

Accurate step-hold tracking of smoothly varying periodic and aperiodic probability

Matthew Ricci¹ · Randy Gallistel²

© The Psychonomic Society, Inc. 2017

Abstract Subjects observing many samples from a Bernoulli distribution are able to perceive an estimate of the generating parameter. A question of fundamental importance is how the current percept—what we think the probability now is—depends on the sequence of observed samples. Answers to this question are strongly constrained by the manner in which the current percept changes in response to changes in the hidden parameter. Subjects do not update their percept trial-by-trial when the hidden probability undergoes unpredictable and un signaled step changes; instead, they update it only intermittently in a step-hold pattern. It could be that the step-hold pattern is not essential to the perception of probability and is only an artifact of step changes in the hidden parameter. However, we now report that the step-hold pattern obtains even when the parameter varies slowly and smoothly. It obtains even when the smooth variation is periodic (sinusoidal) and perceived as such. We elaborate on a previously published theory that accounts for: (i) the quantitative properties of the step-hold update pattern; (ii) subjects' quick and accurate reporting of changes; (iii) subjects' second thoughts about previously reported changes; (iv) subjects' detection of higher-order structure in patterns

of change. We also call attention to the challenges these results pose for trial-by-trial updating theories.

Keywords Bayesian modeling · Decision-making · Memory

Introduction

Perception can be generally described as an estimation problem involving non-stationary stochastic processes. Incoming sense data are random variables drawn from some distribution whose parameters change over time. Non-stationary stochastic processes have both quantitative and structural properties: the data and parameters that generate them are numerical quantities, but changes in parameters across time may be described by a formal model. For example, the intensity of sunlight striking an outdoor observer's eyes is a random variable due to cloud cover; yet the model generating these data has strong higher-order structure, namely, circadian periodicity. Studies of perception should take into account both of these elements.

With this in mind, Gallistel et al. (2014) studied the human perception of a stepwise non-stationary Bernoulli process. In their experiment, which roughly replicated a similar experiment by Robinson (1964), subjects used a computer interface to make thousands of individual draws of red or green circles from a box. Subjects were asked to estimate, draw-by-draw, the hidden parameter p_g of the Bernoulli process, that is, the proportion of green circles in the box. The parameter p_g would silently change on random trials. Subjects were additionally required to signal when they thought these silent changes occurred.

✉ Matthew Ricci
matthew_ricci_1@brown.edu
Randy Gallistel
galliste@rucss.rutgers.edu

¹ Brown University, Providence, RI, USA

² Rutgers University, New Brunswick, NJ, USA

Despite many differences in method and parameters, the experiments of Robinson (1964) and Gallistel et al. (2014) gave similar results: subjects tracked the hidden probability accurately and precisely over the full range of probabilities, and they responded quickly and abruptly to the hidden changes. Moreover, they consciously detected and reported these changes. Subjects sometimes had second thoughts about a change report; after seeing more data, they decided that their most recent report was erroneous, that there had not in fact been a change. This suggests that subjects keep a record of the observed sequence and recode earlier portions of the sequence in the retrospective light thrown by subsequent data.

A particularly surprising result was that subjects did not update their estimates (move the lever or the slider) observation-by-observation. They not uncommonly adjust their estimate by a small amount after a long interval (sometimes more than 100 observations). We call this the "step-hold" pattern in the perception of a probability. The step-hold pattern is theoretically important, because most computational models for the perception of probability assume trial-by-trial delta-rule updating of the percept (Glimcher, 2003; Sugrue et al. 2004, 2005; Behrens et al. 2007; Brown and Steyvers, 2009; Krugel et al. 2009; Wilson et al. 2013). Because the observed outcomes of a Bernoulli process are usually far from the current estimate of the parameter p_g (the percept), trial-by-trial delta-rule updating jerks the estimate around, unless it is also averaged over many trials. However, an average over many trials cannot change abruptly, and large, maximally abrupt adjustments in response to changes in p_g were observed in both experiments. The obvious explanation—reluctance to overtly adjust the lever or slider when the change required by the most recent trial or two is small—is ruled out by the form of the distribution of step heights. The smallest steps, which would be eliminated from the distribution by the hypothesized reluctance, were in fact the most frequent.

Gallistel et al. (2014) explained subjects' step-hold behavior with a Bayesian model that constructs a representation of the history of the Bernoulli parameter p_g in terms of its estimated change-points. For example, suppose that between trials 1 and 41, the model estimates $p_g = .25$, after which it detects the parameter has changed to $p_g = .9$. The representation of the p_g parameter history would then be the sequence of ordered pairs $\{(0, .25), (41, .9)\}$. The current percept is the second element of the most recent entry in the sequence.

The model detects change-points by computing the Kullback–Leibler divergence of its current estimate from the sequence observed since the most recent change point in the parameter history. If and when the probability that the current estimate is valid falls below a threshold, the

model re-estimates p_g . In doing so, it decides which of three possibilities is the most likely explanation for its failure to predict the most recently observed relative frequency of green circles:

1. The current estimate is inaccurate due to the inescapable small sample errors that arise from making a new estimate as soon as a change is detected. In this case, it keeps its estimate of the most recent putative change point but re-estimates the current p in the light of the additional data seen since the initial estimate was made.
2. The current estimate is inaccurate because, in the light of subsequent data, the most recent change point was not in fact a change point. In this case, p_g is re-estimated using the data extending back to the penultimate putative change point, and the most recent putative change point is dropped from the representation of the parameter history.
3. The current estimate is wrong because there has been a new change. In this case, it estimates the locus of that change, adds that change point to its representation of the parameter history, and estimates the new p_g , using only the data after the estimated new change point.

Because the computational model adjusts its estimates of p_g only when it has evidence that the current estimate is invalid—the authors call this the "if it ain't broke (IIAB), don't fix it principle"—it changes its estimate only intermittently, as do human subjects. Henceforth, we call this model IIAB. For an extensive comparison between IIAB's and delta-rule models' ability to capture human behavior, see Gallistel et al. (2014).

IIAB accounted well for subjects' estimation of a stepwise non-stationary process, but it remained unclear how it would generalize to other types of non-stationary stochastic processes, like those whose parameters change continuously or have deterministic structure. The subjects in Gallistel et al. (2014) may have been induced to display step-hold behavior because the true parameter was generated by a step function. In this case, the model would reflect only an experimentally induced strategy rather than a basic property of the probability perception mechanism. Further, the stepwise process used in Gallistel et al. (2014) changed completely at random, so the authors could not ask whether subjects were able to deduce deterministic structure in the process purely from data. They therefore could not confirm the report of Estes (1984), who claimed subjects estimating a sinusoidally changing Bernoulli parameter could explicitly detect periodicity, contrary to the predictions of his delta-rule updating model.

The purpose of the current experiment is to go beyond the comparison to delta-rule models presented in Gallistel et al. (2014) and instead to extend IIAB to new types of

data and emphasize the utility of an explicit change-point memory in the detection of structure. Our subjects estimated the generating parameter of a Bernoulli distribution that changed continuously in one of two ways: p_g either changed smoothly between stationary sections, or varied sinusoidally. We find that the step-hold pattern is seen in every subject even when the hidden probability changes continuously, that is, even when the characteristics of the stochastic process to which subjects are exposed discourages such a strategy. Further, we found that subjects in the periodic condition demonstrated improved performance on a structure-dependent measure compared to those in the aperiodic condition, supporting Estes' conclusion that subjects can detect periodic structure. Finally, we describe the IIAB model in more detail and discuss some advantages models encoding hierarchical structure have over delta-rule models in perception, learning and memory.

Methods

Nine subjects participated in the experiment. Following standard psychophysical assumptions, we consider each subject as a replication. In this case, we have nine replications of all the essential findings. Because we are primarily concerned with effects per trial, rather than per subject, there is large experimental power in the 10,000 trials we ran on each of the nine subjects. We note below wherever between-subject differences occurred and how they can be better captured by IIAB than by delta-rule models.

On a computer monitor, the subjects viewed the user interface shown in Fig. 1. They used a mouse to draw a new sample from the hidden distribution, the "Box of RINGS", by clicking on the "Next" button. Each click of the "Next" button prompted the appearance of a green or red ring to the right of the "Box of RINGS". Subjects were told that the hidden distribution contained some proportion of green

and red rings and that this proportion silently changes. They were not told whether the change would be sudden, gradual, periodic, etc. At their discretion, subjects updated their current estimate of the hidden proportion of green rings, p_g , by adjusting a slider. We made it clear to our subjects that their goal was to estimate the hidden proportion p_g and not the observed proportion, which is the total number of drawn green rings divided by the number of draws. Subjects were told to set the slider to some initial estimate before any rings were observed. The mean initial slider setting was .47, suggesting subjects had an unbiased prior as to the initial proportion of rings. Note that, as in the previous version of this experiment, subjects drew rings at their leisure and updated the slider setting whenever they felt the need.

On the right of the user interface was a box containing 1000 green and red rings accurately representing the subject's current estimate of p_g . Though this was intended as a visual guide to the subjects, most said they ignored it. Unlike in the version reported by Gallistel et al. (2014), subjects were not told to explicitly record their detection of change-points by clicking on boxes marked "I think the box has changed" or "I take that back!". As there were no discrete change points, these requests would not have made sense.

After practicing with the user interface, subjects completed ten sessions of 1,000 trials (draws) each. At the end of each session, subjects were allowed to take a break. Subjects were paid a baseline of \$10 per session and given a bonus corresponding to their accuracy. In Gallistel et al. (2014), there was no performance bonus; in Robinson (1964), subjects were penalized according to their error.

The hidden parameter p_g varied smoothly and periodically for four subjects and smoothly and aperiodically for five. In the first case, p_g was a sine function of trial number, oscillating between 0 and 1 with a period of 200 trials. This oscillation continued for all sessions until the last, at which point the parameter was fixed at .5. In the smooth, aperiodic

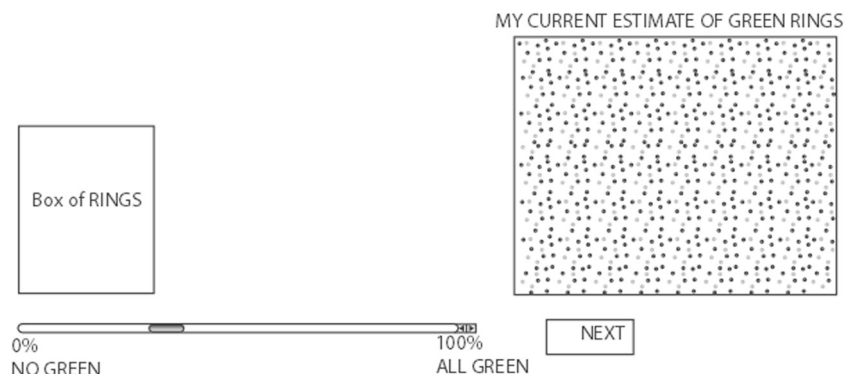


Fig. 1 Grayscale cartoon of the computer screen. Subjects clicked on the NEXT button to draw another red or green circle from the Box of Rings. They used the slider to indicate their current perception of p_g , the fraction of green circles in the box. The large box at upper right showed the proportion of red and green indicated by the slider's position

case, the hidden parameter was generated in two steps. First, p_g was modeled as a step function like that controlling the hidden parameter in Gallistel et al. (2014). The probability of a step change after any trial was .005, so the changes were geometrically distributed with an expected interval between change-points of 200 trials. This aperiodic step function was then smoothed by three Gaussian kernels with different variances. The result was a hidden p_g that was constant on long intervals but then gradually changed in a smooth way (see solid lines in Fig. 3). In both conditions, the value of the hidden parameter changed only by very small amounts between any two trials.

Two of the subjects in the periodic condition mistakenly exited the experiment computer program, effectively deleting a total of four sessions, about 2% of all trials, from our data. We consider this an unsubstantial decrease in total experimental power.

Results

We include two types of results for this experiment. First, we report the “quantitative performance” of the subjects; namely, how accurate are they across trials and what are the distributions of slider movements? We call these “quantitative” as they do not explicitly measure subjects’ ability

to detect non-local properties of the model generating observed data. Next, we describe the “structural performance” of subjects; namely, how quickly do they detect changes in the hidden parameter, and can they detect the current derivative of the generating model or its periodicity?

Because of the large number of samples from each subject (10,000), reported effects are trivially significant ($p < 10^{-6}$). Hence, we only explicitly state effect sizes (Cohen’s d) below.

Quantitative performance

Step-hold updating. Examples of subject slider movements in an early session, together with the samples actually observed by subjects, are displayed in Fig. 2. All nine subjects displayed the step-hold updating pattern originally observed in Robinson (1964) and replicated in Gallistel et al. (2014). They adjusted the slider at irregular intervals, often keeping their estimate constant across many trials (Figs. 3, 4). This confirms (Robinson, 1964)’s finding that he had observed this pattern even in pilot experiments with a continuously varying Bernoulli parameter.

The joint distribution of step widths and step heights for the data pooled across subjects is shown in Fig. 5a, with contrasting distributions from two individual subjects in Fig. 5b and c. One subject (Fig. 5b) produced a bimodal distribution

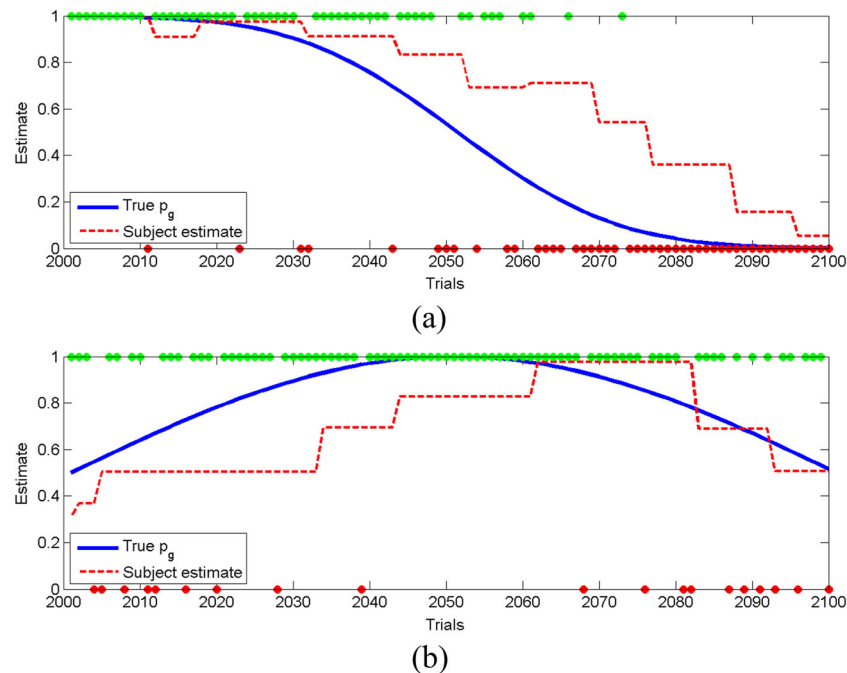


Fig. 2 Third-session slider movements plotted against observed samples. The solid blue line is the true p_g ; the dotted red line is subject estimate; the red and green dots at top and bottom of graph are the actual samples viewed by subjects. Data for the first 100 trials from the session are shown. **a)** After a string of ten green samples, the unexpected red sample on trial 2011 may have caused subject ‘BC’ to adjust

the slider downward, only to cancel the adjustment after a subsequent string of greens. **b)** Notice that, between trials 2000 and 2080, subject ‘JM’ never moves the slider downward despite 12 red samples. Only when two red samples occur back-to-back on trials 2081-2 after a long string of mostly green does the subject begin to adjust the slider, correctly, downward

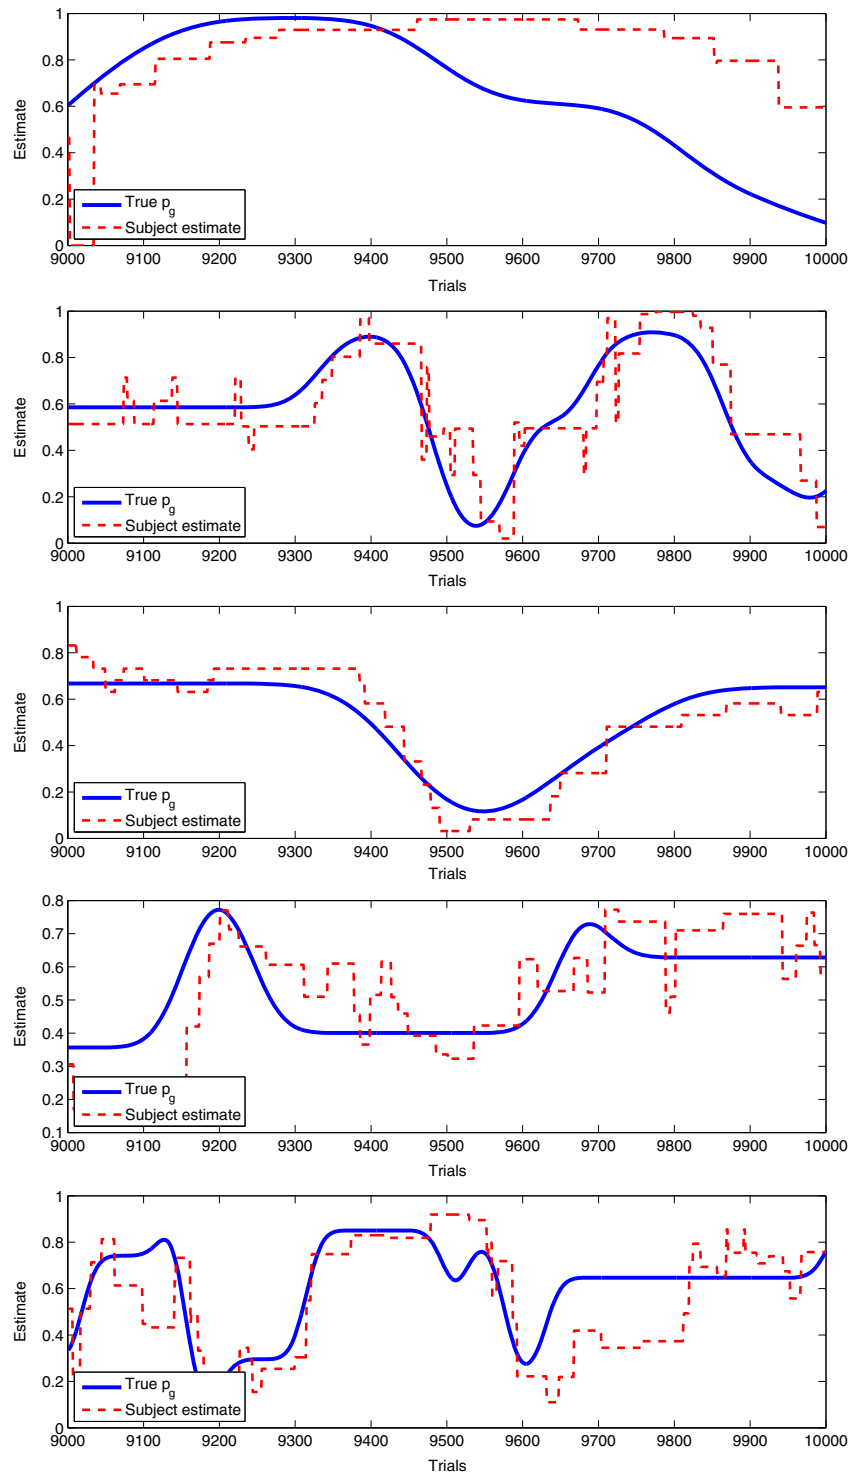


Fig. 3 Aperiodic slider settings. Trial-by-trial slider settings (red dotted lines) and hidden p_g values (blue solid lines) for the last session for the five subjects in the aperiodic condition. Subjects typically moved the slider by small amounts after long intervals (the step-hold behavior)

of step heights, but his data reveals that small step movements were in no sense completely eliminated. The maximal hold time across all subjects was 711 trials, nearly one whole session. Subjects displayed step-hold behavior,

despite the underlying, continuously changing parameter. Further, there was only a slight but significant increase in mean hold times during stationary sections (mean 30.45 during stationary sections; mean 27.00 for non-stationary

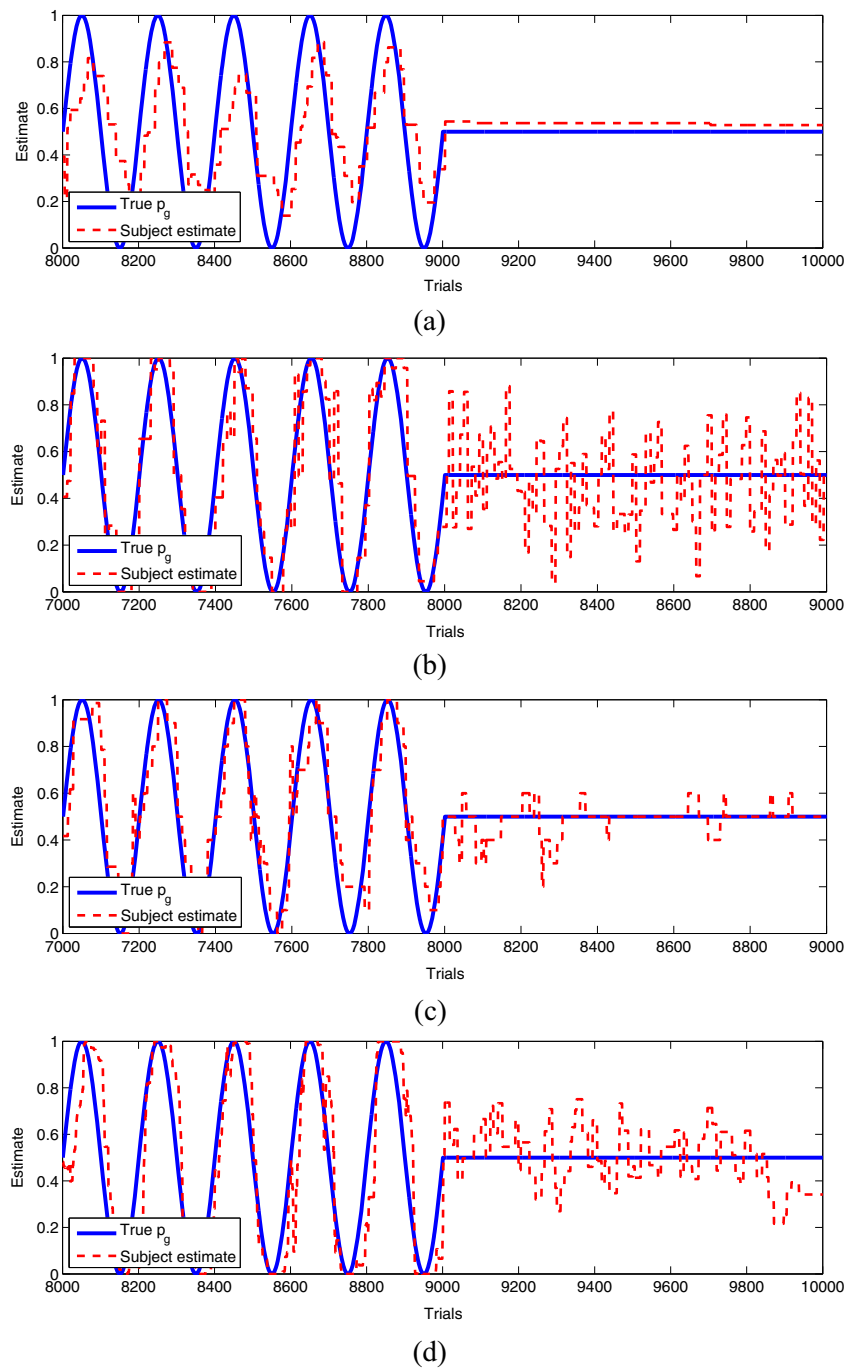


Fig. 4 Periodic slider settings. Trial-by-trial slider settings (*red dotted lines*) and hidden p_g values (*blue solid lines*) over the last two sessions for the four subjects in the periodic condition. The hidden p_g went flat at 0.5 at the beginning of last session (Trial 9000 in this plot). Note that, in the second and fourth panels, subjects continued to vary their estimate widely. The other two subjects largely kept their estimate constant in the final session

sections, $d = .745$). The persistence of the step-hold pattern in the behavioral read-out of the perceived p_g , even when it does not mimic the pattern of changes in the hidden parameter, suggests that step-hold behavior is an inherent property of probabilistic parameter perception in humans, not a volitional strategy that comes into play only when the step-hold pattern in the Bernoulli parameter encourages it.

Accuracy There are two measures of ground truth against which to compare our subjects' performance across all trials. The first ground truth measure is the actual hidden p_g value from the experiment. The second is the parameter estimated by an ideal observer. Here, we take our ideal observer to be the online Bayesian model of Adams and Mackay (2007), which estimates the run-length r of a non-stationary

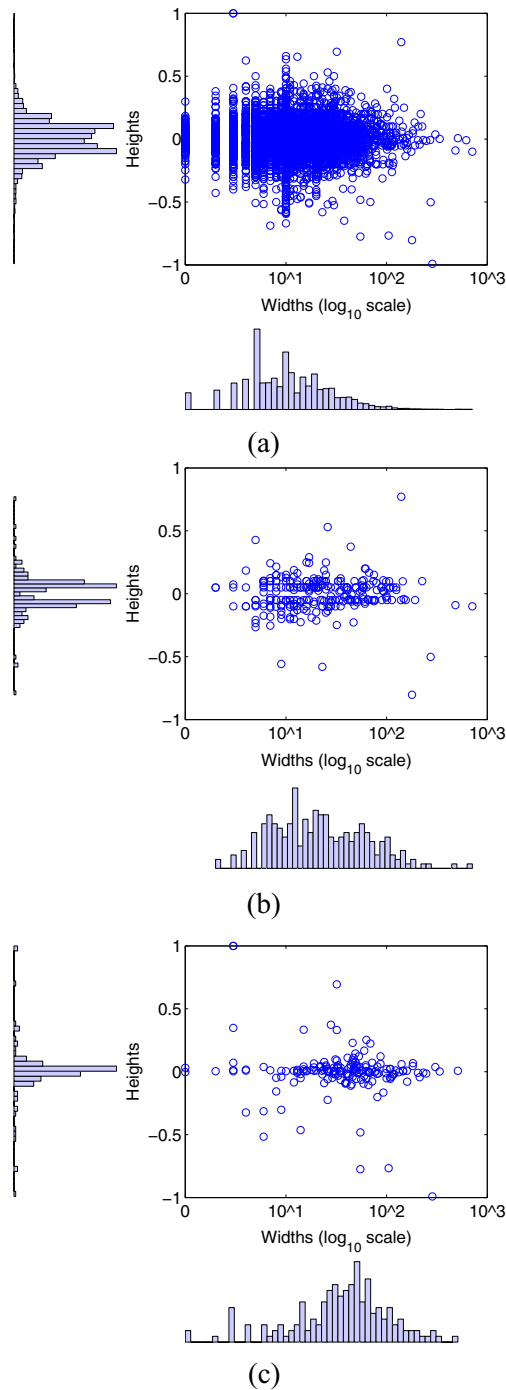


Fig. 5 Joint distributions for step widths and heights. **A.** Across all subjects, heights are bimodally distributed, while the widths distribution is broad and unimodal. **B.** The joint distribution for Subject DD, which mimics the bimodality seen in the pooled distribution. **C.** Joint distribution for Subject BC; this subject’s height distribution is unimodal

stochastic process. At time step t , the algorithm updates a set of t conjugate priors on p_g and r , one for each possible past change-point. Then, by determining the maximally likely run-length at t , it determines the maximally likely value for p_g (details in Adams and Mackay (2007)).

Table 1 Subject accuracies. The difference between aperiodic and periodic subjects, compared to the ideal observer, was insignificant (bold-faced values in final column)

	Aperiodic	Periodic	Cohen’s d
RMS + True p_g	.028	.042	.230
RMS + Ideal	.026	.028	.025
KL + True p_g	.092	.193	.369
KL + Ideal	.085	.095	.046

Additionally, there are two measures of error: the root mean square error across all trials and the mean Kullback–Leibler divergence between the subject’s estimate and ground truth. This second error represents the additional cost, measured in bits, of assuming the distribution has the estimated parameter, when the ground truth is different. We report the performance results, for both ground truth measures and error measures, in Table 1.

Note that the only appreciable effect sizes occur when ground truth is taken as the true p_g . When compared to an ideal observer, however, there is no substantial difference between aperiodic and periodic subjects. This is true for both RMS and KL error measures. KL divergence is an important error measure, since it describes the information theoretic strain undergone by the memory substrate of subjects. The equality of performance between groups compared to the optimum is noteworthy since periodic subjects had a qualitatively more stressful task. The true parameter for periodic subjects was nowhere stationary, so that they could never hold the slider still for long. Indeed, periodic subjects moved the slider an average of 785.25 times in the experiment, compared to only 348.60 times in the aperiodic condition, and aperiodic subjects waited 17.377 trials longer between slider moves, on average, than periodic subjects ($d = .745$).

Additionally, we analyzed whether or not there was an effect of the true p_g on our subjects’ error. This effect, too, depended on the combination of ground truth and error measure in a visually obvious manner (Fig. 6). There is a stark difference in the effect of the true p_g on the two groups for RMS error. Periodic subjects tended incur more error when p_g was close to 0 or 1, resulting in the V shape of Fig. 6b. This difference largely disappears for KL error (Fig. 6c, d). Note that the non-linearity of the KL-divergence tends to make it large near 0 or 1 anyway, resulting in the peaks in the last two bins. Despite this, compared to the ideal observer, there is little appreciable difference between aperiodic and periodic subjects in the effect of p_g .¹ For example, aperiodic subjects had an average KL error of .025 bits in the

¹The weak undulation in the light blue bars of Fig. 6d arises from subjects’ slight tendency to misestimate the crests and troughs of p_g (as we confirmed with some basic signal analysis and is visible in Fig. 4a).

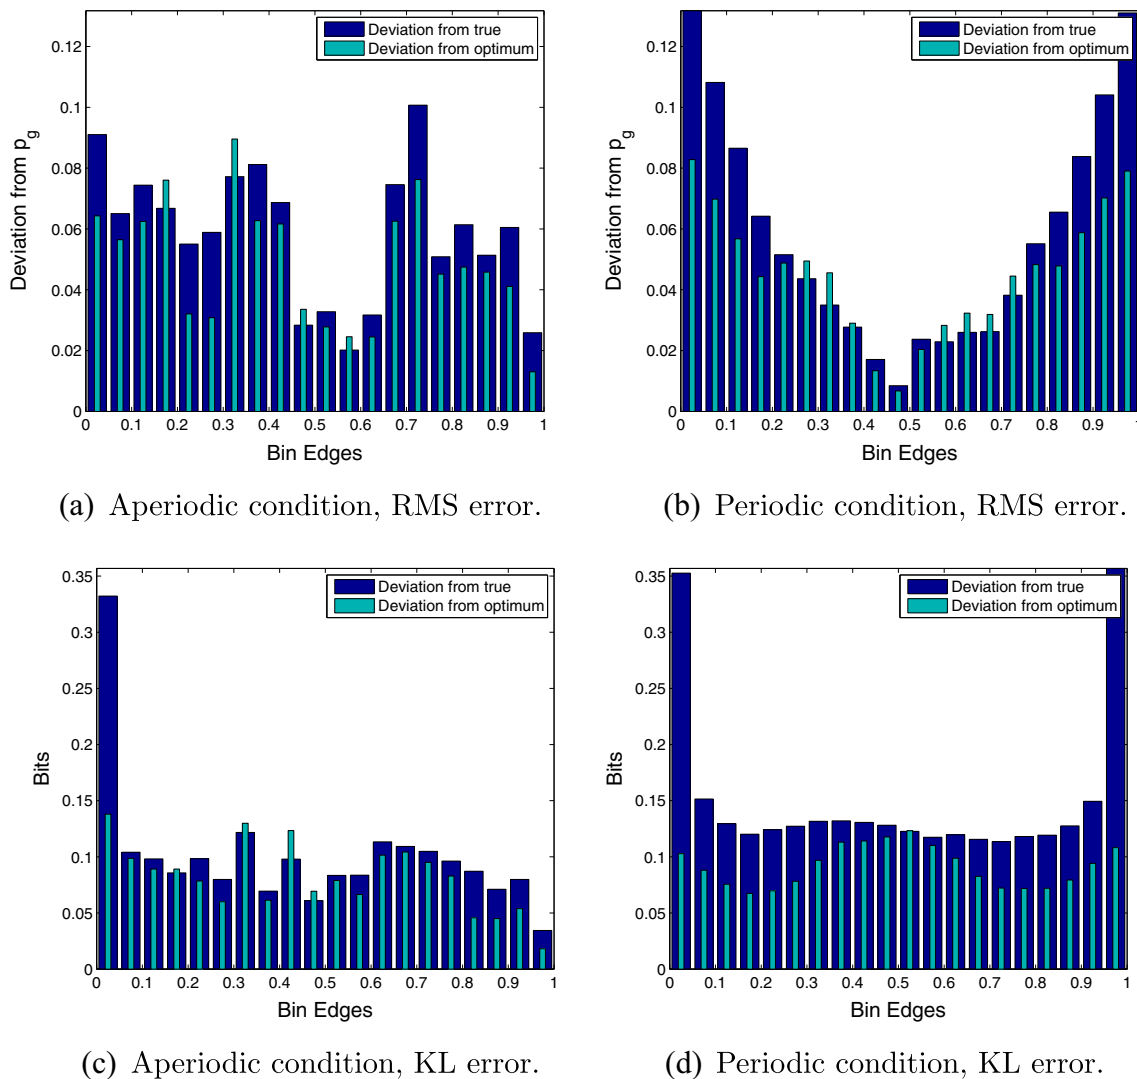


Fig. 6 Binned subject error. Each panel depicts one combination of condition and error type across 20 p_g bins of size .05. Compared to aperiodic subjects, periodic subjects show a distinct V-shaped RMS

error across p_g bins due to the misestimation of parameter crests and troughs. However, this difference is largely diminished according to the more information theoretically meaningful KL measure

p_g bin centered at .95; this means that, subjects wasted 1 bit of memory every 40 trials which happened to feature a true probability in that value range. In the same bin, periodic subjects wasted 1 bit every ten trials, a small difference in absolute terms.

In Gallistel et al. (2014), the authors reported no appreciable effect of the true p_g value on accuracy. At first, the RMS results for periodic subjects in the current experiment seem to run counter to the original finding; they appear to recapitulate some aspects of the substantial literature on estimation bias (Kahneman and Tversky, 1979; Hertwig et al., 2004) demonstrating systematic distortion of probabilistic estimates for rare events. However, the fact that this effect was not borne out in the aperiodic condition suggests that other phenomena might be at play in our case. For example, we found that the average run-lengths of trials in

the aperiodic condition for which the parameter exceeded .9 or fell below .1 were 565.3 and 215, respectively; those values both drop to 41 trials for the periodic condition. It seems likely that subjects in the periodic condition simply had less time to adjust to the extreme p_g values before the parameter returned to moderate values, all the more likely when one considers the change-point detection latencies reported below. If subjects can detect the underlying rate-of-change of the parameter, as we argue below, then there might be a further effect of the p_g derivative that causes periodic subjects to incur more RMS error near crests and troughs: Away from peaks, the derivative of the parameter is close to constant (since a sinusoid here is approximately linear by the small-angle approximation), so subjects can make slider movements at regular intervals. At extreme values, however, the derivative quickly switches sign, so that

subjects must, from stochastic samples alone, sense that the direction of the slider movements must now change. From the point of view of subject strategy, this is a more taxing moment. Again, the distortion of error near extremes does not occur for KL error (except for the boundary bins where the KL-divergence blows up to $\pm\infty$).

Hence, by the information theoretically grounded KL error measure, subjects in both groups showed uniform tendency to incur error across all p_g values. This was also evident when we instead considered median slider estimates compared to ground truth. Across all tested hidden parameters, the mapping from median subject estimate to the true parameter is the identity, plus or minus a quartile. This is consistent with (Robinson, 1964) experiment, the review of the early literature by Peterson and Beach (1967), and our own previous work. For more on the accuracy of subjects near extreme p_g values, see the Discussion of Gallistel et al. (2014).

Finally, we examined each subject's error across sessions to look for an effect of experiment duration on performance. Except for subject 'BC', there was no evident effect of session on performance. 'BC', beginning at session 6, began to fluctuate in performance somewhat wildly. Almost uniformly, periodic subjects incurred greater error across sessions than did aperiodic subjects, again, with the exception of subject 'BC'. There was no significant effect of experiment duration on performance, either from fatigue or from adjustment of strategy. Additionally, we measured time taken per trial and found neither an effect of experiment duration nor a correlation with error.

Structural performance

Change-point detection In Robinson (1964) and Gallistel et al. (2014), change-points were trials at which the hidden parameter made discrete jumps. In the current paradigm, changes in the hidden parameter were smooth. We define change-points in this setting to be those trials at which the hidden parameter reaches an extremum. Change-points are either isolated peaks or valley bottoms in slider settings or the boundaries of stationary periods. We define a subject's change-point detection latency as the number of trials after a change-point that it takes for the subject to adjust the slider in the direction of the new parameter value. The median latency of the median subject was 29 trials. Average latency for subjects given aperiodic hidden parameters was longer than that of subjects in the periodic setting (41.2 trials for aperiodic versus 31.5 trials for periodic, $d = .499$). Aperiodic Change-points sometimes occurred in close succession or only shifted p_g a small amount, making them in principle undetectable before the next change occurred. Nonetheless, the average percentage of change-points detected across all subjects was high (92.36%). The

four subjects in the periodic paradigm detected each change-point, while the five aperiodic subjects detected 86.25% ($d = .589$). Further, there was no significant interaction between change-point detection latency averaged over subjects and session number. In other words, detection was as speedy in early sessions as it was in later sessions.

Detection of underlying structure In the aperiodic condition, the underlying parameter p_g had no deterministic structure across trials. Therefore, only subjects in the periodic condition might have perceived the regular structure of the underlying parameter. Earlier work by Estes (1984) tested subjects' sensitivity to periodicity in the generating parameter of a Bernoulli distribution by first conditioning them to the periodic parameter (period was 80 trials) and then suddenly fixing the parameter for many trials at .5. When his subjects continued to move the slider sinusoidally, (Estes, 1984) concluded they had explicitly encoded the periodicity of the earlier trials.

Unlike in Estes' experiment, our subjects did not continue to move the slider periodically after the parameter flatlined in the final session (Fig. 4). Indeed, as we postulate that subjects are trying to minimize the KL divergence between their estimate and the true distribution, continuing sinusoidal slider movement would be a bad strategy. Two subjects (Fig. 4b, d) seemed to carry the volatility of slider movement from the first 9 sessions to the final session, but signal analysis revealed no periodicity. However, during debriefing, all 4 subjects spontaneously remarked that the probability changed periodically. We take these unprompted declarations as a confirmation of Estes' finding that subjects can detect the periodic structure underlying the data.

Besides the declaration of the subjects, their ability to detect periodicity is evident in their performance data. A sinusoidally varying p_g consists of alternating increasing and decreasing portions. Thus, if subjects are sensitive to the global model generating the data, they could use this knowledge to better detect the derivative of p_g . To test for this effect, we compared the tendency of subjects to move the slider in the correct direction between the aperiodic and periodic conditions. For example, moving the slider up when the true p_g was increasing is considered a correct trial by this measure. We calculated the average correct slider movements across four regimes: for every trial, for all trials on which the subject moved the slider, for those trials on which the true p_g moved, and finally when both the slider and true p_g moved (Fig. 7).

The step-hold behavior of subjects means that, overwhelmingly, all subjects tacitly estimate the derivative of p_g as 0. Therefore, when we calculated correct slider movements across all trials, we found higher performance in the aperiodic condition ($d = .612$), in which subjects benefited from the many trials of true stationarity. However, the

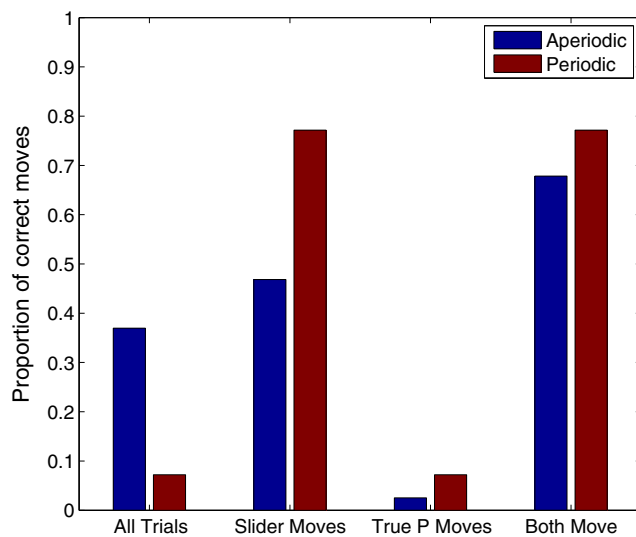


Fig. 7 Performance on structure-dependent measures. *All trials*) Aperiodic subjects move the slider in the correct direction more often across all trials. This is because they benefit from the frequent stationarity of the parameter. *Slider moves*) If a subject decided to move the slider, he/she moved it in the correct direction significantly more often if he/she was in the periodic condition. *True P moves*) Whenever the true p_g was moving, periodic subjects moved their slider in the correct direction more, but not significantly. *Both move*) On those trials during which both the slider and the true p_g moved, aperiodic and periodic subjects were both very accurate, though periodic subjects performed marginally better

opposite obtained when we restricted the calculation to only those trials on which subjects moved the slider ($d = .609$). That is, whenever subjects moved the slider, they tended to move it in the correct direction more in the periodic condition, with large effect. The other two regimes, true p_g moving and both moving, gave moderate effect sizes ($d = .297$ and $d = .198$, respectively), though periodic subjects did have higher means. We take the fact that the two groups deviated on derivative detection measure as evidence that subjects can detect the higher order structure generating the data.

Discussion

Our results lend further support to the conclusion that the step-hold pattern seen in subjects' slider settings (or, in Robinson's case, lever settings) accurately reflects the characteristics of the underlying process for forming a perception of a Bernoulli probability. They imply that the computational process that yields the percept does not change the percept each trial. Step-hold behavior is seen even when the change in p_g on any trial is very small, and even when subjects realize that the changes are gradual and predictable

Preparatory to discussing their theoretical implications, we summarize the properties of the perceptual process so

far revealed by the small literature that tracks the perception of an unfolding non-stationary Bernoulli probability observation by observation (Robinson, 1964; Gallistel et al., 2014):

- The percept is not updated following each observation; it may go unchanged for hundreds of observations, even when the hidden parameter changes smoothly and by very small amounts between observations (Figs. 3, 4, 5; see also (Robinson, 1964), p. 11, and Figs. 5 and 11 of Gallistel et al. (2014), pp. 102,105).
- The distribution of update magnitudes (step heights) across all subjects peaks around the smallest possible update under most circumstances (Fig. 5; see also Fig. 11 of Gallistel et al. (2014), p.105).
- However, updates spanning most of the possible range (0 to 1) frequently occur following large changes in the hidden parameter (Fig. 5; see also Figs. 5 and 11 of Gallistel et al. (2014), p. 102, 105).
- To a first approximation, the function mapping from the hidden parameter to the perceived parameter is the identity (see also Fig. 6 of Gallistel et al. (2014), p. 102).
- The accuracy of the perceived parameter relative to the parameter estimated by an ideal observer is generally good. After any given observation, the median percept is sufficiently close to the underlying truth that it would take about 100 observations to detect the error (Figs. 17 and 18 of Gallistel et al. (2014), p. 114).
- When measured by its Kullback-Leibler divergence from the ideal observer's parameter, the accuracy of the perceived parameter is approximately the same over all but the most extreme values for the hidden parameter (Fig. 6; see also Fig. 18 of Gallistel et al. (2014), p.114).
- Substantial changes in the hidden parameter are reliably and rapidly perceived; they are events in their own right (Gallistel et al., 2014).
- The perceptual process is appropriately sensitive to the prior odds of a change in the parameter, that is, to the volatility: The relative-likelihood threshold for the detection of a change in a sequence of any given length is lower when the volatility is high (Robinson, 1964; Gallistel et al., 2014).
- Subjects have second thoughts about previously perceived changes in the hidden parameter ((Gallistel et al., 2014)). After more observations—sometimes many more observations (Fig. 9 of Gallistel et al. (2014), p.104)—they conclude that their most recent perception of a change was erroneous.
- Smooth sinusoidal changes in the hidden parameter are perceived as periodic (present paper; see also (Estes, 1984)).

We divide our discussion of the theoretical implications into two parts. In the first, we show how the model of the perceptual process proposed in Gallistel et al. (2014) explains the results. In the second, we discuss the challenges that the results pose for models that assume trial-by-trial updating of the percept, with no record of the sequence of observations that generated the current percept.

The IIAB model

In IIAB (Fig. 8), the current percept arises from a computation that constructs a compact history of the stochastic process that is assumed to have generated the observed outcomes. There are two motivations for constructing such a model of the stochastic process: it minimizes long term memory load by providing the basis for a lossless compression of the sequence of generating distributions already observed, and it best predicts the outcomes not yet observed. The model that best achieves both of these goals is the model that best adjudicates the trade-off between the complexity of the representation and the accuracy with which it captures the observed sequence (see (Grunwald et al., 2005), Chapters 1 & 2). In a change-point model, the more change-points added to it, the more complex it becomes.

However, adding change points also makes it more accurate, further reducing the cost of storing the observed sequence of outcomes using that model. A model of the process that constructs the change-point representation must address the problem of deciding in real time whether the increased accuracy due to an added change-point is worth the increased complexity of the representation. In IIAB, this decision is mediated by Bayesian model selection, because it takes model complexity into account in a principled way.

It is computationally much simpler to decide whether the current estimate of the hidden parameter adequately explains recent observations than it is to decide whether those observations justify increasing the complexity of the parameter history with a new change-point or reducing it by dropping an earlier change point. Therefore, (Gallistel et al., 2014) assume a first stage that computes a measure of how poorly the current estimate of the hidden parameter is doing (left half, Fig. 8). If the current estimate is doing well, there is no further computation. This first stage explains the step-hold pattern: much more often than not, the current estimate is doing fine (“If it ain’t broke...”), so there is no reason to revise it (“...don’t fix it.”). The model generates a distribution of step widths that is a reasonable approximation to the

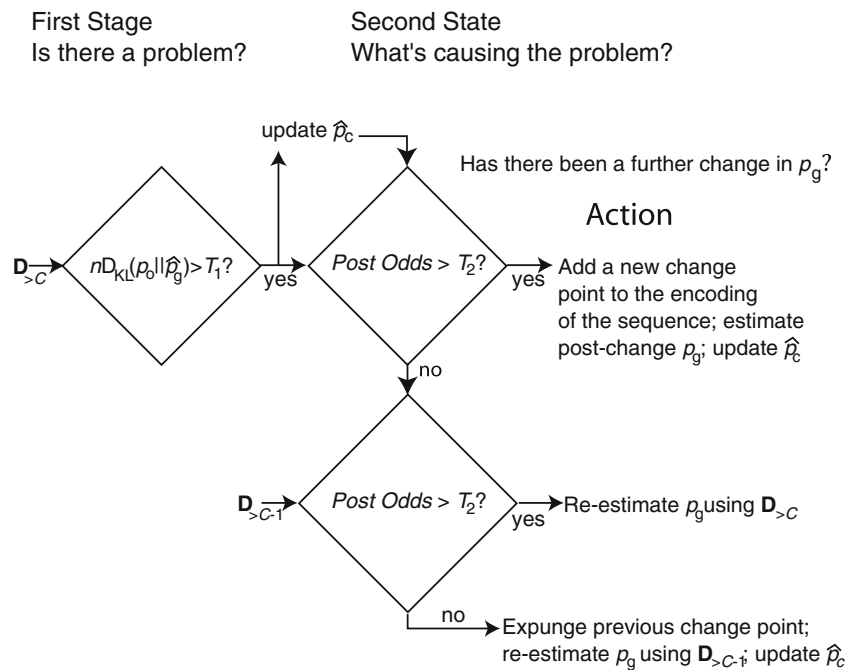


Fig. 8 *The IIAB model.* In the first stage, the sequence of data since the last change-point, $D_{>c}$, is used to calculate the empirical frequency of green rings against which is compared, in the Kullback-Leibler sense, the current estimate. If the KL divergence times the number of observations n in $D_{>c}$ exceeds a threshold T_1 , the model proceeds to the second stage. Meanwhile, the estimated probability of a change-point, \hat{p}_c is updated in a Bayesian way. The second stage adjudicates between three options using Bayesian model selection. If the posterior odds of a change-point having happened are greater than a threshold

T_2 , this change-point is added at the maximally likely spot in the sequence $D_{>c}$. If not, then it provisionally removes the last change-point, recalculates the posterior odds to see if they now exceed T_2 , and checks if the change-point can be replaced at a different location in the sequence $D_{>c-1}$ of data since the penultimate change-point. If the posterior odds still do not exceed T_2 , then the provisionally removed change-point is permanently expunged and Bayesian updates are performed on the estimated parameter \hat{p}_g and the estimated probability of a change-point \hat{p}_c . Figure taken from (Gallistel et al., 2014)

distribution generated by subjects ((Gallistel et al., 2014), Fig. 15, p. 112)

Only when the first stage decides that the estimate of the current value of the hidden parameter is broken does a second stage become active (right half, Fig. 8). It uses Bayesian model selection to decide among three explanations:

1. There has been no further change, but the current estimate of p_g needs to be improved in the light of the data obtained since it was first made. These changes in the estimate are generally small, because they are corrections to the small-sample errors, based on a larger sample. These small corrections are relatively numerous. That is why the distribution of step heights produced by the model generally has a single mode at the smallest corrections, as do the distributions generated by subjects ((Gallistel et al., 2014), Fig. 15, p. 112). However, depending on the thresholds governing transitions between the two stage of IIAB, the model can produce both bimodal and unimodal distributions of step heights, like the subject data in Figs. 5b and c respectively.
2. There has been a further change in p_g , in which case, a new change point is added to the evolving model of the process history, and p_g is re-estimated using only the data since this newly added change is estimated to have occurred. When this occurs, the model makes arbitrarily large one-trial jumps in its estimate of the current probability, because that new estimate is based only on the portion of the sequence observed since the estimated location of the most recent change in the p_g .
3. The change point most recently added to the representation of the process history is not justified in the light of the data seen since it was added. In that case, it is removed from the model of the process history, and p_g is re-estimated from the observations stretching back to the penultimate change point in the estimated history of the process. When this occurs, the model has second thoughts; it retroactively revises its representation of the history of the process.

The mapping from the current value of the hidden parameter to the model's estimate approximates the identity over the full range of p , as is the case for the subjects'. And, the model's estimates, like the subjects', are approximately equally accurate over the full range. The model's estimates are more accurate than the subjects', but, the model is implemented with a double-precision floating point representation of all the quantities, that is, with $1/2^{53}$ precision. By contrast, the Weber fraction for adult human subjects' representations of numerosity are on the order of $\pm 12.5\%$ (Halberda and Feigenson, 2008), which implies approximately $1/2^4$ precision.

The model detects changes with hit rates and false alarm rates similar to those of the subjects ((Gallistel et al., 2014), Fig. 8, p. 103) and with similar post-change latencies ((Gallistel et al., 2014), Fig. 7, p. 103). Its second thoughts about the changes it detects occur at latencies comparable to the latencies at which subjects report their second thoughts ((Gallistel et al., 2014), Fig. 9, p. 104).

The model estimates the probability of a change, that is, the volatility, and it uses that estimate to compute the prior odds. In the basic Bayesian inference formula, the prior odds scale the Bayes Factor. Thus, in the model, increased volatility (as reflected in the estimate of the prior odds) increases the sensitivity to within-sequence evidence for a change (as reflected in the Bayes Factor). This explains qualitatively the subject's sensitivity to the prior odds. It explains it too well, however, in that the model converges on an accurate estimate of the prior odds more rapidly than subjects do.

Although the model constructs a representation of parameter history, it is not explicitly sensitive to higher-order structure. Our subjects, on the other hand, revealed their sensitivity to this structure in both their improved performance on structure-dependent measures and by their explicit detection of periodicity. In fact, even the ability of subjects in Gallistel et al. (2014) to retrospectively decide that one of their change-points was a mistake indicates they had computational access to the parameter history. Presumably, our subjects' ability to detect periodicity rested on just this computational access.

It is easy to see how IIAB could be improved by adding computational access to the parameter history. For example, given the two points in the parameter history $\{(t_1, p_1), (t_2, p_2)\}$, one could calculate the slope of the secant line between t_1 and t_2 , $m = \frac{p_2 - p_1}{t_2 - t_1}$. This simple computation indicates that, between trials t_1 and t_2 , the parameter seems to be changing at a rate m . If one assumes a sufficiently smooth underlying parameter, one might allow m to bias the future estimate of the current parameter. When this functionality is added to IIAB,² it can regularize slider movement and decrease reaction time to sudden changes (Fig. 9b) This is only possible with a memory of past change-points.

Because the model treats changes in the hidden parameter as events in their own right, it is inherently recursive, that is, it will bring to bear on these perceived events that same probability-estimating process that generated the perceptions of the changes. Recursive application of IIAB builds a hierarchical representation of the parameter in

²To introduce sensitivity to the derivative, we adjusted the "effective" number of green rings seen by the model since the last change point as a function of the derivative. The effective number was scaled between 0 and the current run-length with a Gompertz function, a type of asymmetrical sigmoid.

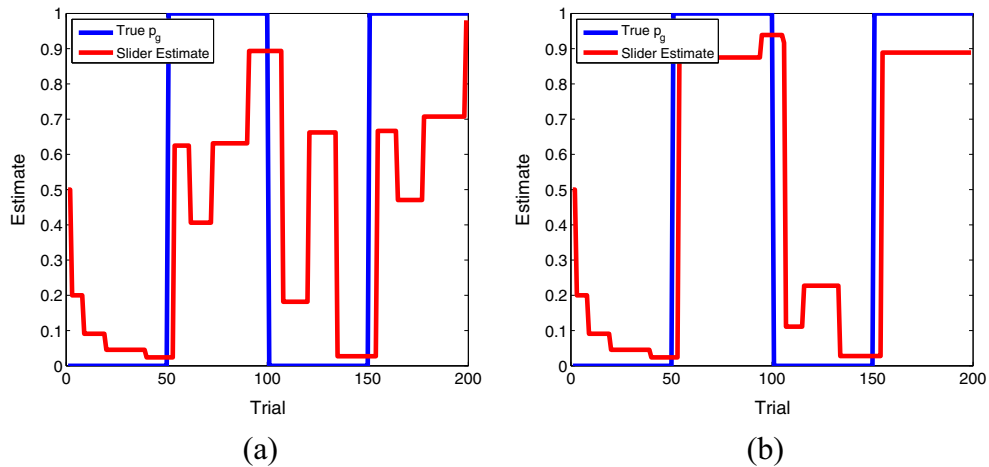


Fig. 9 Derivative sensitivity improves IIAB performance on toy data. **a)** When the model has no derivative sensitivity, it reacts slowly and incompletely to sudden changes (e.g., trials 50, 100). **b)** When the

memory (a two-level structure created by IIAB is shown in Fig. 10). At the bottom of the structure is an encoding of the observed sequence. One level up is an encoding of the parameter-history string. At a second level is an encoding of a parameter of that history string, namely, the frequency with which changes occur. Higher levels would encode changes of change-points, etc. Robinson's (1964)

model is biased to think that the parameter will continue to change at the rate indicated by a local derivative, it adjusts its estimate more rapidly and completely

results suggest that included in the second level is an encoding of the distribution of change magnitudes (step heights).

A hierarchical organization of events makes possible greater data compression and more powerful prediction. The detection of higher-order structure explains both Estes' result and Robinson's finding that his subjects sensed the

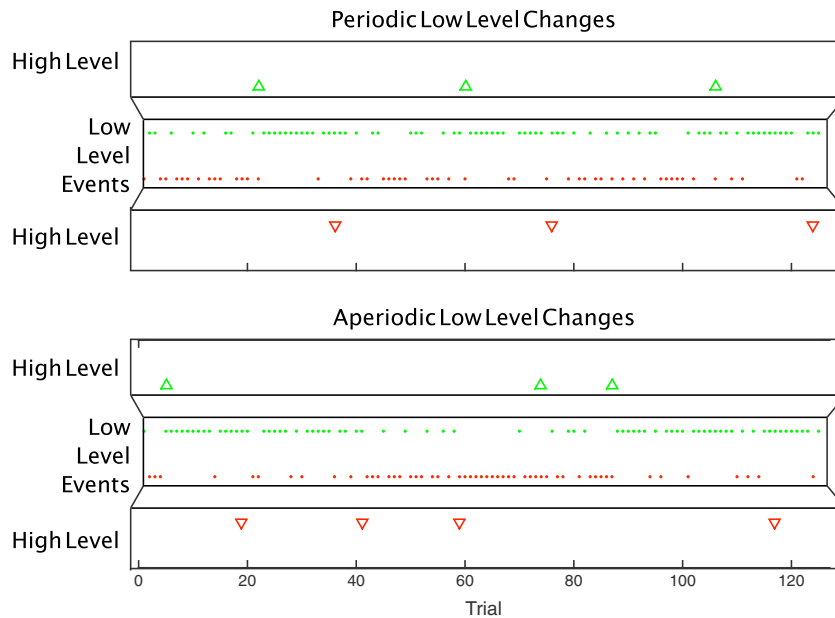


Fig. 10 Event hierarchy. At the bottom level (Low Level Events) is the stream of Bernoulli events (*draws from the urn*) in which either a green (g) or red (r) circle is drawn, as indicated by the *dots* on the g and r lines. In the *upper panel*, the hidden parameter, p_g , changed every 20 draws between two levels, .25 and .75; thus, they occur periodically. In the *lower panel*, there was a .05 probability of such a change from one of these levels to the other after each draw; thus, the changes occur aperiodically. The perceived changes are themselves perceptual experiences. The model of change detection determines where these changes are perceived to have occurred. These perceptions are subject

to error; sometimes a change is not perceived and sometimes one is perceived when none occurred. And, the perceived locus of a change often deviates somewhat from the draw on which the hidden change in fact occurred. The loci of the perceived up and down changes constitute a second level event stream, as indicated by the *upward* and *downward pointing triangles* on the up and down lines at the 2nd level of the event hierarchy. These triangles are more regularly spaced when the bottom-level event stream changed periodically than when it changed aperiodically

difference between his un signaled blocks of small-change and large-change problem sets.

The hierarchical organization outlined above may allow the detection of higher-order structure, but, unless the set of possible higher-order structures is constrained in some way, detection may be infeasible. Hierarchical representation gives access to local derivative information, but it does not offer a simple way to use this local information to deduce the global model generating the data. For example, our subjects claimed not just that the parameter consisted of increasing and decreasing portions, but that the parameter was “periodic.” They had discovered a way to map the hierarchical structure of the parameter history to a formal data-generating model, a sine wave. As a global model, the sine wave determines all p_g 's across trials, past and present. The hierarchical memory structure alone does not uniquely determine a generating model, and therefore requires some additional constraints. We consider the elucidation of these constraints a key challenge for future work.

The challenges for trial-by-trial-updating models

At this point in theory development, it is not possible to compare the performance of the numerous trial-by-trial-updating models of probability perception, like (Yu and Dayan, 2005; Wilson et al., 2010), or even Kalman filters, to the performance of the IIAB model, because none of the other extant models known to us attempts to explain many of the above-listed properties of the process that generates a subject's perception of the current probability.³ All of the trial-by-trial models known to us attempt only to explain the tracking of the probability, and they all implicitly assume that the subject has in memory only an estimate of the current probability and the current volatility. None of them posits a record in memory of the sequence on which the currently perceived probability is based, nor a record of the history of that hidden parameter. The IIAB model's assumption that subjects have a record of the sequence of outcomes, which is at the foundation of the model, is also its most controversial assumption. It is, we believe, the assumption that most theorists are, understandably, the most reluctant to make.

None of the extant trial-by-trial-updating models has been applied to the data on subjects' observation-by-observation perception of a non-stationary hidden probability. To apply them, we would have to make additional

assumptions, assumptions that the authors of a given model may not embrace. For example, it is easy to get a trial-by-trial, delta-rule updating models to exhibit step-hold behavior by adding a threshold between the running average produced by the delta-rule updating, which changes after almost every observation, and the current percept. Only when the running average deviates from the current percept by a supra-threshold amount, does the current percept change. Or, under another interpretation of what is mathematically the same assumption: maybe the step-hold pattern does not reflect a property of the underlying percept, but only a property of the decision process leading to a change in the setting of the slider or the lever, which is the experimentally observed subject behavior. Gallistel et al., 2014 ran simulations of a variety of assumptions of this sort and with many different values for the output threshold. Their simulations demonstrated the reality of an intuitively obvious problem: when the threshold is set high enough to produce steps remotely as wide as those produced by subjects, it eliminates or greatly reduces the steps with small heights, but these small steps are in fact the ones that subjects most frequently make. Thus, the assumption of a threshold on the output is probably not one that the authors of a trial-by-trial updating model would want to make. The question therefore remains: What assumption does one want to make that will explain the fact that subjects do not update their percept observation by observation even though each observation has a non-trivial impact on the estimate based on either a running average (generated by delta-rule updating) or on the mean of the Bayesian posterior.

For a second example: None of the extant models explains the fact that subjects perceive the changes themselves. The models focus only on the subjects' ability to track the changes. A seemingly simple way to imbue delta-rule models with the ability to perceive the changes is to assume a fast and a slow running average. So long as the two averages give roughly comparable values for the estimated parameter, the subject perceives the average with the longer decay time because it will be more accurate when there has not been a recent change. When, however, that estimate differs from the estimate delivered by the fast average (the one with the rapid decay) by a supra-threshold amount, a change is perceived to have occurred, and the current percept of the parameter is then based on the fast average, the one least influenced by the more distant past. It remains based on the fast average until the difference between the slow and fast average falls below the threshold. Gallistel et al., 2014 ran simulations of delta-rule updating models when augmented by this assumption. In their simulations, these models always produced outlying dips in the distribution of step heights, which dips have never been observed in any subject. Thus, this is probably not an assumption that authors of delta-rule updating models

³IIAB does share some important similarities with these other models, even though they addresses fundamentally different questions. Like IIAB, (Yu & Dayan, 2005) only uses recent parameter history for estimation to make inference tractable. Both (Wilson et al., 2010), Kalman filters and IIAB estimate process volatility.

would want to embrace in order to explain the fact that the changes are themselves perceptible events. The question therefore remains: How does one want to explain the fact that a step change in the hidden parameter of a Bernoulli process is itself a perceived event. Moreover, the volatility results suggest that the probability of a change event is also perceived. In subsequent work, it would be interesting to verify this by asking subjects to indicate observation by observation their perception of the current probability and the probability of a change in that probability.

A Bayesian tracking model for the ideal observer (Adams & Mackay, 2007) can produce abrupt changes in the estimates of the hidden parameter. However, the Adams and Mackay model—which was not intended as a psychological model—has the following property: At any given time, it has an estimate of parameter based only on the most recent outcome, an estimate based only on the 2 most recent outcomes, an estimate based only on the 3 most recent outcomes, and so on backwards through the observed sequence. Moreover, it has an estimate of the likelihood that there was a change before the most recent outcome, and an estimate of the likelihood that there was a change before the second most recent outcome, and so on backward through the sequence for many outcomes. Thus, it has a form of the sequence-memory assumption that is the most objectionable feature of the IIAB model. And, like all trial-by-trial-updating models, its estimate of the current parameter changes after almost every observation.

The fact that subjects have second thoughts about previously perceived changes is another challenge. These second thoughts often arise many trials after reporting those perceptions. To us, these second thoughts are perhaps the strongest evidence in favor of the seemingly implausible assumption that subjects have some record, however rough, of the observed sequence of outcomes. Why should the underlying process not simply generate yet another change perception in order to explain the discrepancy between what was perceived back then, when the preceding change was reported, and what observations since then suggest? It seems that the underlying process weighs the evidence from the observations that postdate that earlier perception along with the observations that led to that earlier perception. But how can it do that if it has no record of those earlier observations? Thus, we take this to be another important challenge.

Finally, like Estes (1984) we view the evidence that subjects can recognize higher order structure in the observed sequence of outcomes as a challenge to any model that assumes no record of the sequence of outcomes. If the brain has no record of the sequence history, how can it decide on a stochastic model for that history? Future work could probe subjects' ability to classify parameter histories purely from noisy samples and could investigate the depth of

hierarchical organization available to humans' probability perception mechanism.

References

- Adams, R. P., & Mackay, D. J. C. (2007). Bayesian online changepoint detection. arXiv:0710.3742
- Behrens, T. E. J., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. *Nat. Neurosci.*, *10*(9), 1214–21. doi:10.1038/nn1954. <http://www.ncbi.nlm.nih.gov/pubmed/17676057>
- Brown, S. D., & Steyvers, M. (2009). Detecting and predicting changes. *Cogn. Psychol.*, *58*(1), 49–67.
- Estes, W. K. (1984). Global and local control of choice behavior by cyclically varying outcome probabilities. *J. Exp. Psychol. Learn. Mem. Cogn.*, *10*(2), 258–270. doi:10.1037/0278-7393.10.2.258
- Gallistel, C. R., Krishan, M., Liu, Y., Miller, R., & Latham, P. E. (2014). The perception of probability. *Psychol. Rev.*, *121*, 96–123. doi:10.1037/a0035232. <http://www.ncbi.nlm.nih.gov/pubmed/24490790>
- Glimcher, P. W. (2003). The neurobiology of visual-saccadic decision making. *Annu. Rev. Neurosci.*, *26*(1), 133–179. doi:10.1146/annurev.neuro.26.010302.081134
- Grunwald, P. D., Myung, I. J., & Pitt, M. A. (2005). *Advances in minimum description length theory and applications*. Cambridge, MA: MIT Press.
- Halberda, J., & Feigenson, L. (2008). Developmental change in the acuity of the “number sense”: The approximate number system in 3-, 4-, 5-, and 6-year-olds and adults. *Dev. Psychol.*, *44*(5), 1457–65. doi:10.1037/a0012682. <http://www.ncbi.nlm.nih.gov/pubmed/18793076>
- Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychol. Sci.*, *15*(8), 534–539.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263–292. arXiv:1011.1669v3. doi:10.2307/1914185. <http://www.jstor.org/stable/1914185>
- Krugel, L. K., Biele, G., Mohr, P. N. C., Li, S.-C., & Heekeren, H. R. (2009). Genetic variation in dopaminergic neuromodulation influences the ability to rapidly and flexibly adapt decisions. *Proc. Natl. Acad. Sci. U.S.A.*, *106*(42), 17951–6. doi:10.1073/pnas.0905191106
- Peterson, C. R., & Beach, L. R. (1967). Man as an intuitive statistician. *Psychol. Bull.*, *68*(1), 29–46. doi:10.1037/h0024722
- Robinson, G. H. (1964). Continuous estimation of a time-varying probability. *Ergonomics*, *7*(1), 7–21. doi:10.1080/00140136408930721
- Sugrue, L. P., Corrado, G. S., & Newsome, W. T. (2004). Matching behavior and the representation of value in the parietal cortex. *Science (80-.)*, *304*(5678), 1782–1787.
- Sugrue, L. P., Corrado, G. S., & Newsome, W. T. (2005). Choosing the greater of two goods: neural currencies for valuation and decision making. *Nat. Rev. Neurosci.*, *6*(5), 363–75. doi:10.1038/nrn1666. <http://www.ncbi.nlm.nih.gov/pubmed/15832198>
- Wilson, R. C., Nassar, M. R., & Gold, J. I. (2010). Bayesian online learning of the hazard rate in change-point problems. *Neural Comput.*, *22*(9), 2452–76. http://www.mitpressjournals.org/doi/full/10.1162/NECO_a.00007
- Wilson, R. C., Nassar, M. R., & Gold, J. I. (2013). A mixture of delta-rules approximation to Bayesian inference in change-point problems. *PLoS Comput. Biol.*, *9*(7). doi:10.1371/journal.pcbi.1003150
- Yu, A. J., & Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron*, *46*(4), 681–692. doi:10.1016/j.neuron.2005.04.026