’), or whether it decreases this proportion (‘repulsive effect’). The two costs presented above reproduce qualitatively the main patterns in subjects’ data, but they make distinct predictions as to the modulations of the recency effect as a function of the history of stimuli, beyond the last stimulus. We show that the responses of subjects exhibit an elaborate, and at times counter-intuitive, pattern of attractive and repulsive effects, and we compare these to the predictions of our models. Our results suggest that the brain infers a stimulus generative probability, but under a constraint on the precision of its internal representations; the *inferred* generative process may be more general than the actual one, and include higher-order statistics (e.g., transition probabilities), in contrast with the Bernoulli-distributed stimulus used in the experiment.

We present the behavioral task and we examine the predictions of subjects — in particular, how they vary with the stimulus generative probability, and how they depend, at each trial, on the preceding stimulus. We then introduce our framework of inference under constraint, and the two costs we consider, from which we derive two families of models. We examine the behavior of these models and the extent to which they capture the behavioral patterns of subjects. The models make different qualitative predictions about the sequential effects of past observations, which we confront to subjects’ data. We find that the predictions of subjects are qualitatively consistent with a model of inference of conditional probabilities, in which precise posteriors are costly.

Results

Subjects’ predictions of a stimulus increase with the stimulus probability

In a computer-based task, subjects are asked to predict which of two rods the lightning will strike. On each trial, the subject first selects by a key press the left- or right-hand-side rod presented on screen. A lightning symbol (which is here the stimulus) then randomly strikes either of the two rods. The trial is a success if the lightning strikes the rod selected by the subject (Fig. 1a). The location of the lightning strike (left or right) is a Bernoulli random variable whose parameter p (the *stimulus generative probability*) we manipulate across blocks of 200 trials: in each block, p is a multiple of .05 chosen between .05 and .95. Changes of block are explicitly signaled to the subjects. Moreover, in order to capture the ‘stationary’ behavior of subjects, which presumably prevails after ample exposure to the stimulus, each




Figure 1: The prediction task. (a) In each successive trial, the subject is asked to predict which of two rods the lightning will strike. (1) A trial begins. (1') The subject chooses the right-hand-side rod (bold lines), but the lightning strikes the left one. (2) 1.25 seconds after the beginning of the preceding trial, a new trial begins. (2') The subject chooses the right-hand-side rod, and this time the lightning strikes the rod chosen by the subject. The rod and the connected battery light up (yellow), indicating success. (3) A new trial begins. (3') If after 1 second the subject has not made a prediction, a red cross bars the screen, indicating a timeout. (4) A new trial begins. (4') The subject chooses the left-hand-side rod, and the lightning strikes the same rod. In all cases, the duration of a trial is 1.25 seconds. (b) The probability of a lightning strike on the right (respectively left) rod on a given trial is held fixed within a block and denoted p (respectively $1 - p$). To explore the probability range exhaustively, the values of p are evenly spaced from .05 to .95 by increments of .05, and from block to block we sample p without replacement among those values. In order to probe the subjects' predictions in a stationary regime, within each block the first 200 trials consist in passive observation and the 200 following trials are active trials (whose structure is depicted in a).

block is preceded by 200 passive trials in which the stimuli (sampled with the probability chosen for the block) are successively shown with no action from the subject (Fig. 1b). We provide further details on the task in [Methods](#).

The behavior of subjects varies with the stimulus generative probability, p . We denote by 0 and 1 the left and right outcomes of the random strikes, and by $\bar{p}(1)$ the proportion of trials in which a given subject predicts the right-hand-side rod. In the equiprobable condition ($p = .5$), the subjects predict either side on about half the trials ($\bar{p}(1) = .496$; standard error of the mean (sem): .008; p-value of t-test of equality with .5: .59). In the non-equiprobable conditions, the optimal behavior is to predict 1 on none of the trials ($\bar{p}(1) = 0$) if $p < .5$, or

on all trials ($\bar{p}(1) = 1$) if $p > .5$. The proportion of predictions 1 adopted by the subjects also increases as a function of the stimulus generative probability (Pearson correlation coefficient between p and $\bar{p}(1)$, subjects pooled: .97; p-value: 3.3e-6), but not as steeply: it lies between the stimulus generative probability p , and the optimal response 0 (if $p < .5$) or 1 (if $p > .5$; Fig. 2a).

First-order sequential effects: attractive influence of the most recent stimulus on subjects' predictions

The sequences presented to subjects correspond to independent, Bernoulli-distributed random events. Having shown that the subjects' predictions follow (in a non optimal fashion) the stimulus generative probability, we now test whether they also exhibit the non-independence of consecutive trials featured by the Bernoulli process. Under this hypothesis and in the stationary regime, the proportion of predictions 1 conditional on the preceding stimulus being a 1, $\bar{p}(1|1)$, should be no different than the proportion of predictions 1 conditional on the preceding stimulus being a 0, $\bar{p}(1|0)$. In other words, conditioning on the preceding stimulus should have no effect. In subjects' responses, however, these two conditional proportions are markedly different for all stimulus generative probabilities (Fisher exact test: all p-values $< 1e-10$; Fig. 2a). Both quantities increase as a function of the stimulus generative probability, but the proportions of predictions 1 conditional on a 1 are consistently greater than the proportions of predictions 1 conditional on a 0 (Fig. 2b). (We note that because the stimulus is either 0 or 1, it follows that, symmetrically, the proportions of predictions 0 conditional on a 0 are consistently greater than the proportions of predictions 0 conditional on a 1.) In other words, the preceding stimulus has an 'attractive' sequential effect.

A framework of costly inference

The attractive effect of the preceding stimulus on subjects' responses suggests that the subjects have not correctly inferred the Bernoulli statistics of the process generating the stimuli. We investigate the hypothesis that their ability to infer the underlying statistics of the stimuli is hampered by cognitive constraints. We assume that these constraints can be understood as a cost, bearing on the representation, by the brain, of the subject's beliefs about the statistics. Specifically, we derive an array of models from a framework of inference under costly posteriors [26], which we now present. We consider a model subject who is presented on each trial t with a stimulus $x_t \in \{0, 1\}$ and who uses the sequence of stimuli $x_{1:t} = (x_1, \dots, x_t)$ to infer the stimulus generative probability, over which she holds the belief distribution \hat{P}_t . A Bayesian observer equipped with this belief \hat{P}_t and observing a new observation x_{t+1} would obtain its updated belief P_{t+1} through Bayes' rule. However, a cognitive cost $C(P)$ hinders our model subject's ability to represent probability distributions P . Thus she approximates the Bayesian posterior P_{t+1} through another distribution \hat{P}_{t+1} that minimizes a loss function L defined as

$$L(\hat{P}_{t+1}) = D(\hat{P}_{t+1}; P_{t+1}) + \lambda C(\hat{P}_{t+1}), \quad (1)$$

where D is a measure of distance between two probability distributions, and $\lambda \geq 0$ is a coefficient specifying the relative weight of the cost. (We are not proposing that subjects

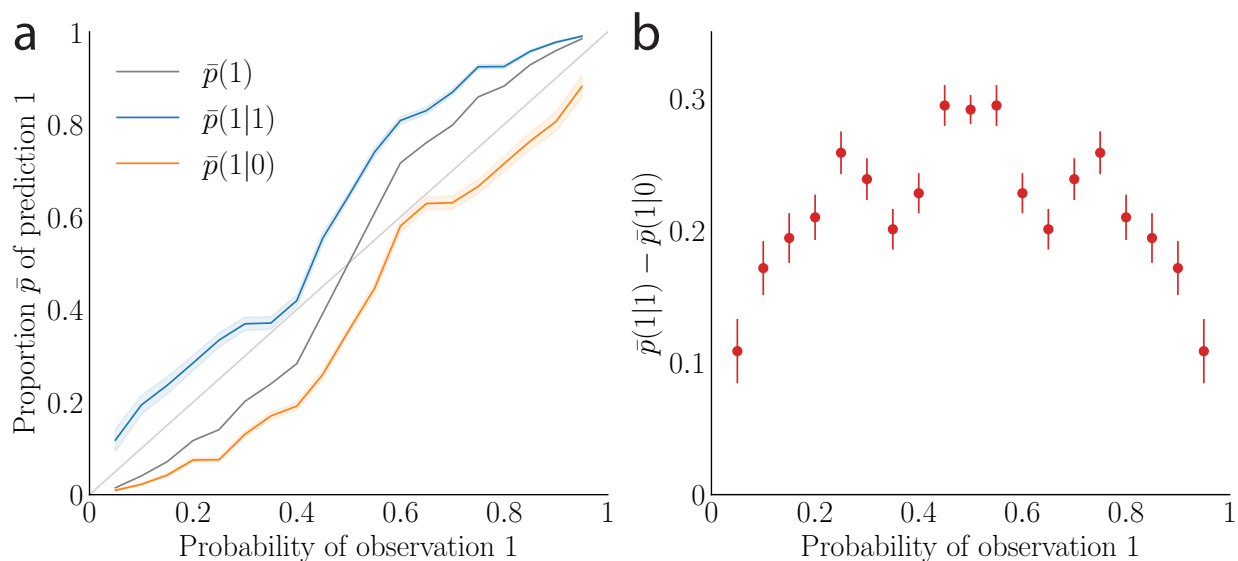


Figure 2: Across all stimulus generative probabilities, subjects are more likely than average to make a prediction equal to the preceding observation. (a) Proportion of predictions 1 in subjects responses as a function of the stimulus generative probability, conditional on observing a 1 (blue line) or a 0 (orange line) on the preceding trial, and unconditional (grey line). Note that ‘left’ and ‘right’ predictions are pooled together as long as the probabilities of the corresponding observations are the same; thus the grey line is symmetric about the middle point, and the blue and orange lines are reflections of each other. The widths of the shaded areas indicate the standard error of the mean. (b) Difference between the proportion of predictions 1 conditional on the preceding observation being a 1, and the proportion of predictions 1 conditional on the preceding observation being a 0. This difference is positive across all stimulus generative probabilities (all t-test p-values $< 1e-15$), i.e., observing a 1 at the preceding trial increases the probability of predicting a 1. Bars indicate the standard error of the mean.

actively minimize this quantity, but rather that the brain’s inference process is an effective solution to this optimization problem.) Below, we use the Kullback-Leibler divergence for the distance (i.e., $D(\hat{P}_{t+1}; P_{t+1}) = D_{KL}(\hat{P}_{t+1}||P_{t+1})$). If $\lambda = 0$, the solution to this minimization problem is the Bayesian posterior; if $\lambda \neq 0$, the cost distorts the Bayesian solution in ways that depend on the form of the cost borne by the subject (we detail further below the two kinds of costs we investigate).

In our framework, the subject assumes that the m preceding stimuli ($x_{t-m+1:t}$ with $m \geq 0$) and a vector of parameters q together determine the distribution of the stimulus at trial $t+1$, $p(x_{t+1}|x_{t-m+1:t}, q)$. Although in our task the stimuli are Bernoulli-distributed (thus they do not depend on preceding stimuli) and a single parameter determines the probability of the outcomes (the *stimulus generative probability*), the subject may admit the possibility that more complex mechanisms govern the statistics of the stimuli, e.g., transition probabilities between consecutive stimuli. Therefore, the vector q may contain more than one parameter and the number m of preceding stimuli assumed to influence the probability of the following stimulus, which we call the ‘Markov order’, may be greater than 0.

Below, we call ‘Bernoulli observer’ any model subject who assumes that the stimuli are Bernoulli-distributed ($m = 0$); in this case the vector q consists of a single parameter that determines the probability of observing a 1, which we also denote by q for the sake of concision. The bias and variability in the inference of the Bernoulli observer is studied in Ref. [26]. We call ‘Markov observer’ any model subject who posits that the probability of the stimulus depends on the preceding stimuli ($m > 0$). In this case, the vector q contains the 2^m conditional probabilities of observing a 1 after observing each possible sequence of m stimuli. For instance, with $m = 1$ the vector q is the pair of parameters (q_0, q_1) where q_i is the probability of observing a 1 after observing the stimulus i . In the absence of a cost, the belief over the parameter(s), $\hat{P}_t(q)$, eventually converges towards the parameter vector that is consistent with the generative Bernoulli statistics governing the stimulus (except if the prior precludes this parameter vector). Below, we assume a uniform prior.

The ways in which the Bayesian inference is distorted by our model subject depend on the nature of the cost that weighs on the inference process. Although many assumptions could be made on the kind of constraint that hinders human inference, and on the cost it would entail in our framework, here we examine two costs that stem from two possible principles: that the cost is a function of the beliefs held by the subject, or that it is a function of the environment that the subject is inferring. We detail, below, these two costs.

Unpredictability cost A first hypothesis is that the subjects favor, in their inference, parameter vectors q that correspond to more predictable outcomes. We quantify the outcome unpredictability by the Shannon entropy [27] of the outcome implied by the vector of parameters q , which we denote by $H(X; q)$. The cost associated with the distribution $\hat{P}(q)$ is the expectation of this quantity averaged over beliefs i.e.,

$$C_u(\hat{P}) = \mathbb{E}_{\hat{P}}[H(X; q)] = \int H(X; q) \hat{P}(q) dq, \quad (2)$$

which we call the ‘unpredictability cost’. For a Bernoulli observer, a posterior concentrated on extreme values of the Bernoulli parameter (toward 0 or 1), thus representing more predictable environments, comes with a lower cost than a posterior concentrated on values of the Bernoulli parameter close to 0.5, which correspond to the most unpredictable environments (Fig. 3a). The loss function (Eq. (1)) under this cost is minimized by the posterior

$$\hat{P}_{t+1}(q) \propto \hat{P}_t(q) e^{-\lambda H(X; q)}. \quad (3)$$

In the Bernoulli case ($m = 0$), assuming a uniform prior ($\hat{P}_0(q) \propto 1$), the posterior after t stimuli has the form $\hat{P}_t(q) \propto [\varphi(q; r_t)]^t$, where r_t is the proportion of stimuli 1 observed up to trial t , i.e. $r_t = \sum_{i=1}^t x_i/t$, and

$$\varphi(q; r_t) = q^{r_t} (1 - q)^{1-r_t} e^{-\lambda H(X; q)}. \quad (4)$$

This function has a global maximum $q^*(r_t)$ and as the number of presented stimuli t grows the posterior \hat{P}_t becomes concentrated around this maximum. The proportion of 1s, r_t , naturally converges to the stimulus generative probability p , thus our subject’s inference converges towards the value $q^*(p)$ which is different from the true value p in the non-equiprobable case

($p \neq .5$). The equiprobable case ($p=.5$) is singular, in that with a weak cost ($\lambda < 1$) the inferred probability is unbiased ($q^*(p) = .5$), while with a strong cost ($\lambda > 1$) the inferred probability does not converge but instead alternates between two values above and below .5; see Ref. [26]. In other words, except in the equiprobable case, the inference converges but it is biased, i.e., the posterior peaks at an incorrect value of the stimulus generative probability (Fig. 3c). This value is closer to the extremes (0 and 1) than the stimulus generative probability, i.e., it implies an environment more predictable than the actual one.

In the case of a Markov observer ($m > 0$), the posterior also converges to a vector of parameters q which implies not only a bias but also that the conditional probabilities of a 1 (conditioned on different stimulus histories) are not equal. The prediction of the next stimulus being 1 on a given trial depends on whether the preceding stimulus was a 0 or a 1: this model therefore predicts sequential effects. We further examine below the behavior of this model in the cases of a Bernoulli observer and of different Markov observers. We refer the reader interested in more details on the Markov models, including their mathematical derivations, to the [Methods](#) section.

In short, with the unpredictability-cost models, when $p \neq .5$, the inference process converges to an asymptotic posterior $q^*(p)$ which does not itself depend on the history of the stimulus, but that is biased. In particular, for Markov observers ($m > 0$), the asymptotic posterior corresponds to an erroneous belief about the dependency of the stimulus on the recent stimulus history, which results in sequential effects in behavior. We now turn to another family of models in which the posterior itself depends on the recent stimulus history, and thus does not asymptotically converge.

Precision cost A different hypothesis about the inference process of subjects is that the brain mobilizes resources to represent probability distributions, and that more ‘precise’ distributions require more resources. We write the cost associated with a distribution, $\hat{P}(q)$, as the negative of its entropy,

$$C_p(\hat{P}) = -H[\hat{P}(q)] = \int \hat{P}(q) \ln \hat{P}(q) dq, \quad (5)$$

which is a measure of the amount of certainty in the distribution. Wider (less concentrated) distributions provide less information about the probability parameter and are thus less costly than narrower (more concentrated) distributions (Fig. 3b). As an extreme case, the uniform distribution is the least costly.

With this cost, the loss function (Eq. (1)) is minimized by the distribution equal to the Bayesian posterior raised to the exponent $1/(\lambda + 1)$, and renormalized, i.e.,

$$\hat{P}_{t+1}(q) \propto \left[\hat{P}_t(q) p(x_{t+1} | x_{t-m+1:t}, q) \right]^{1/(\lambda+1)}. \quad (6)$$

Since λ is strictly positive, the exponent is positive and lower than 1. As a result, the solution ‘flattens’ the Bayesian posterior, and in the extreme case of an unbounded cost ($\lambda \rightarrow \infty$) the posterior is the uniform distribution. Furthermore, in the expression of our model subject’s posterior, the likelihood $p(x_{t+1} | x_{t-m+1:t}, q)$ is raised after k trials to the exponent $1/(\lambda+1)^{k+1}$, it thus decays to zero as the number k of new stimuli increases. One can interpret this effect

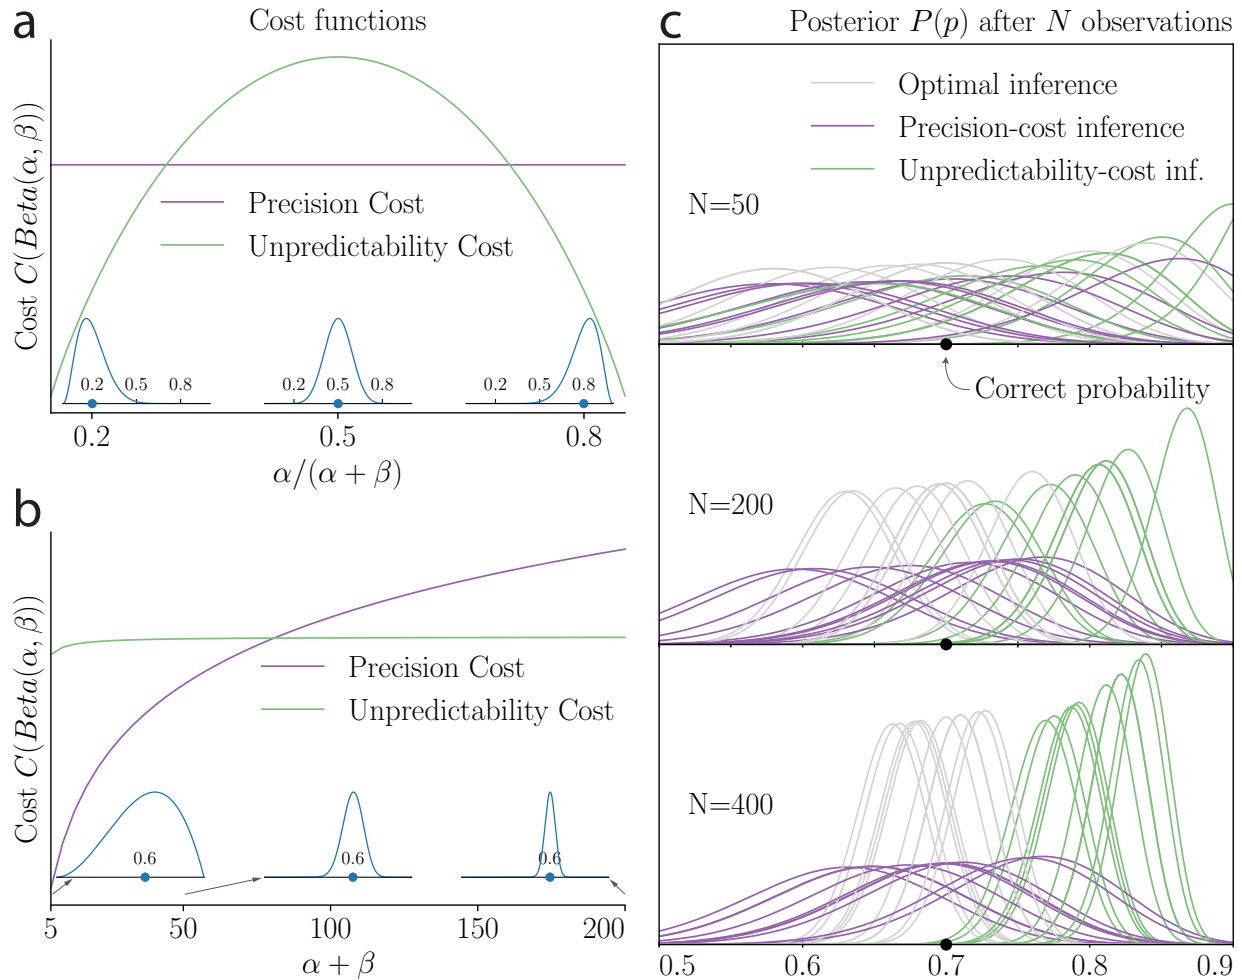


Figure 3: Illustration of the Bernoulli-observer models, with unpredictability and precision costs. (a) Precision cost (purple) and unpredictability cost (green lines) of a Beta distribution with parameters α and β , as functions of the mean of the distribution, $\alpha/(\alpha + \beta)$, and keeping the entropy constant. The precision cost is the negative of the entropy and it is thus constant, regardless of the mean of the distribution. The unpredictability cost is larger when the mean of the distribution is closer to 0.5 (i.e., when unpredictable environments are likely, under the distribution). *Insets:* Beta distributions with mean .2, .5 and .8, and constant entropy. (b) Costs as functions of the sample size parameter, $\alpha + \beta$. A larger sample size implies a higher precision and lower entropy, thus the precision cost increases as a function of the sample size, whereas the unpredictability cost is less sensitive to changes in this parameter. *Insets:* Beta distributions with mean .6 and sample size parameter, $\alpha + \beta$, equal to 5, 50 and 200. (c) Posteriors $P(p)$ of an optimal observer (gray), a precision-cost observer (purple) and an unpredictability-cost observer (green lines), after the presentation of ten sequences of $N = 50, 200$, and 400 observations sampled from a Bernoulli distribution of parameter 0.7. The posteriors of the optimal observer narrow as evidence is accumulated, and the different posteriors obtained after different sequences of observations are drawn closer to each other and to the correct probability. The posteriors of the unpredictability-cost observer also narrow and group together, but around a probability larger (less unpredictable) than the correct probability. Precise distributions are costly to the precision-cost observer and thus the posteriors do not narrow after long sequences of observations. Moreover, as uncertainty is maintained throughout inference, the posteriors fluctuate with the recent history of the stimuli and they do not converge.

as gradually forgetting past observations. The posterior $\hat{P}_t(q)$ is thus dominated by the recent history of the stimuli. It does not converge, but instead fluctuates with the recent stimulus history (Fig. 3c). Hence, this model implies predictions about subsequent stimuli that depend on the stimulus history, i.e., it predicts sequential effects.

Overview of the inference models Although the two families of models derived from the two costs both potentially generate sequential effects, they do so by giving rise to qualitatively different inference processes. Under the unpredictability cost, the inference converges to a posterior that, in the Bernoulli case ($m = 0$), implies a biased estimate of the stimulus generative probability, while in the Markov case ($m > 0$) it implies the belief that there are serial dependencies in the stimuli: predictions therefore depend on the recent stimulus history. By contrast, the precision cost prevents beliefs from converging. As a result, the subject’s predictions vary with the recent stimulus history. This inference process amounts to an exponential discount of remote observations, or equivalently, to the overweighting of recent observations.

To investigate in more detail the sequential effects that these two costs produce, we implement two families of inference models derived from the two costs. Each model is characterized by the type of cost (unpredictability cost or precision cost), and by the assumed Markov order (m): we examine the case of a Bernoulli observer ($m = 0$) and three cases of Markov observers (with $m = 1, 2$, and 3). We thus obtain $2 \times 4 = 8$ models of inference. Each of these models has one parameter λ controlling the weight of the cost.

Response-selection strategy We assume that the subject’s response on a given trial depends on the inferred posterior according to a generalization of ‘probability matching’ implemented in other studies [28–30]. In this response-selection strategy, the subject predicts a 1 with the probability $\bar{p}_t^\kappa / (\bar{p}_t^\kappa + (1 - \bar{p}_t)^\kappa)$, where \bar{p}_t is the expected probability of observing a 1 derived from the posterior, i.e., $\bar{p}_t \equiv \int p(x_{t+1} = 1 | x_{t-m+1:t}, q) \hat{P}_t(q) dq$. The single parameter κ controls the randomness of the response: with $\kappa = 0$ the subject chooses 0 and 1 with equal probability; with $\kappa = 1$ the response-selection strategy corresponds to probability matching, i.e., the subject chooses 1 with probability \bar{p}_t ; and as κ increases toward infinity the choices become optimal, i.e., the subject responds 1 if the expected probability of observing a 1, \bar{p}_t , is greater than .5, and 0 if it is lower than .5 (if $\bar{p}_t = .5$ the subject chooses 0 or 1 with equal probability).

In our investigations, we also implement three other response-selection strategies. First, a strategy based on a ‘softmax’ function that smoothes the optimal decision rule; it does not yield, however, a behavior substantially different from that of the generalized probability-matching response-selection strategy. Second, we examine a strategy in which the model subject chooses the optimal response with a probability that is fixed across conditions, which we fit onto her choices. No subject is best-fitted by this strategy. Third, another possible strategy proposed in the game-theory literature [31] is ‘win-stay, lose-shift’: it prescribes to repeat the same response as long as it proves correct and to change otherwise. In the context of our binary-choice prediction task, it is indistinguishable from a strategy in which the model subject chooses a prediction equal to the outcome that last occurred. This strategy is a special case of our Bernoulli observer hampered by a precision cost whose weight λ is

large combined with the optimal response-selection strategy ($\kappa \rightarrow \infty$). Since the generalized probability-matching strategy parameterized by the exponent κ appears either more general, better than or indistinguishable from those other response-selection strategies, we selected it to obtain the results presented below.

Model fitting favors Markov-observer models

Each of our eight models has two parameters: the factor weighting the cost, λ , and the exponent of the generalized probability-matching, κ . We fit the parameters of each model to the responses of each subject, by maximizing their likelihoods. We find that 60% of subjects are best fitted by one of the unpredictability-cost models, while 40% are best fitted by one of the precision-cost models. When pooling the two types of cost, 65% of subjects are best fitted by a Markov-observer model. We implement a ‘Bayesian model selection’ procedure [32], which takes into account, for each subject, the likelihoods of all the models (and not only the maximum) in order to obtain a Bayesian posterior over the distribution of models in the general population (see [Methods](#)). The derived expected probability of unpredictability-cost models is 57% (and 43% for precision-cost models) with an exceedance probability (i.e., probability that unpredictability-cost models are more frequent in the general population) of 78%. The expected probability of Markov-observer models, regardless of the cost used in the model, is 70% (and 30% for Bernoulli-observer models) with an exceedance probability (i.e., probability that Markov-observer models are more frequent in the general population) of 98%. These results indicate that the responses of subjects are generally consistent with a Markov-observer model, although the stimuli used in the experiment are Bernoulli-distributed. As for the unpredictability-cost and the precision-cost families of models, Bayesian model selection does not provide decisive evidence in favor of either model, indicating that they both capture some aspects of the responses of the subjects. Below, we examine more closely the behaviors of the models, and point to qualitative differences between the predictions resulting from each model family.

Before turning to these results, we validate the robustness of our model-fitting procedure with two additional analyses. First, we estimate a confusion matrix (see [Methods](#)) to examine the possibility that the model-fitting procedure could misidentify the models which generated test sets of responses. We find that the best-fitting model corresponds to the true model in at least 70% of simulations (the chance level is $12.5\% = 1/8$ models), and actually more than 90% for the majority of models (Fig. 9).

Second, we seek to verify whether the best-fitting cost factor, λ , that we obtain for each subject is consistent across the range of probabilities tested. Specifically, we fit separately the models to the responses obtained in the blocks of trials whose stimulus generative probability was ‘medium’ (between 0.3 and 0.7, included) on the one hand, and to the responses obtained when the probability was ‘extreme’ (below 0.3, and above 0.7) on the other hand; and we compare the values of the best-fitting cost factors λ in these two cases. More precisely, for the precision-cost family, we look at the decay time, equal to $1/\ln(1 + \lambda)$, which is the characteristic time over which the model subject ‘forgets’ past observations. With both families of models, we find that on a logarithmic scale the parameters in the medium- and extreme-probabilities cases are significantly correlated across subjects (Pearson’s r , precision-cost models: 0.58, p-value: $7e-3$; unpredictability-cost models: $r = .74$, p-value: $2e-4$). In

other words, if a subject is best fitted by a large cost factor in medium-probabilities trials, he or she is likely to be also best fitted by a large cost factor in extreme-probabilities trials. This indicates that our models capture idiosyncratic features of subjects that generalize across conditions instead of varying with the stimulus probability (Fig. 10; see [Methods](#)).

Models of costly inference reproduce the attractive effect of the most recent stimulus

We now examine the behavioral patterns resulting from the models. Looking at the responses of the models that best fit each subject, we find that the proportion of predictions 1, $\bar{p}(1)$, increases smoothly as a function of the stimulus generative probability p ; it is below it at small probabilities ($\bar{p}(1) < p$ when $p < .5$) and above it at high probabilities ($\bar{p}(1) > p$ when $p > .5$; Fig. 4a). The proportions of predictions 1 conditional on the preceding stimulus being a 1, $\bar{p}(1|1)$, are larger at all probabilities than their counterparts conditional on a 0, $\bar{p}(1|0)$ (Fig. 4b). In other words, the predictions of the model subjects are biased towards the preceding stimulus. The models of constrained inference, thus, reproduce the attractive sequential effect observed in subjects' responses.

In order to investigate the sequential effects exhibited by the two model families (precision-cost models and unpredictability-cost models), we first examine separately the behaviors of the subjects whose responses are best-fitted by each of the two costs. The subjects in both cases exhibit the attractive sequential effect, i.e., the proportion of predictions 1 following the observation of a 1 is greater than the same proportion after observing a 0. The behaviors differ quantitatively in the non-equiprobable conditions: with the subjects best-fitted by unpredictability-cost models, if the preceding stimulus (e.g. 1) corresponds to the least likely outcome (e.g. $p < .5$), then the proportion of predictions equal to the preceding stimulus is close to the stimulus generative probability (e.g., $\bar{p}(1|1) \approx p$), while with the subjects best-fitted by precision-cost models, this proportion is larger than the stimulus generative probability ($\bar{p}(1|1) > p$; Fig. 4c, first column). The best-fitting models derived from the two costs reproduce these behavioral patterns. Notably, the models derived from both costs yield the attractive sequential effect. Moreover, they differ in the proportion of predictions equal to the preceding stimulus, when the latter is least likely, consistently with the subjects' predictions that they fit (Fig. 4c).

In addition, we investigate for each subject whose responses are best fit by a precision-cost model, the best-fitting unpredictability-cost model, and vice-versa. We find that the precision-cost models reproduce the behavioral patterns of the subjects that are best fit by unpredictability-cost models (in particular, the proportion of predictions equal to the least likely outcome), while the converse is not verified, i.e., the models in the unpredictability-cost family do not reproduce well the behavior of the subjects that are best fit by precision-cost models (Fig. 4c, middle and last columns). In short, the precision-cost family of models seems more able to capture the heterogeneity in the behavior of the subjects.

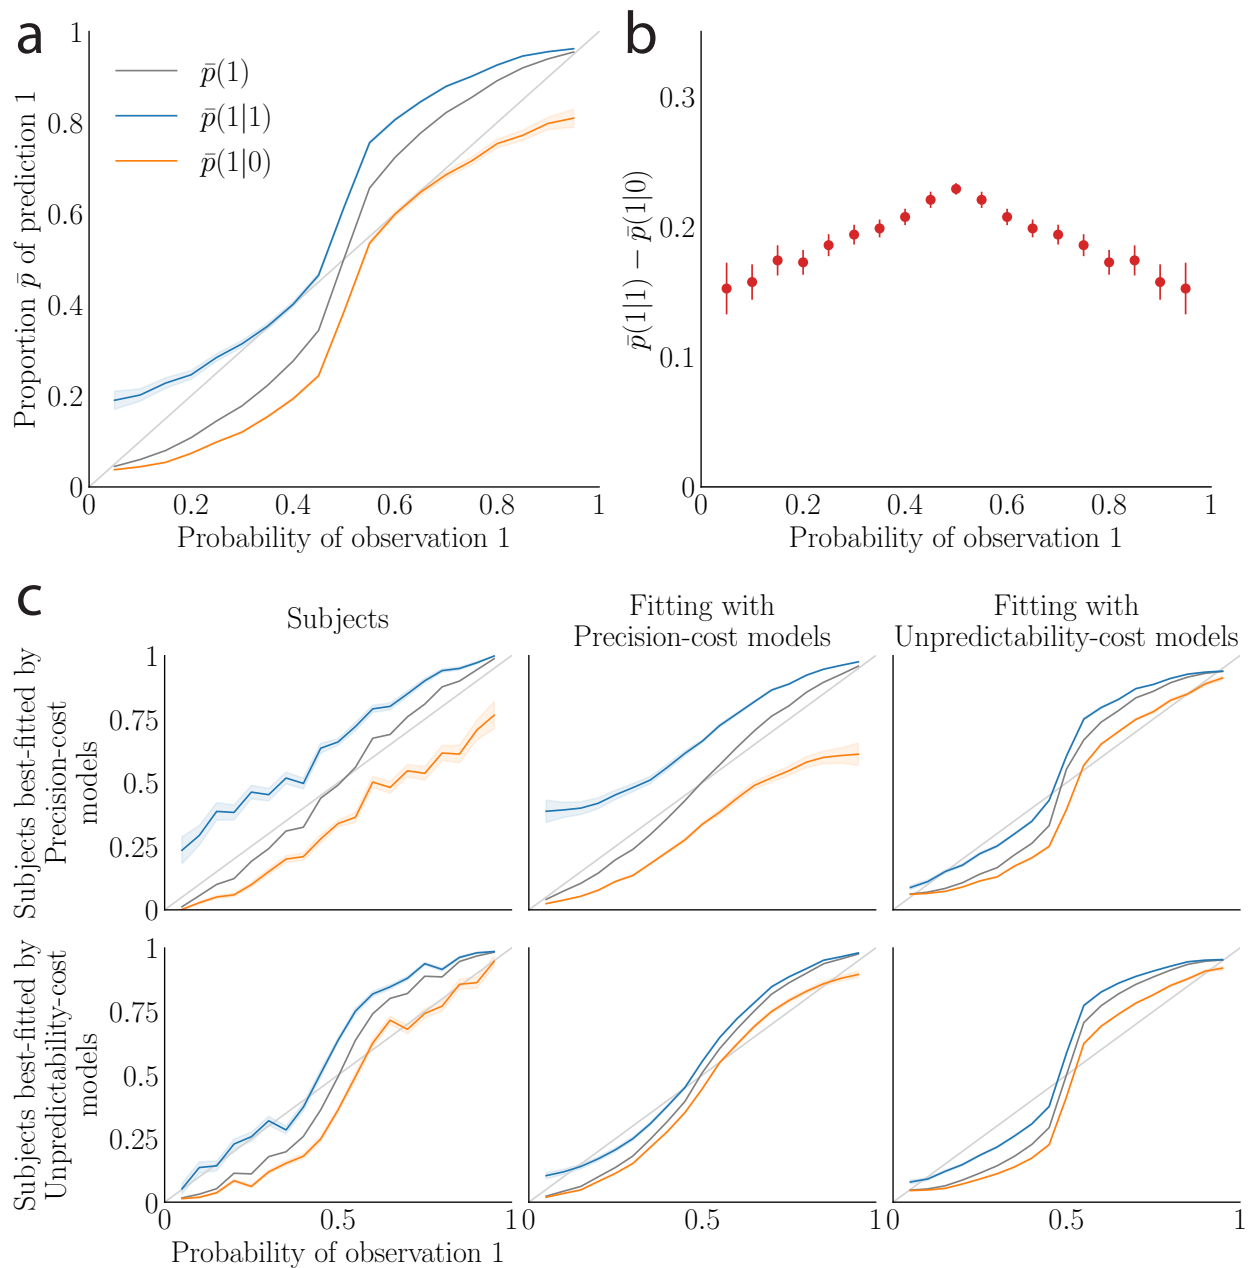


Figure 4: The best-fitting precision-cost and unpredictability-cost models reproduce the subjects' attractive sequential effects. (a) Proportion of predictions 1 in the model subjects' responses as a function of the stimulus generative probability, conditional on observing a 1 (blue line) or a 0 (orange line) on the preceding trial, and unconditional (grey line). The responses of each model best-fitting each subject are pooled together. (b) Difference, in the model subjects' responses, between the proportion of predictions 1 conditional on the preceding observation being a 1, and the proportion of predictions 1 conditional on the preceding observation being a 0. This difference is positive across all stimulus generative probabilities, indicating an attractive sequential effect (i.e., observing a 1 at the preceding trial increases the model subjects' probability of predicting a 1). (c) The subjects are divided in two groups: those whose responses are best-fitted by a prediction-cost models (top row) and those whose responses are best-fitted by an unpredictability-cost model (bottom row). The sequential effects of these two groups are shown in the panels in the left column. The panels in the other two columns show the sequential effects in the responses of the prediction-cost models fitted to each group of subjects (middle column), and of the unpredictability-cost models fitted to each group of subjects (right column).

Beyond the most recent stimulus: patterns of higher-order sequential effects

Notwithstanding the quantitative differences just presented, both families of models yield qualitatively similar attractive sequential effects: the model subjects' predictions are biased towards the preceding stimulus. Does this pattern also apply to the longer history of the stimulus, i.e., do more distant trials also influence the model subjects' predictions? To investigate this hypothesis, we examine the difference between the proportion of predictions 1 after observing a sequence of length n that starts with a 1, minus the proportion of predictions 1 after the same sequence, but starting with a 0, i.e. $\bar{p}(1|1x) - \bar{p}(1|0x)$, where x is a sequence of length $n - 1$, and $1x$ and $0x$ denote the same sequence preceded by a 1 and by a 0. This quantity enables us to isolate the influence of the n -to-last stimulus on the current prediction. If the difference is positive, the effect is 'attractive'; if it is negative, the effect is 'repulsive' (in this latter case, the presentation of a 1 decreases the probability that the subjects predicts a 1 in a later trial, as compared to the presentation of a 0); and if the difference is zero there is no sequential effect stemming from the n -to-last stimulus. The case $n = 1$ corresponds to the immediately preceding stimulus, whose effect we have shown to be attractive, i.e. $\bar{p}(1|1) - \bar{p}(1|0) > 0$, in the responses both of the best-fitting models and of the subjects (Figs. 2b and 4b).

We investigate the effect of the n -to-last stimulus on the behavior of the two families of models, with $n = 1, 2$, and 3. We present here the main results of this investigation; we refer the reader to [Methods](#) for a more detailed analysis. With unpredictability-cost models of Markov order m , there are non-vanishing sequential effects stemming from the n -to-last stimulus only if the Markov order is greater than or equal to the distance from this stimulus to the current trial, i.e. if $m \geq n$. In this case, the sequential effects are attractive (Fig. 5).

With precision-cost models, the n -to-last stimuli yield non-vanishing sequential effects regardless of the Markov order, m . With $n = 1$, the effect is attractive, i.e., $\bar{p}(1|1) - \bar{p}(1|0) > 0$. With $n = 2$ (second-to-last stimulus), the effect is also attractive, i.e., in the case of the pair of sequences '11' and '01', $\bar{p}(1|11) - \bar{p}(1|01) > 0$ (Fig. 5a). By symmetry, the difference is also positive for the other pair of relevant sequences, '10' and '00' (e.g., we note that $\bar{p}(1|10) = 1 - \bar{p}(0|10)$, and that $\bar{p}(0|10)$ when the probability of a 1 is p is equal to $\bar{p}(1|01)$ when the probability of a 1 is $1 - p$).

As for the third-to-last stimulus ($n = 3$), it can be followed by four different sequences of length two, but we only need to examine two of these four; the other two are obtained by symmetry (as presented just above, the probabilities of predictions 1 conditional on two of these sequences are functions of the probabilities of 1 conditional on the other two sequences). We find that for the precision-cost models, with all the Markov orders we examine (from 0 to 3), the probability of predicting a 1 after observing the sequence '111' is greater than that after observing the sequence '011', i.e., $\bar{p}(1|111) - \bar{p}(1|011) > 0$, that is, there is an attractive sequential effect of the third-to-last stimulus if the sequence following it is '11' (and, by symmetry, if it is '00'; Fig. 5b). So far, thus, we have found only attractive effects. However, the results are less straightforward when the third-to-last stimulus is followed by the sequence '01' (or '10'). In this case, for a Bernoulli observer ($m = 0$), the effect is also attractive: $\bar{p}(1|101) - \bar{p}(1|001) > 0$. With Markov observers ($m \geq 1$), over a range of stimulus generative probability p , the effect is repulsive: $\bar{p}(1|101) - \bar{p}(1|001) < 0$, that is, the

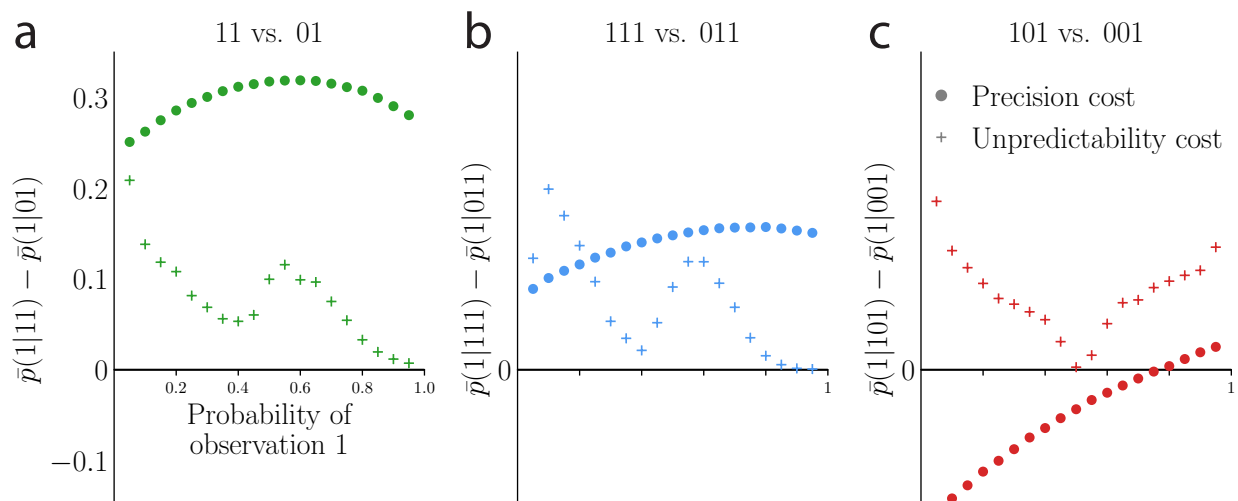


Figure 5: Higher-order sequential effects: the precision-cost model of a Markov observer predicts a repulsive effect of the third-to-last stimulus. Sequential effect of the second-to-last (a) and third-to-last (b,c) stimuli, in the responses of the precision-cost model of a Markov observer with $m = 1$ (points), and of the unpredictability-cost model of a Markov observer with $m = 3$ (crosses). (a) For the two models, difference between the proportion of prediction 1 conditional on observing ‘11’, and conditional on observing ‘01’, i.e., $\bar{p}(1|11) - \bar{p}(1|01)$, as a function of the stimulus generative probability. With both models, this difference is positive, indicating an attractive sequential effect of the second-to-last stimulus. (b) Difference between the proportion of prediction 1 conditional on observing ‘111’, and conditional on observing ‘011’, i.e., $\bar{p}(1|111) - \bar{p}(1|011)$. The positive difference indicates an attractive sequential effect of the third-to-last stimulus in this case. (c) Difference between the proportion of prediction 1 conditional on observing ‘101’, and conditional on observing ‘001’, i.e., $\bar{p}(1|101) - \bar{p}(1|001)$. With the precision-cost model of a Markov observer, the negative difference when the stimulus generative probability is lower than .8 indicates a repulsive sequential effect of the third-to-last stimulus in this case, while when the probability is greater than .8, and with the unpredictability-cost model of a Markov observer, the positive difference indicates an attractive sequential effect of the third-to-last stimulus.

presentation of a 1 *decreases* the probability that the model subject predicts a 1 three trials later, as compared to the presentation of a 0 (Fig. 5c). The occurrence of the repulsive effect in this particular case is a distinctive trait of the precision-cost models of Markov observers ($m \geq 1$); we do not obtain any repulsive effect with any of the unpredictability-cost models, nor with the precision-cost model of a Bernoulli observer ($m = 0$).

Subjects’ predictions exhibit higher-order repulsive effects

We now examine the sequential effects in subjects’ responses, beyond the attractive effect of the preceding stimulus ($n = 1$; discussed above). With $n = 2$ (second-to-last stimulus), for the majority of the 19 stimulus generative probabilities p , we find attractive sequential effects: the difference $\bar{p}(1|11) - \bar{p}(1|01)$ is significantly positive (Fig. 6a; p-values $< .01$ for 11 stimulus

generative probabilities, $< .05$ for 13 probabilities). With $n = 3$ (third-to-last stimulus), we also find significant attractive sequential effects in subjects' responses for some of the stimulus generative probabilities, when the third-to-last stimulus is followed by the sequence '11' (Fig. 6b; p-values $< .01$ for four probabilities, $< .05$ for seven probabilities). When it is instead followed by the sequence '01', we find that for eight stimulus generative probabilities, all between .25 and .75, there is a significant repulsive sequential effect: $\bar{p}(1|101) - \bar{p}(1|001) < 0$ (p-values $< .01$ for six probabilities, $< .05$ for eight probabilities). Thus, in these cases, the occurrence of a 1 as the third-to-last stimulus increases (in comparison with the occurrence of a 0) the proportion of the *opposite* prediction, 0. For the remaining stimulus generative probabilities, this difference is in most cases also negative although not significantly different from zero (Fig. 6c).

The repulsive sequential effect of the third-to-last stimulus in subjects' predictions only occurs when the third-to-last stimulus is a 1 followed by the sequence '01' (or a 0 followed by '10'). It is also only in this case that the repulsive effect appears with the precision-cost models of a Markov observer (while it never appears with the unpredictability-cost models). This qualitative difference suggests that the precision-cost models offers a better account of sequential effects in subjects. However, model-fitting onto the overall behavior presented above showed that a fraction of the subjects is better fitted by the unpredictability-cost models. We investigate, thus, the presence of a repulsive effect in the predictions of the subjects best fitted by the precision-cost models, and of those best fitted by the unpredictability-cost models. For the subjects best fitted by the precision-cost models, we find (expectedly) that there is a significant repulsive sequential effect of the third-to-last stimulus ($\bar{p}(1|101) - \bar{p}(1|001) < 0$; p-values $< .01$ for two probabilities, $< .05$ for four probabilities; Fig. 6d, left panel). For the subjects best fitted by the unpredictability-cost models (a family of model that does not predict any repulsive sequential effects), we also find, perhaps surprisingly, a significant repulsive effect of the third-to-last stimulus (p-values $< .01$ for three probabilities, $< .05$ for five probabilities), which demonstrates the robustness of this effect (Fig. 6d, right panel). Thus, in spite of the results of the model-selection procedure, some sequential effects in subjects' predictions support only one of the two families of model. Regardless of the model that best fits their overall predictions, the behavior of the subjects is consistent only with the precision-cost family of models with Markov order equal to or greater than 1, i.e., with a model of inference of conditional probabilities hampered by a cognitive cost weighing on the precision of belief distributions.

Discussion

We investigated the hypothesis that sequential effects in human predictions result from cognitive constraints hindering the inference process carried out by the brain. We devised a framework of constrained inference, in which the model subject bears a cognitive cost when updating its belief distribution upon the arrival of new evidence: the larger the cost, the more the subject's posterior differs from the Bayesian posterior. The models we derive from this framework make specific predictions. First, the proportion of forced-choice predictions for a given stimulus should increase with the stimulus generative probability. Second, most of those models predict sequential effects: predictions also depend on the recent stimulus

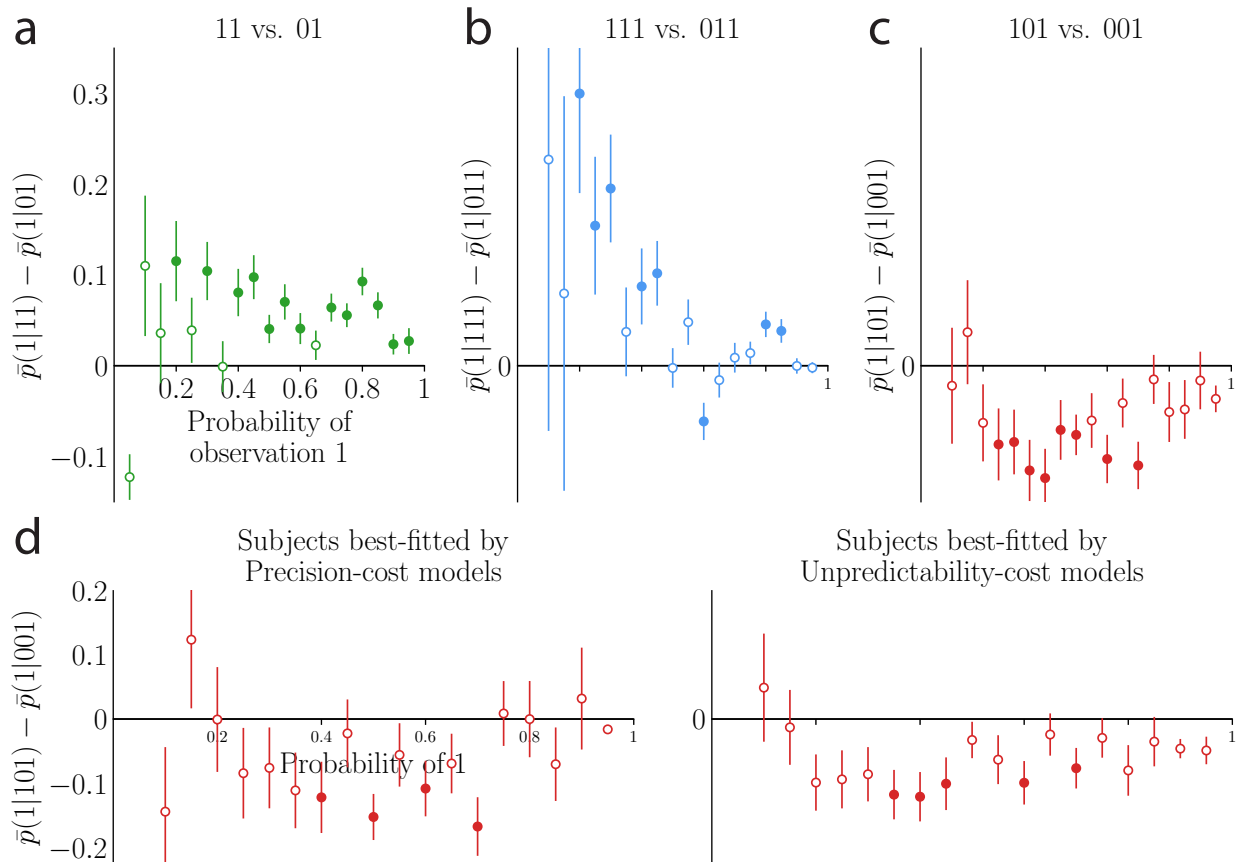


Figure 6: Patterns of attractive and repulsive sequential effects in subjects' responses. (a) Difference between the proportion of prediction 1 conditional on observing '11', and conditional on observing '01', i.e., $\bar{p}(1|11) - \bar{p}(1|01)$, as a function of the stimulus generative probability. With a majority of stimulus generative probabilities, this difference is positive, indicating an attractive sequential effect of the second-to-last stimulus. (b) Difference between the proportion of prediction 1 conditional on observing '111', and conditional on observing '011', i.e., $\bar{p}(1|111) - \bar{p}(1|011)$. With a majority of stimulus generative probabilities, this difference is positive, indicating an attractive sequential effect of the third-to-last stimulus in this case. (c) Difference between the proportion of prediction 1 conditional on observing '101', and conditional on observing '001', i.e., $\bar{p}(1|101) - \bar{p}(1|001)$. With a majority of stimulus generative probabilities, this difference is *negative*, indicating a repulsive sequential effect of the third-to-last stimulus in this case. (d) Same as (c), with subjects split in two groups: the subjects best-fitted by precision-cost models (left) and the subjects best-fitted by unpredictability-cost models (right). In all panels, filled points indicate that the p-value of the Fisher exact test is below 0.05.

history. Models with different types of cognitive cost resulted in different patterns of attractive and repulsive effects of the past few stimuli on predictions. To compare the predictions of constrained inference with human behavior, we asked subjects to predict each next outcome in sequences of binary stimuli. We manipulated the stimulus generative probability in blocks of trials, exploring exhaustively the probability range from .05 to .95 by increments of

.05. We found that subjects' predictions depend on both the stimulus generative probability and the recent stimulus history. Sequential effects exhibited both attractive and repulsive components which were modulated by the stimulus generative probability. This behavior was qualitatively accounted for by a model of constrained inference in which the subject infers the transition probabilities underlying the sequences of stimuli and bears a cost that increases with the precision of the posterior distributions. Our study proposes a novel theoretical account of sequential effects in terms of optimal inference under cognitive constraints and it uncovers the richness of human behavior over a wide range of stimulus generative probabilities.

The notion that human decisions can be understood as resulting from a constrained optimization has gained traction across several fields, including neuroscience, cognitive science, and economics. In neuroscience, a voluminous literature that started with Attneave [33] and Barlow [34] investigates the idea that perception maximizes the transmission of information, under the constraint of costly and limited neural resources [35–37]; related theories of 'efficient coding' account for the bias and the variability of perception [38–41]. In cognitive science and economics, 'bounded rationality' is a precursory concept introduced in the 1950s by Herbert Simon, who defines it as "rational choice that takes into account the cognitive limitations of the decision maker — limitations of both knowledge and computational capacity" [42]. For Gigerenzer, these limitations promote the use of heuristics, which are "fast and frugal" ways of reasoning, leading to biases and errors in humans and other animals [43, 44]. A range of more recent approaches can be understood as attempts to specify formally the limitations in question, and the resulting trade-off. The 'resource-rational analysis' paradigm aims at a unified theoretical account that reconciles principles of rationality with realistic constraints about the resources available to the brain when it is carrying out computations [45]. In this approach, biases result from the constraints on resources, rather than from 'simple heuristics' (see Ref. [46] for an extensive review). For instance, in economics, theories of 'rational inattention' propose that economic agents optimally allocate resources (a limited amount of attention) to make decisions, thereby proposing new accounts of empirical findings in the economic literature [47–52].

Our study puts forward a 'resource-rational' account of sequential effects. Traditional accounts since the 1960s attribute these effects to a belief in sequential dependencies between successive outcomes [14, 16] (potentially "acquired through life experience" [18]), and more generally to the incorrect models that people assume about the processes generating sequences of events (see Ref. [22] for a review). This traditional account was formalized, in particular, by models in which subjects carry out a statistical inference about the sequence of stimuli presented to them, and this inference assumes that the parameters underlying the generating process are subject to changes [8, 9, 11, 12]. In these model, sequential effects are thus understood as resulting from a rational adaptation to a changing world. Human subjects indeed dynamically adapt their learning rate when the environment changes [53–55], and they can even adapt their inference to the statistics of these changes [30, 56]. However, in our task and in many previous studies in which sequential effects have been reported, the underlying statistics are in fact not changing across trials. The models just mentioned thus leave unexplained why subjects' behavior, in these tasks, is not rationally adapted to the unchanging statistics of the stimulus.

What underpins our main hypothesis is a different kind of rational adaptation: one,

instead, to the “cognitive limitations of the decision maker”, which we assume hinder the inference carried out by the brain. We show that rational models of inference under a cost yield rich patterns of sequential effects. When the cost varies with the precision of the posterior (measured here by the negative of its entropy, Eq. (5)), the resulting optimal posterior is proportional to the product of the prior and the likelihood, each raised to an exponent $1/(\lambda+1)$ (Eq. (6)). Many previous studies on biased belief updating have proposed models that adopt the same form except for the different exponents applied to the prior and to the likelihood [57–59]. Here, with the precision cost, both quantities are raised to the same exponent and we note that in this case the inference of the subject amounts to an exponentially-decaying count of the patterns observed in the sequence of stimuli, which is sometimes called ‘leaky integration’ in the literature [8–10, 12]. The models mentioned above, that posit a belief in changing statistics, indeed are well approximated by models of leaky integration [8, 12], which shows that the exponential discount can have different origins. Reference [12] shows that the precision-cost, Markov-observer model with $m = 1$ (named ‘local transition probability model’ in this study) accounts for a range of other findings, in addition to sequential effects, such as biases in the perception of randomness and patterns in the surprise signals recorded through EEG and fMRI. Here we reinterpret these effects as resulting from an optimal inference subject to a cost, rather than from a suboptimal erroneous belief in the dynamics of the stimulus’ statistics. In our modeling approach, the minimization of a loss function (Eq. (1)) formalizes a trade-off between the distance to optimality of the inference, and the cognitive constraints under which it is carried out. We stress that our proposal is not that the brain actively solves this optimization problem online, but instead that it is endowed with an inference algorithm (whose origin remains to be elucidated) which is effectively a solution to the constrained optimization problem.

By grounding the sequential effects in the optimal solution to a problem of constrained optimization, our approach opens avenues for exploring the origins of sequential effects, in the form of hypotheses about the nature of the constraint that hinders the inference carried out by the brain. With the precision cost, more precise posterior distributions are assumed to take a larger cognitive toll. The intuitive assumption that it is costly to be precise finds a more concrete realization in neural models of inference with probabilistic population codes: in these models, the precision of the posterior is proportional to the average activity of the population of neurons and to the number of neurons [60, 61]. More neural activity and more neurons arguably come with a metabolic cost, and thus more precise posteriors are more costly in these models. Imprecisions in computations, moreover, was shown to successfully account for decision variability and adaptive behavior in volatile environments [62, 63]. The unpredictability cost, which we introduce, yields models that also exhibit sequential effects (for Markov observers), and that fit several subjects better than the precision-cost models. The unpredictability cost relies on a different hypothesis: that the brain prefers predictable environments over unpredictable ones. Human subjects exhibit a preference for predictive information indeed [64, 65], while unpredictable stimuli have been shown not only to increase anxiety-like behavior [66], but also to induce more neural activity [66–68] — a presumably costly increase, which may result from the encoding of larger prediction errors [66, 69]. We note that both costs (precision and unpredictability) can predict sequential effects, even though neither carries *ex ante* an explicit assumption that presupposes the existence of sequential effects. They both reproduce the attractive recency effect of the last stimulus

exhibited by the subjects. They make quantitatively different predictions (Fig. 4c); we also find this diversity of behaviors in subjects.

The precision cost, as mentioned above, yields leaky-integration models which can be summarized by a simple algorithm in which the observed patterns are counted with an exponential decay. The psychology and neuroscience literature proposes many similar ‘leaky integrators’ or ‘leaky accumulators’ models [8, 12, 70–80]. In connectionist models of decision-making, for instance, decision units in abstract network models have activity levels that accumulate evidence received from input units, and which decay to zero in the absence of input [71, 72, 74, 76, 79]. In other instances, perceptual evidence [77, 78, 80] or counts of events [8, 12, 75] are accumulated through an exponential temporal filter. Although it is less clear whether the unpredictability-cost models lend themselves to a similar algorithmic simplification, they consist in a distortion of Bayesian inference, for which various neural-network models have been proposed [81–84].

Turning to the experimental results, we note that in spite of the rich literature on sequential effects, the majority of studies have focused on *equiprobable* Bernoulli environments, in which the two possible stimuli both had a probability equal to 0.5, as in tosses of a fair coin [6–11, 18, 25]. In environments of this kind, the two stimuli play symmetric roles and all sequences of a given length are equally probable. In contrast, in biased environments one of the two possible stimuli is more probable than the other. Although much less studied, this situation breaks the regularities of equiprobable environments and is arguably very frequent in real life. In our experiment, we explore stimulus generative probabilities from .05 to .95, thus allowing to investigate the behavior of subjects in a wide spectrum of Bernoulli environments: from these with “extreme” probabilities (e.g., $p = .95$) to these only slightly different from the equiprobable case (e.g., $p = .55$) to the equiprobable case itself ($p = .5$). The subjects are sensitive to the imbalance of the non-equiprobable cases: while they predict 1 in half the trials of the equiprobable case, a probability of just .55 suffices to prompt the subjects to predict 1 in about in 60% of trials, a significant difference ($\bar{p}(1) = .602$; sem: .008; p-value of t-test of equality with .5: 1.7e-11).

The well-known ‘probability matching’ hypothesis [85–87] suggests that the proportion of predictions 1 matches the stimulus generative probability: $\bar{p}(1) = p$. This hypothesis is not supported by our data. We find that in the non-equiprobable conditions these two quantities are significantly different (all p-values $< 1e-11$, when $p \neq .5$). More precisely, we find that the proportion of prediction 1 is more extreme than the stimulus generative probability (i.e., $\bar{p}(1) > p$ when $p > .5$, and $\bar{p}(1) < p$ when $p < .5$; Fig. 2a). This result is consistent with the observations made by W. Edwards [14, 88] and with the conclusions of a more recent review [86].

In addition to varying with the stimulus generative probability, the subjects’ predictions depend on the recent history of stimuli. Recency effects are common in the psychology literature; they were reported from memory [89] to causal learning [90] to inference [59, 91, 92]. In prediction tasks, both attractive recency effects, also called ‘hot-hand fallacy’, and repulsive recency effects, also called ‘gambler’s fallacy’, have been reported [13, 14, 18–20, 22]. The observation of both effects within the same experiment has been reported in a visual identification task [93] and in risky choices (‘wavy recency effect’ [23, 24]). As to the heterogeneity of these results, several explanations have been proposed; two important factors seem to be the perceived degree of randomness of the predicted variable and whether

it relates to human performance [18–20, 22]. In any event, most studies focus exclusively on the influence of ‘runs’ of identical outcomes on the upcoming prediction, e.g., in our task, on whether three 1s in a row increases the proportion of predictions 1. With this analysis, Edwards [14] in a task similar to ours concluded to an attractive recency effect (which he called ‘probability following’). Although our results are consistent with this observation (in our data three 1s in a row do increase the proportion of predictions 1), we provide a more detailed picture of the influence of each stimulus preceding the prediction, whether it is in a ‘run’ of identical stimuli or not, which allows us to exhibit the non-trivial finer structure of the recency effects that is often overlooked.

Up to two stimuli in the past, the recency effect is attractive: observing a 1 at trial $t - 2$ or at trial $t - 1$ induces, all else being equal, a higher proportion of predictions 1 at trial t (in comparison to observing a 0; Figs. 2, 6a). The influence of the third-to-last stimulus is more intricate: it can yield either an attractive or a repulsive effect, depending on the second-to-last and the last stimuli. For a majority of probability parameters, p , while a 1 followed by the sequence ‘11’ has an attractive effect (i.e., $p(1|111) > p(1|011)$), a 1 followed by the sequence ‘01’ has a repulsive effect (i.e., $p(1|101) < p(1|001)$; Fig. 6b, c). How can this reversal be intuited? Only one of our models, the precision-cost model with a Markov order 1 ($m = 1$), reproduces this behavior; we show how it provides an interpretation for this result. From the update equation of this model (Eq. (6)), it is straightforward to show that the posterior of the model subject (a Dirichlet distribution of order 4) is determined by four quantities, which are exponentially-decaying counts of the four two-long patterns observed in the sequence of stimuli: ‘00’, ‘01’, ‘10’, and ‘11’. The higher the count of a pattern, the more likely the model subject deems this pattern to happen again. In the equiprobable case ($p = 0.5$), after observing the sequence ‘111’, the count of ‘11’ is higher than after observing ‘011’, thus the model subject believes that ‘11’ is more probable, and accordingly predicts 1 more frequently, i.e., $p(1|111) > p(1|011)$. As for the sequences ‘101’ and ‘001’, both result in the same count of ‘11’, but the former results in a higher count of ‘10’ — in other words, the short sequence ‘101’ suggests that a 1 is usually followed by a 0, but the sequence ‘001’ does not — and thus the model subject predicts more frequently a 0, i.e., less frequently a 1 ($p(1|101) < p(1|001)$).

In short, the ability of the precision-cost model of a Markov observer to capture the repulsive effect found in behavioral data suggests that human subjects extrapolate the local statistical properties of the presented sequence of stimuli in order to make predictions, and that they pay attention not only to the ‘base rate’ — the marginal probability of observing a 1, unconditional on the recent history — as a Bernoulli observer would do, but also to the statistics of more complex patterns, including the repetitions and the alternations, thus capturing the transition probabilities between consecutive observations. References [9, 10, 12] similarly argue that sequential effects result from an imperfect inference of the base rate and of the frequency of repetitions and alternations. Reference [94] argues that the knowledge of transition probabilities is a central mechanism in the brain’s processing of sequences (e.g., in language comprehension), and infants as young as five months were shown to be able to track both base rates and transition probabilities (see Ref. [95] for a review). Learning of transition probabilities has also been observed in rhesus monkeys [96].

Finally, the deviations from perfect inference, in the precision-cost model, originate in the constraints faced by the brain when performing computation with probability distributions.

In spite of the success of the Bayesian framework, we note that human performance in various inference tasks is often suboptimal [30, 55, 97, 98]. Our approach suggests that the deviations from optimality in these tasks may be explained by the cognitive constraints at play in the inference carried out by humans.

Methods

Task and subjects

The computer-based task was programmed using the Python library PsychoPy [99]. The experiment comprised ten blocks of trials, which differed by the stimulus generative probability, p , used in all the trials of each block. The probability p was chosen randomly among the ten values ranging from .50 to .95 by increments of .05, excluding the values already chosen; and with probability 1/2 the stimulus generative probability $1 - p$ was used instead. Each block started with 200 passive trials, in which the subject was only asked to look at the 200 stimuli sampled with the block's probability and successively presented. No action from the subject was required for these passive trials. The subject was then asked to predict, in each of 200 trials, the next location of the stimulus. Subjects provided their responses by keypress. The task was presented as a game to the subjects: the stimulus was a lightning symbol, and predicting correctly whether the lightning would strike the left or the right rod resulted in the electrical energy of the lightning being collected in a battery (Fig. 1). Twenty subjects participated in the experiment. All subjects completed the ten blocks of trials, except one subject who did not finish the experiment and was excluded from the analysis. Written consent was obtained from the participants before the experiment.

Sequential effects of the models

We run simulations of the eight models and look at the predictions they yield. To reproduce the conditions faced by the subjects, which included 200 passive trials, we start each stimulation by showing to the model subject 200 randomly sampled stimuli (without collecting predictions at this stage). We then show an additional 200 samples, and obtain a prediction from the model subject after each sample. The sequential effects of the most recent stimulus, with the different models, are shown in Fig. 7. With the precision-cost models, the posterior distribution of the model subject does not converge, but fluctuates instead with the recent history of the stimulus. This results in attractive sequential effects (Fig. 7a), including for the Bernoulli observer, who assumes that the probability of a 1 does not depend on the most recent stimulus. With the unpredictability-cost models, the posterior of the model subject does converge. With Markov observers, it converges toward a parameter vector q that implies that the probability of observing a 1 depends on the most recent stimulus, resulting in the presence of sequential effects of the most recent stimulus (Fig. 7b, second to fourth row). With a Bernoulli observer, the posterior of the model subject converges toward a value of the stimulus generative probability that does not depend on the stimulus history. As more evidence is accumulated, the posterior narrows around this value but not without some fluctuations that depend on the sequence of stimuli presented. In consequence the model

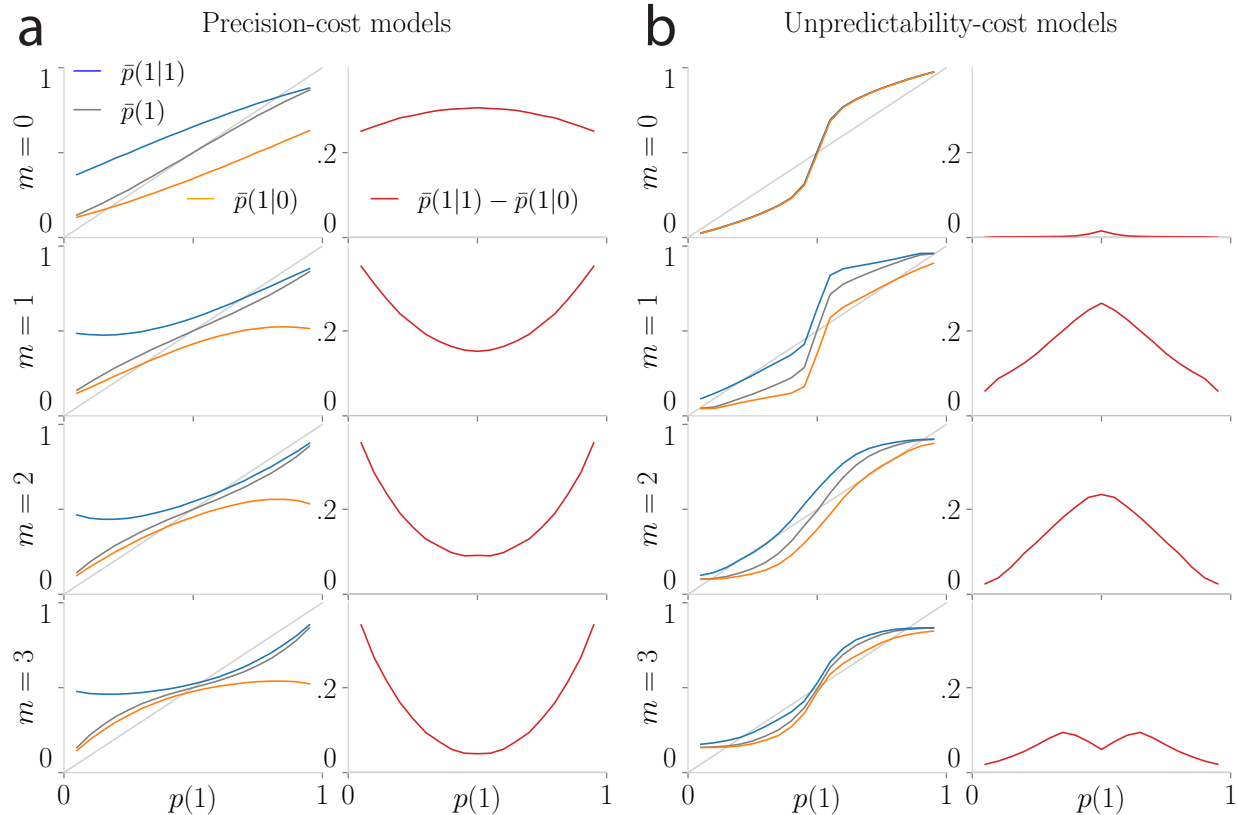


Figure 7: Sequential effects of the most recent stimulus in precision-cost and unpredictability-cost models. (a) Precision-cost models. (b) Unpredictability-cost models. First row: Bernoulli observers ($m = 0$). Second to fourth rows: Markov observers ($m = 1, 2$, and 3). First column (each panel): proportion of predictions 1 in the models' responses as a function of the stimulus generative probability, conditional on the preceding observation being a 1 (blue line) or a 0 (orange line), and unconditional (grey line). Second column (each panel): difference between the proportion of predictions 1 conditional on the preceding observation being a 1, and the same proportion conditional on a 0. A positive difference indicates an attractive sequential effect of the most recent stimulus.

subject's estimate of the stimulus generative probability is also subject to fluctuations, and depends on the history of stimuli (including the most recent stimulus), although the width of the fluctuations tend to zero as more stimuli are observed. After the 200 stimuli of the passive trials, the sequential effects of the most recent stimulus resulting from this transient regime appear small in comparison to the sequential effects obtained with the other models (Fig. 7b, first row).

Turning to higher-order sequential effects, we look at the influence on predictions of the second- and third-to-last stimulus (Fig. 8). As mentioned, only precision-cost models of Markov observers yield repulsive sequential effects, and these occur only when the third-to-last-stimulus is followed by '01' (or, symmetrically, by '10'). They do not occur with the second-to-last stimulus, nor with the third-to-last-stimulus when it is followed by '11' (or, symmetrically, by '00'; Fig. 8a); and they do not occur in any case with the unpredictability-

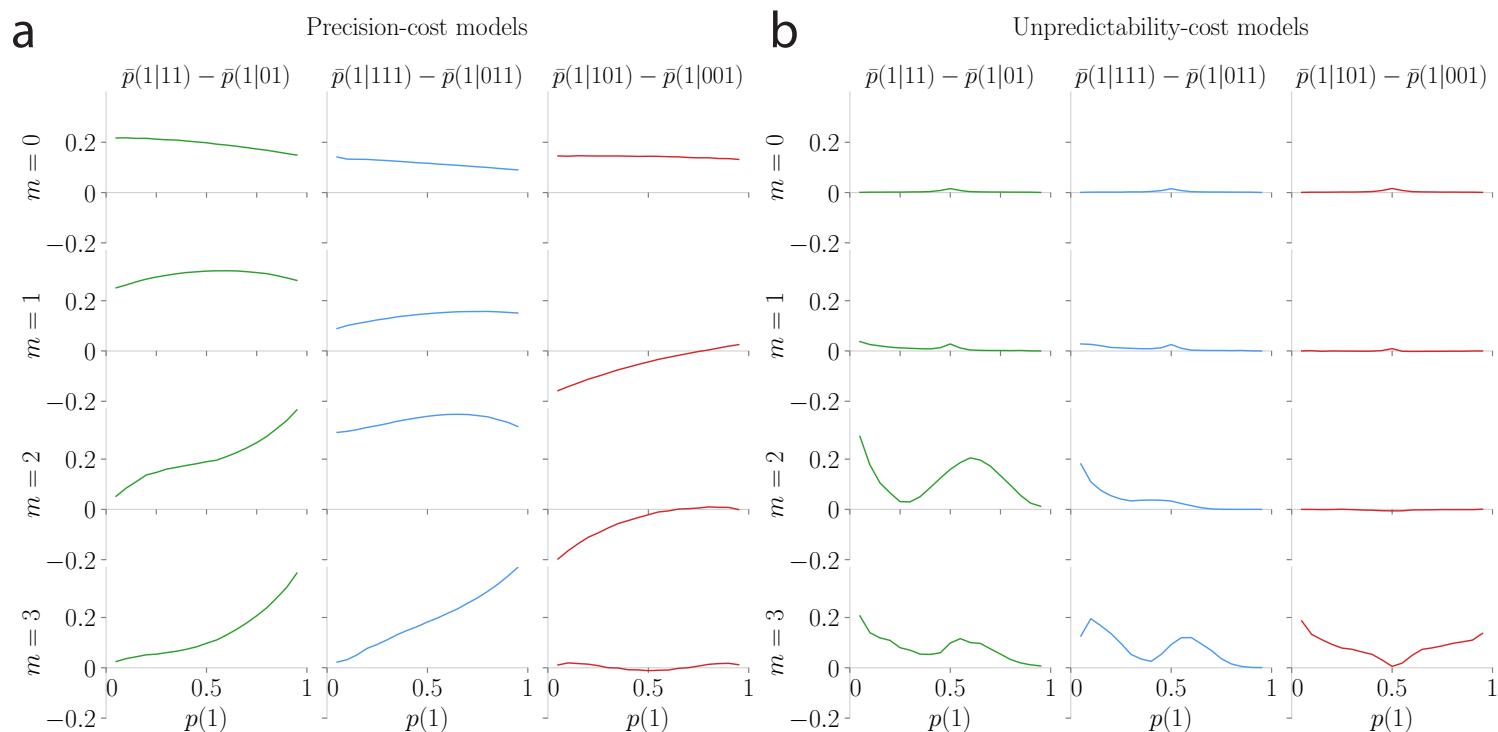


Figure 8: Sequential effects of the second- and third-to-last stimuli in precision-cost and unpredictability-cost models. (a) Precision-cost models. (b) Unpredictability-cost models. *First row:* Bernoulli observers ($m = 0$). *Second to fourth rows:* Markov observers ($m = 1, 2,$ and 3). *First column (each panel):* difference between the proportion of predictions 1 in the model subject’s responses, conditional on the two preceding observations being the sequence ‘11’, and the same proportion conditional on the sequence ‘01’. A positive difference indicates an attractive sequential effect of the second-to-last stimulus. *Second column (each panel):* difference between the proportion of predictions 1 in the model subject’s responses, conditional on the three preceding observations being the sequence ‘111’, and the same proportion conditional on the sequence ‘011’. *Third column (each panel):* difference between the proportion of predictions 1 in the model subject’s responses, conditional on the three preceding observations being the sequence ‘101’, and the same proportion conditional on the sequence ‘001’. The precision-cost models of Markov observers are the only models that yield a negative difference, i.e., a repulsive sequential effect of the third-to-last stimulus, in this case.

cost models (Fig. 8b).

Bayesian Model Selection

We implement the Bayesian model selection (BMS) procedure described in Ref. [32]. Given M models, this procedure aims at deriving a probabilistic belief on the distribution of these models among the general population. This unknown distribution is a categorical distribution, parameterized by the probabilities of the M models, denoted by $r = (r_1, \dots, r_M)$, with $\sum r_m = 1$. With a finite sample of data, one cannot determine with infinite preci-

sion the values of the probabilities r_m . The BMS, thus, computes an approximation of the Bayesian posterior over the vector r , as a Dirichlet distribution parameterized by the vector $\alpha = (\alpha_1, \dots, \alpha_M)$, i.e., the posterior distribution

$$p(r|\alpha) = \frac{1}{Z(\alpha)} \prod_{m=1}^M r_m^{\alpha_m - 1}. \quad (7)$$

Computing the parameters α_k of this posterior makes use of the log-evidence of each model for each subject, i.e., the logarithm of the joint probability, $p(y|m)$, of a given subject's responses, y , under the assumption that a given model, m , generated the responses. We use the model's maximum likelihood to obtain an approximation of the model's log-evidence, as [100]

$$\ln p(y|m) \simeq \max_{\theta} [\ln p(y|m, \theta)] - \frac{d}{2} \ln N, \quad (8)$$

where θ denotes the parameters of the model, $p(y|m, \theta)$ is the likelihood of the model when parameterized with θ , d is the dimension of θ , and N is the size of the data, i.e., the number of responses. (The well-known Bayesian Information Criterion [101] is equal to this approximation of the model's log-evidence, multiplied by $-1/2$.)

In our case, there are $M = 8$ models, each with $d = 2$ parameters: $\theta = (\lambda, \kappa)$. The posterior distribution over the parameters of the categorical distribution of models in the general population, $p(r|\alpha)$, allows for the derivation of several quantities of interest; following Ref. [32], we derive two types of quantities. First, given a family of models, i.e., a set $\mathcal{M} = \{m_i\}$ of different models (for instance, the prediction-cost models, or the Bernoulli-observer models), the expected probability of this class of model, i.e., the expected probability that the behavior of a subject randomly chosen in the general population follows a model belonging to this class, is the ratio

$$\frac{\sum_{m \in \mathcal{M}} \alpha_m}{\sum_{m=1}^K \alpha_m}. \quad (9)$$

We compute the expected probability of the precision-cost models (and the complementary probability of the unpredictability-cost models), and the expected probability of the Bernoulli-observer models (and the complementary probability of the Markov-observer models; see [Results](#)).

Second, we estimate, for each family of models \mathcal{M} , the probability that it is the most likely, i.e., the probability of the inequality

$$\sum_{m \in \mathcal{M}} r_m > 1/2, \quad (10)$$

which is called the 'exceedance probability'. We compute an estimate of this probability by sampling one million times the Dirichlet belief distribution (Eq. (7)), and counting the number of samples in which the inequality is verified. We estimate in this way the exceedance probability of the precision-cost models (and the complementary probability of the unpredictability-cost models), and the exceedance probability of the Bernoulli-observer models (and the complementary probability of the Markov-observer models; see [Results](#)).

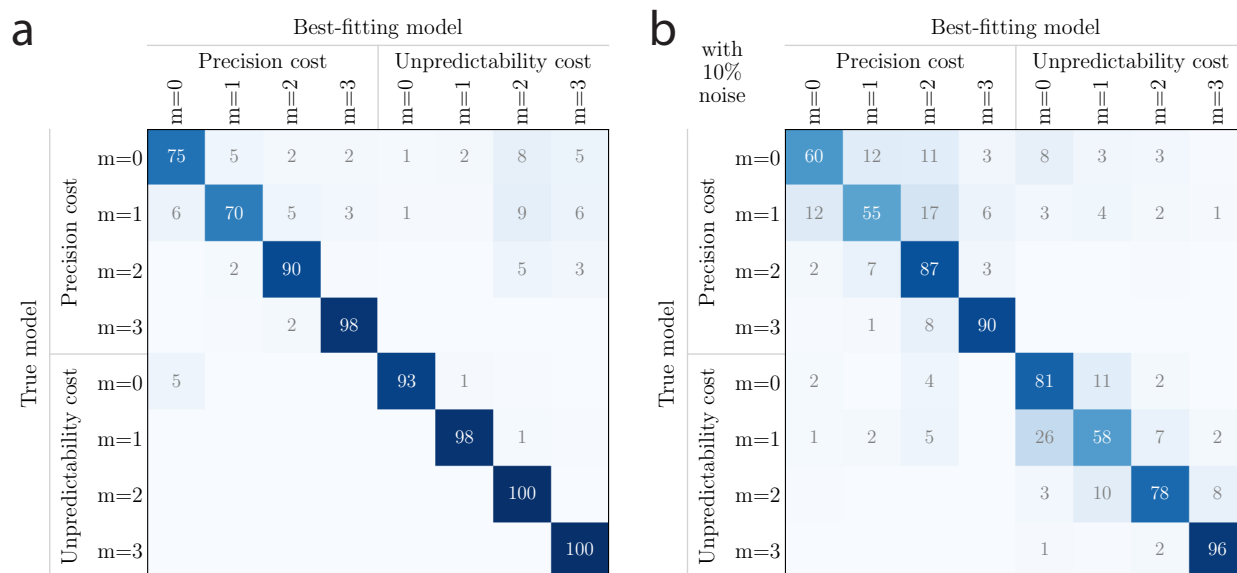


Figure 9: Model-fitting confusion matrix. (a) For each row models (‘true model’), percentage of simulated datasets of 200 responses that were best fitted by column models (‘best-fitting model’). Example: when fitting data generated by the precision-cost model with $m = 3$, the best-fitting model was the correct model on 98% of the fits, and the precision-cost model with $m = 2$ on 2% of the fits. (b) Same as a, with 10% of responses (randomly chosen in each simulated dataset) replaced by the opposite responses.

Robustness of the model fitting

To evaluate the ability of the model-fitting procedure to correctly identify the model that generated a given set of responses, we compute a confusion matrix of the eight models. For each model, we simulate 200 runs of the task (each with 200 passive trials followed by 200 trials in which a prediction is obtained), with values of λ and κ close to values typically obtained when fitting the subjects’ responses (for prediction-cost models, $\lambda \in \{0.03, 0.7, 2, 15\}$; for unpredictability-cost models, $\lambda \in \{0.7, 2\}$; and $\kappa \in \{0.7, 1.5, 2\}$ for both families of models). We then fit each of the eight models to each of these simulated datasets, and count how many times each model best fit each dataset (Fig. 9a). To further test the robustness of the model-fitting procedure, we randomly introduce errors in the simulated responses: for 10% of the responses, randomly chosen in each dataset, we substitute the response by its opposite (i.e., 0 for 1, and 1 for 0), and compute a confusion matrix using these new responses (Fig. 9b). In both cases, the model-fitting procedure identifies the correct model a majority of times (i.e., the best-fitting model is the model that generated the data; Fig. 9).

Finally, to examine the robustness of the weight of the cost, λ , we consider for each subject its best-fitting model in each family (the precision-cost family and the unpredictability-cost family), and we fit separately each model to the subject’s responses obtained in trials in which the stimulus generative probability was medium ($p \in \{.3, .35, .4, .45, .5, .55, .6, .65, .7\}$) and those in which it was extreme ($p \in \{.05, .1, .15, .2, .25, .75, .8, .85, .9, .95\}$). The Figure 10 shows the correlation between the best-fitting parameters obtained in these two cases.

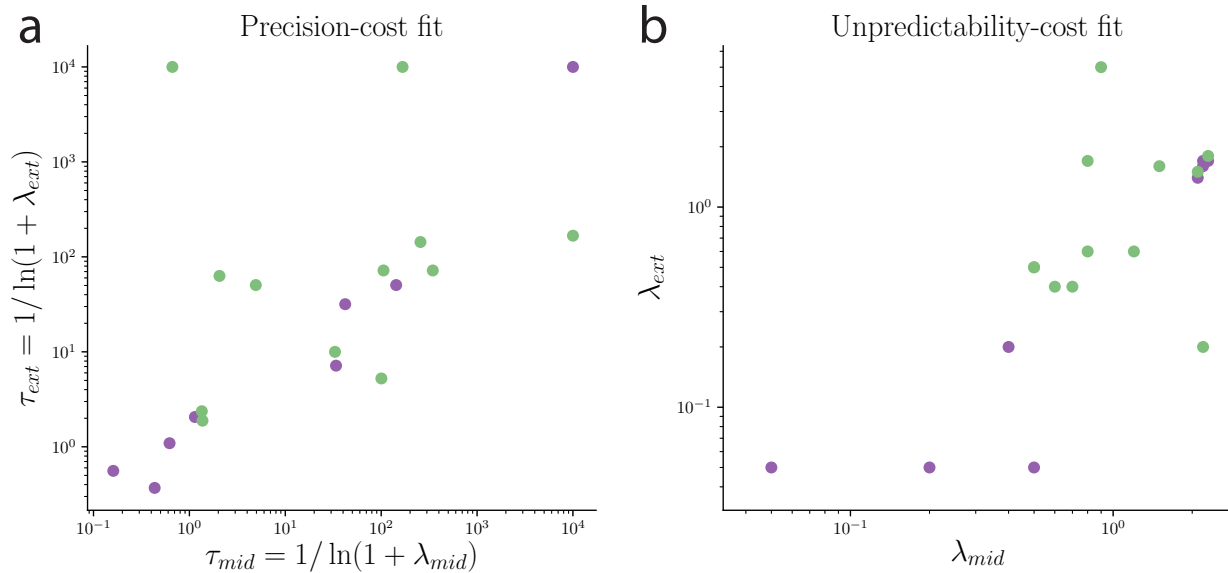


Figure 10: Stability of the cost-weight parameter across medium and extreme values of the stimulus generative probability. Log-log plot of the best-fitting parameters of the subjects when fitting the data obtained in trials with extreme values of the stimulus generative probability (i.e., p or $1 - p$ in $\{.75, .8, .85, .9, .95\}$), as a function of the best-fitting parameters of the subjects when fitting the data obtained in trials with medium values of the stimulus generative probability (i.e., p or $1 - p$ in $\{.5, .55, .6, .65, .7\}$), with (a) precision-cost models, and (b) unpredictability-cost models. *Purple dots*: subjects best-fitted by prediction-cost models. *Green dots*: subjects best-fitted by unpredictability-cost models. For the precision-cost models, we plot the characteristic decay time, $\tau = 1/\ln(1 + \lambda)$.

Unpredictability cost for Markov observers

Here we derive the expression of the unpredictability cost for Markov observers as a function of the elements of the parameter vector q . For an observer of Markov order 1 ($m = 1$), the vector q has two elements, which are the probability of observing a 1 at a given trial conditional on the preceding outcome being a 0, and the probability of observing a 1 at a given trial conditional on the preceding outcome being a 1, which we denote by q_0 and q_1 , respectively. The Shannon entropy, $H(X; q)$, implied by the vector q , is the average of the conditional entropies implied by each conditional probability, i.e.,

$$H(X; q) = p_0 H(X; q_0) + p_1 H(X; q_1), \quad (11)$$

where p_0 and p_1 are the *unconditional* probabilities of observing a 0 and a 1, respectively, and

$$H(X; q_i) = -q_i \ln q_i - (1 - q_i) \ln(1 - q_i), \quad (12)$$

where i is 0 or 1. The unconditional probabilities p_0 and p_1 are functions of the conditional probabilities q_0 and q_1 . We find:

$$p_0 = \frac{1 - q_1}{1 + q_0 - q_1} \text{ and } p_1 = \frac{q_0}{1 + q_0 - q_1}. \quad (13)$$

The entropy $H(X; q)$ implied by the vector q is obtained by substituting these quantities in Eq. (11).

Similarly, for $m = 2$ and 3 , the 2^m elements of the vector q are the parameters q_{ij} and q_{ijk} , respectively, where $i, j, k \in \{0, 1\}$, and where q_{ij} is the probability of observing a 1 at a given trial conditional on the two preceding outcomes being the sequence ‘ ij ’, and q_{ijk} is the probability of observing a 1 at a given trial conditional on the three preceding outcomes being the sequence ‘ ijk ’. The Shannon entropy, $H(X; q)$, implied by the vector q , is here also the average of the conditional entropies implied by each conditional probability, as

$$H(X; q) = \sum_{ij} p_{ij} H(X; q_{ij}), \text{ for } m = 2, \quad (14)$$

$$\text{and } H(X; q) = \sum_{ijk} p_{ijk} H(X; q_{ijk}), \text{ for } m = 3, \quad (15)$$

where p_{ij} and p_{ijk} are the unconditional probabilities of observing the sequence ‘ ij ’, and of observing the sequence ‘ ijk ’, respectively. These unconditional probabilities are functions of the conditional probabilities. We find, for $m = 2$:

$$\begin{aligned} p_{00} &= \left(1 + \frac{2q_{00}}{1 - q_{10}} + \frac{q_{01}}{1 - q_{11}} \frac{q_{00}}{1 - q_{10}} \right)^{-1}, \\ p_{01} &= \frac{q_{00}}{1 - q_{10}} p_{00}, \\ p_{10} &= p_{01}, \\ \text{and } p_{11} &= \frac{q_{01}}{1 - q_{11}} \frac{q_{00}}{1 - q_{10}} p_{00}. \end{aligned}$$

For $m = 3$, we find the relations:

$$\begin{aligned} p_{001} &= p_{000} \frac{q_{000}}{1 - q_{100}}, \\ p_{010} &= p_{000} \frac{q_{000}}{1 - q_{100}} \frac{1 - q_{001}(1 - q_{110}) - q_{101}q_{110}}{1 - q_{010}(1 - q_{101}) - q_{101}q_{110}}, \\ p_{011} &= p_{010} \frac{q_{101}q_{010}}{1 - q_{101}q_{110}} + p_{001} \frac{q_{001}}{1 - q_{101}q_{110}}, \\ p_{100} &= p_{001}, \\ p_{101} &= p_{010} + p_{011} - p_{001}, \\ p_{110} &= p_{011}, \\ \text{and } p_{111} &= p_{011} \frac{q_{011}}{1 - q_{111}}. \end{aligned}$$

Together with the normalization constraint $\sum_{ijk} p_{ijk} = 1$, these relations allow determining the eight unconditional probabilities p_{ijk} , and thus the expression of the Shannon entropy.

Acknowledgements

We thank Doron Cohen and Michael Woodford for inspiring discussions. This work was supported by the Alfred P. Sloan Foundation through grant G-2020-12680 and the CNRS through UMR8023. A.P.C. was supported by a Ph.D. fellowship of the Fondation Pierre-Gilles de Gennes pour la Recherche.

References

- [1] Christopher Summerfield and Floris P. De Lange. Expectation in perceptual decision making: Neural and computational mechanisms. *Nature Reviews Neuroscience*, 15(11):745–756, 2014.
- [2] W. E. Hick. On the Rate of Gain of Information. *Quarterly Journal of Experimental Psychology*, 4(1):11–26, 1952.
- [3] Ray Hyman. Stimulus information as a determinant of reaction time. *Journal of Experimental Psychology*, 1953.
- [4] Paul Bertelson. Serial choice reaction-time as a function of response versus signal-and-response repetition. *Nature*, 206(4980):217–218, 1965.
- [5] Sylvan Kornblum. Choice reaction time for repetitions and non-repetitions: A re-examination of the information hypothesis. *Acta Psychologica*, 27(0):178–187, 1967.
- [6] Eric Soetens, L. C. Boer, and J. E. Hueting. Expectancy or automatic facilitation? Separating sequential effects in two-choice reaction time. *Journal of Experimental Psychology: Human Perception and Performance*, 11(5):598–616, 1985.
- [7] Raymond Y Cho, Leigh E Nystrom, Eric T Brown, Andrew D Jones, Todd S Braver, and Jonathan D Cohen. Mechanisms underlying dependencies of performance on stimulus history in a two-alternative forced-choice task. *Cognitive, Affective, & Behavioral Neuroscience*, 2(4):283–299, 2002.
- [8] Angela J Yu and Jonathan D Cohen. Sequential effects: Superstition or rational behavior? *Advances in neural information processing systems*, 21:1873–1880, 2008.
- [9] Matthew H. Wilder, Matt Jones, and Michael C. Mozer. Sequential effects reflect parallel learning of multiple environmental regularities. *Advances in Neural Information Processing Systems 22 - Proceedings of the 2009 Conference*, pages 2053–2061, 2009.
- [10] Matt Jones, Tim Curran, Michael C Mozer, and Matthew H Wilder. Sequential effects in response time reveal learning mechanisms and event representations. *Psychological Review*, 120(3):628–666, 2013.
- [11] Shunan Zhang, Crane He Huang, and Angela J Yu. Sequential effects : A Bayesian analysis of prior bias on reaction time and behavioral choice. *Proceedings of the 36th Annual Conference of the Cognitive Science Society*, pages 1844–1849, 2014.

- [12] Florent Meyniel, Maxime Maheu, and Stanislas Dehaene. Human Inferences about Sequences: A Minimal Transition Probability Model. *PLoS Computational Biology*, 12(12):1–26, 2016.
- [13] Murray E. Jarvik. Probability learning and a negative recency effect in the serial anticipation of alternative symbols. *Journal of Experimental Psychology*, 41(4):291–297, 1951.
- [14] Ward Edwards. Probability learning in 1000 trials. *Journal of experimental psychology*, 62(4):385–394, 1961.
- [15] Gary H. McClelland and Beverly H. Hackenberg. Subjective probabilities for sex of next child: U.S. College students and Philippine villagers. *Journal of Population Behavioral, Social, and Environmental Issues*, 1(2):132–147, 1978.
- [16] Leroy Matthews and Warren Sanders. Effects of Causal and Noncausal Sequences of Information on Subjective Prediction. *Psychological Reports*, 54(1):211–215, 1984.
- [17] Thomas Gilovich, Robert Vallone, and Amos Tversky. The hot hand in basketball: On the misperception of random sequences. *Cognitive Psychology*, 17(3):295–314, 1985.
- [18] Peter Ayton and Ilan Fischer. The hot hand fallacy and the gambler’s fallacy: Two faces of subjective randomness? *Memory & Cognition*, 32(8):1369–1378, dec 2004.
- [19] Bruce D Burns and Bryan Corpus. Randomness and inductions from streaks: “Gambler’s fallacy” versus “hot hand”. *Psychonomic Bulletin & Review*, 11(1):179–184, feb 2004.
- [20] Rachel Croson and James Sundali. The gambler’s fallacy and the hot hand: Empirical data from casinos. *Journal of Risk and Uncertainty*, 30(3):195–209, 2005.
- [21] Michael Bar-Eli, Simcha Avugos, and Markus Raab. Twenty years of “hot hand” research: Review and critique. *Psychology of Sport and Exercise*, 7(6):525–553, nov 2006.
- [22] An T Oskarsson, Leaf Van Boven, Gary H McClelland, and Reid Hastie. What’s next? Judging sequences of binary events. *Psychological Bulletin*, 135(2):262–285, 2009.
- [23] Ori Plonsky, Kinneret Teodorescu, and Ido Erev. Reliance on small samples, the wavy recency effect, and similarity-based learning. *Psychological Review*, 122(4):621–647, 2015.
- [24] Ori Plonsky and Ido Erev. Learning in settings with partial feedback and the wavy recency effect of rare events. *Cognitive Psychology*, 93:18–43, 2017.
- [25] Dinis Gökaydin and Anastasia Ejova. Sequential effects in prediction. *Proceedings of the Annual Conference of the Cognitive Science Society*, pages 397–402, 2017.
- [26] Arthur Prat-Carrabin, Florent Meyniel, Misha Tsodyks, and Rava Azeredo da Silveira. Biases and variability from costly Bayesian inference. *Entropy*, 23(5):603, may 2021.

- [27] C E Shannon. A Mathematical Theory of Communication. *The Bell System Technical Journal*, pages 1–55, 1948.
- [28] Peter W. Battaglia, Daniel Kersten, and Paul R. Schrater. How haptic size sensations improve distance perception. *PLoS Computational Biology*, 7(6):e1002080, jun 2011.
- [29] Angela J Yu and He Huang. Maximizing Masquerading as Matching in Human Visual Search Choice Behavior. *Decision*, 1(4):275–287, 2014.
- [30] Arthur Prat-Carrabin, Robert C. Wilson, Jonathan D. Cohen, and Rava Azeredo da Silveira. Human inference in changing environments with temporal structure. *Psychological Review*, sep 2021.
- [31] Martin Nowak and Karl Sigmund. A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner’s Dilemma game. *Nature*, 364(6432):56–8, jul 1993.
- [32] Klaas Enno Stephan, Will D Penny, Jean Daunizeau, Rosalyn J Moran, and Karl J Friston. Bayesian model selection for group studies. *NeuroImage*, 46(4):1004–1017, jul 2009.
- [33] Fred Attneave. Some informational aspects of visual perception. *Psychological Review*, 61(3):183–193, 1954.
- [34] H. B. Barlow. Possible Principles Underlying the Transformations of Sensory Messages. In Walter A. Rosenblith, editor, *Sensory Communication*, chapter 13, pages 217–234. The MIT Press, Cambridge, MA, sep 1961.
- [35] Simon Laughlin. A simple coding procedure enhances a neuron’s information capacity. *Zeitschrift fur Naturforschung - Section C Journal of Biosciences*, 36(9-10):910–912, 1981.
- [36] S B Laughlin, R R de Ruyter Van Steveninck, and J C Anderson. The metabolic cost of neural information. *Nature neuroscience*, 1(1):36–41, 1998.
- [37] Eero P Simoncelli and Bruno A Olshausen. Natural image statistics and neural representation, mar 2001.
- [38] Deep Ganguli and Eero P. Simoncelli. Neural and perceptual signatures of efficient sensory coding. *ArXiv e-prints*, pages 1–24, feb 2016.
- [39] Xue Xin Wei and Alan A. Stocker. A Bayesian observer model constrained by efficient coding can explain ’anti-Bayesian’ percepts. *Nature Neuroscience*, 18(10):1509–1517, 2015.
- [40] Xue-Xin Wei and Alan A Stocker. Lawful relation between perceptual bias and discriminability. *Proceedings of the National Academy of Sciences*, 114(38):10244–10249, 2017.
- [41] Arthur Prat-Carrabin and Michael Woodford. Bias and variance of the Bayesian-mean decoder. *Advances in Neural Information Processing Systems*, 34, 2021.

- [42] Herbert Alexander Simon. Bounded Rationality. In *Models of Bounded Rationality : Empirically Grounded Economic Reason*, pages 291–294. The MIT Press, 1997.
- [43] Gerd Gigerenzer and Daniel G Goldstein. Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103(4):650–669, oct 1996.
- [44] Gigerenzer Gerd and Selten Reinhard. *Bounded Rationality : The Adaptive Toolbox.*, volume 1st MIT Press pbk. ed. The MIT Press, 2002.
- [45] Thomas L Griffiths, Falk Lieder, and D Goodman. Rational Use of Cognitive Resources : Levels of Analysis Between the Computational and the Algorithmic. *Topics in Cognitive Science*, 7:217–229, 2015.
- [46] Falk Lieder and Thomas L. Griffiths. Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, 2019.
- [47] Christopher A. Sims. Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690, 2003.
- [48] Michael Woodford. Information-Constrained State-Dependent Pricing. 2009.
- [49] Andrew Caplin, Mark Dean, and John Leahy. Rational Inattention, Optimal Consideration Sets, and Stochastic Choice. *The Review of Economic Studies*, 86(3):1061–1094, 2019.
- [50] Xavier Gabaix. Behavioral Inattention. Technical report, National Bureau of Economic Research, Cambridge, MA, dec 2017.
- [51] Rava Azeredo da Silveira and Michael Woodford. Noisy Memory and Over-Reaction to News. *AEA Papers and Proceedings*, 109:557–561, may 2019.
- [52] Rava Azeredo da Silveira, Yeji Sung, and Michael Woodford. Optimally Imprecise Memory and Biased Forecasts. Technical report, National Bureau of Economic Research, Cambridge, MA, nov 2020.
- [53] Elise Payzan-LeNestour, Simon Dunne, Peter Bossaerts, and John P O’Doherty. The neural representation of unexpected uncertainty during value-based decision making. *Neuron*, 79(1):191–201, jul 2013.
- [54] Florent Meyniel and Stanislas Dehaene. Brain networks for confidence weighting and hierarchical inference during probabilistic learning. *Proceedings of the National Academy of Sciences of the United States of America*, 114(19):E3859–E3868, 2017.
- [55] Matthew R Nassar, Robert C Wilson, B Heasley, and Joshua I Gold. An Approximately Bayesian Delta-Rule Model Explains the Dynamics of Belief Updating in a Changing Environment. *Journal of Neuroscience*, 30(37):12366–12378, 2010.

- [56] Timothy E J Behrens, Mark W Woolrich, Mark E Walton, and Matthew F S Rushworth. Learning the value of information in an uncertain world. *Nature Neuroscience*, 10(9):1214–1221, 2007.
- [57] David M. Grether. Bayes Rule as a Descriptive Model: The Representativeness Heuristic. *The Quarterly Journal of Economics*, 95(3):537, 1980.
- [58] Kaosu Matsumori, Yasuharu Koike, and Kenji Matsumoto. A Biased Bayesian Inference for Decision-Making and Cognitive Control. *Frontiers in Neuroscience*, 12(October):1–16, oct 2018.
- [59] Daniel J. Benjamin. Errors in probabilistic reasoning and judgment biases. In *Handbook of Behavioral Economics*, volume 2, pages 69–186. Elsevier B.V., 2019.
- [60] Wei Ji Ma, Jeffrey M Beck, Peter E Latham, and Alexandre Pouget. Bayesian inference with probabilistic population codes. *Nature Neuroscience*, 9(11):1432–1438, 2006.
- [61] H. S. Seung and H. Sompolinsky. Simple models for reading neuronal population codes. *Proceedings of the National Academy of Sciences of the United States of America*, 90(22):10749–10753, 1993.
- [62] Charles Findling, Vasilisa Skvortsova, Rémi Dromnelle, Stefano Palminteri, and Valentin Wyart. Computational noise in reward-guided learning drives behavioral variability in volatile environments. *Nature Neuroscience*, 22(December), 2019.
- [63] Charles Findling, Nicolas Chopin, and Etienne Koechlin. Imprecise neural computations as a source of adaptive behaviour in volatile environments. *Nature Human Behaviour*, 5(1):99–112, 2021.
- [64] Hirokazu Ogawa and Katsumi Watanabe. Implicit learning increases preference for predictive visual display. *Attention, Perception, and Psychophysics*, 73(6):1815–1822, 2011.
- [65] Sabrina Trapp, Amitai Shenhav, Sebastian Bitzer, and Moshe Bar. Human preferences are biased towards associative information. *Cognition and Emotion*, 29(6):1054–1068, 2015.
- [66] Cyril Herry, Dominik R. Bach, Fabrizio Esposito, Francesco Di Salle, Walter J. Perrig, Klaus Scheffler, A. Luthi, and Erich Seifritz. Processing of Temporal Unpredictability in Human and Animal Amygdala. *Journal of Neuroscience*, 27(22):5958–5966, may 2007.
- [67] Hanneke E.M. Den Ouden, Karl J. Friston, Nathaniel D. Daw, Anthony R. McIntosh, and Klaas E. Stephan. A dual role for prediction error in associative learning. *Cerebral Cortex*, 19(5):1175–1185, 2009.
- [68] Arjen Alink, Caspar M. Schwiedrzik, Axel Kohler, Wolf Singer, and Lars Muckli. Stimulus predictability reduces responses in primary visual cortex. *Journal of Neuroscience*, 30(8):2960–2966, 2010.

- [69] Wolfram Schultz and Anthony Dickinson. Neuronal Coding of Prediction Errors. *Annual Review of Neuroscience*, 23(1):473–500, mar 2000.
- [70] Philip L. Smith. Psychophysically principled models of visual simple reaction time. *Psychological Review*, 102(3):567–593, 1995.
- [71] R M Roe, J R Busemeyer, and J T Townsend. Multialternative decision field theory: a dynamic connectionist model of decision making. *Psychological review*, 108(2):370–392, 2001.
- [72] M Usher and J L McClelland. The time course of perceptual choice: the leaky, competing accumulator model., 2001.
- [73] Erik P. Cook and John H.R. Maunsell. Dynamics of neuronal responses in macaque MT and VIP during motion detection. *Nature Neuroscience*, 5(10):985–994, 2002.
- [74] Xiao Jing Wang. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron*, 36(5):955–968, 2002.
- [75] Leo P Sugrue, Greg S Corrado, and William T. Newsome. Matching Behavior and the representation of value in the parietal cortex. *Science*, 304(October):457–461, 2004.
- [76] Rafal Bogacz, Eric Brown, Jeff Moehlis, Philip Holmes, and Jonathan D Cohen. The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, 113(4):700–765, 2006.
- [77] Roozbeh Kiani, Timothy D Hanks, and Michael N Shadlen. Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *J Neurosci*, 28(12):3017–3029, 2008.
- [78] Juan Gao, Rebecca Tortell, and James L. McClelland. Dynamic integration of reward and stimulus information in perceptual decision-making. *PLoS ONE*, 6(3), 2011.
- [79] Konstantinos Tsetsos, Juan Gao, James L. McClelland, and Marius Usher. Using time-varying evidence to test models of decision dynamics: Bounded diffusion vs. The leaky competing accumulator model. *Frontiers in Neuroscience*, 6(JUN):1–17, 2012.
- [80] Ori Ossmy, Rani Moran, Thomas Pfeffer, Konstantinos Tsetsos, Marius Usher, and Tobias H. Donner. The timescale of perceptual evidence integration can be adapted to the environment. *Current Biology*, 23(11):981–986, 2013.
- [81] S Deneve, P E Latham, and A Pouget. Efficient computation and cue integration with noisy population codes. *Nature Neuroscience*, 4(8):826–831, aug 2001.
- [82] Wei Ji Ma, Jeffrey M. Beck, and Alexandre Pouget. Spiking networks for Bayesian inference and choice. *Current Opinion in Neurobiology*, 18(2):217–222, 2008.
- [83] Deep Ganguli and Eero P. Simoncelli. Efficient sensory encoding and Bayesian inference with heterogeneous neural populations. *Neural Computation*, 26(10):2103–2134, 2014.

- [84] Rodrigo Echeveste, Laurence Aitchison, Guillaume Hennequin, and Máté Lengyel. Cortical-like dynamics in recurrent circuits optimized for sampling-based probabilistic inference. *Nature Neuroscience*, 23(9):1138–1149, 2020.
- [85] R J Herrnstein. Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior*, 4(3):267–272, jul 1961.
- [86] Nir Vulkan. An economist’s perspective on probability matching. *Journal of Economic Surveys*, (December):1–22, 2000.
- [87] Wolfgang Gaissmaier and Lael J. Schooler. The smart potential behind probability matching. *Cognition*, 109(3):416–422, 2008.
- [88] Ward Edwards. Reward probability, amount, and information as determiners of sequential two-alternative decisions. *Journal of Experimental Psychology*, 52(3):177–188, 1956.
- [89] Hermann Ebbinghaus. *Memory: A contribution to experimental psychology*. Teachers College Press, New York, 1913.
- [90] Darrell J. Collins and David R. Shanks. Momentary and integrative response strategies in causal judgment. *Memory and Cognition*, 30(7):1138–1147, 2002.
- [91] James Shanteau. Descriptive versus normative models of sequential inference judgment. *Journal of Experimental Psychology*, 93(1):63–68, 1972.
- [92] Robin M. Hogarth and Hillel J. Einhorn. Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology*, 24(1):1–55, 1992.
- [93] Adrien Chopin and Pascal Mamassian. Predictive properties of visual adaptation. *Current Biology*, 22(7):622–626, 2012.
- [94] Stanislas Dehaene, Florent Meyniel, Catherine Wacogne, Liping Wang, and Christophe Pallier. The Neural Representation of Sequences: From Transition Probabilities to Algebraic Patterns and Linguistic Trees. *Neuron*, 88(1):2–19, 2015.
- [95] Jenny R. Saffran and Natasha Z. Kirkham. Infant Statistical Learning. *Annual Review of Psychology*, 69:181–203, 2018.
- [96] Travis Meyer and Carl R. Olson. Statistical learning of visual transitions in monkey inferotemporal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 108(48):19401–19406, 2011.
- [97] Yingyao Hu, Yutaka Kayaba, and Matthew Shum. Nonparametric learning rules from bandit experiments: The eyes have it! *Games and Economic Behavior*, 81(1):215–231, 2013.
- [98] Luigi Acerbi, Sethu Vijayakumar, and Daniel M. Wolpert. On the Origins of Suboptimality in Human Probabilistic Inference. *PLoS Computational Biology*, 10(6), 2014.

- [99] Jonathan W Peirce. Generating stimuli for neuroscience using PsychoPy. *Frontiers in Neuroinformatics*, 2(January):1–8, 2009.
- [100] Vijay Balasubramanian. Statistical Inference, Occam’s Razor, and Statistical Mechanics. *Neural Computation*, 368:349–368, 1997.
- [101] Gideon Schwarz. Estimating the Dimension of a Model. *The Annals of Statistics*, 6(2):461–464, mar 1978.