Published in final edited form as:

J Exp Anal Behav. 2017 July ; 108(1): 39–72. doi:10.1002/jeab.261.

THEORETICAL IMPLICATIONS OF QUANTITATIVE PROPERTIES OF INTERVAL TIMING AND PROBABILITY ESTIMATION IN MOUSE AND RAT

Aaron Kheifets¹, David Freestone², and C.R. Gallistel¹

¹Department of Psychology, Rutgers University

²Department of Psychology, Brown University

Abstract

In three experiments with mice (*Mus musculus*) and rats (*Rattus norvegicus*), we used a switch paradigm to measure quantitative properties of the interval-timing mechanism. We found that: 1) Rodents adjusted the precision of their timed switches in response to changes in the interval between the short and long feed latencies (the temporal goalposts). 2) The variability in the timing of the switch response was reduced or unchanged in the face of large trial-to-trial random variability in the short and long feed latencies. 3) The adjustment in the distribution of switch latencies in response to changes in the relative frequency of short and long trials was sensitive to the asymmetry in the Kullback-Leibler divergence. The three results suggest that durations are represented with adjustable precision, that they are timed by multiple timers, and that there is a trial-by-trial (episodic) record of feed latencies in memory.

Keywords

coefficient of variation; endogenous variability; exogenous variability; Kullback-Leibler divergence; timing theories; mouse; rat

We report three experiments measuring quantitative properties of the interval-timing and probability-estimation mechanisms in the mouse and the rat. They ask questions that bear strongly on theories of response timing. They use the switch paradigm, which was first used by Fetterman and Killeen (1995), with pigeons. We have modified it for use with rodents (Balci, Freestone, & Gallistel, 2009; Balci et al., 2008). In this procedure, the subject judges whether and when to switch from one feeding hopper to another as time elapses within a trial. On some trials, it gets fed at a hopper that pays off at a relatively short latency after trial initiation (hereafter, the short hopper); on trials when the short hopper fails to pay off at the expected latency, the subject must switch from it to a hopper that pays off after a longer latency (hereafter, the long hopper). To get a pellet on every trial, the subject must time its switches so that they fall between the short and long pay-off latencies. We refer to these

payoff latencies as the temporal goalposts. Switching prematurely to the long hopper on a short trial costs the subject its pellet, as does switching too late on a long trial.

Control of Precision

The first experiment asks whether mice can alter the trial-to-trial variability in the timing of their switches in response to the narrowing and re-widening of the temporal goalposts.

Response timing precision is the inverse of the coefficient of variation in the distribution of timed switches (σ/u , hereafter the cv). This ratio, which is constant over a wide range of interval durations, is also known as the Weber fraction, because it may be thought of as the percent difference in duration at which two intervals may be reliably discriminated.

In most extant theories of timing, the cv cannot be adjusted in response to changing task demands because the source of the variability in the timing of responses is noise intrinsic to the process that times an elapsing interval. The variable results of a noisy timing process determine the variability in the timing of the observed responses.

Scalar variability in the timer (that is, a constant cv) is a postulate in scalar expectancy theory (Gibbon, 1977; hereafter SET). The appeal of the behavioral theory of timing (Killeen & Fetterman, 1988; hereafter BeT), an influential alternative to SET, rested in part on the fact that it derived the scalar variability in the timing of responses from its assumption about the nature of the response timing process. The essence of this assumption—that the experienced reinforcement rate scales the rate at which a Poisson pulse generator drives a sequence of states to a final state—is also a feature of some contemporary neurally oriented timing theories (for example, Simen, Balci, deSouza, Cohen, & Holmes, 2011). Other neurally oriented theories attribute the variability to variation in the intrinsic dynamics of neural processes (Fiala, Grossberg, & Bullock, 1996; Grossberg & Schmajuk, 1989; Karmarkar & Buonomano, 2007). In none of these theories is it apparent how the Weber fraction in the timing of a response could be adjusted. We show in our first experiment that mice do adjust the variability in the timing of their switch responses when we narrow and re-widen the temporal goalposts.

Distinguishing Endogenous from Exogenous Variability

The second experiment asks whether large trial-by-trial variation in the intervals being timed increases variation in the timing of the animal's responses. In the above-mentioned timing theories, the variability in the distributions of response latencies is jointly determined by the variability intrinsic to the timing mechanism (the noise in the internal timing process) and by the variability in the timed intervals (noise in the world). In most timing experiments, the interval(s) being timed do not vary randomly from trial to trial. Therefore, attention is rarely called to the impact that exogenous variability must have on the results of a timing process. The joint dependence on internal and external variability is inherent in any model in which the duration of an externally programmed interval determines the extent to which an internal process progresses. The variability in the externally programmed intervals is independent of the variability in the rate at which the internal process progresses. Therefore, the two variances must sum to determine the trial-to-trial variance in the state attained by the internal

process at the moment of reinforcement. In most theories, the trial-by-trial variability in the attained state—the end result of the timing process—will be reflected in the variability in the timing of responses on subsequent trials.

The mathematical fact that the variance in the recorded latencies (or terminal states) must be the sum of the variance in the timer and the variance in the intervals timed is most readily seen in SET: The internal process is a timer analogous to a stop watch. The rate at which the watch runs is assumed to vary from trial to trial. The externally programmed reinforcement latency stops the watch. (There is an analogous assumption about the process-stopping effect of reinforcement in all the other theories.) In SET, the reading on the watch when it is stopped by reinforcement is recorded in memory. Thus, what is recorded on any given trial depends jointly on how fast the watch ran and on the duration of the reinforcement latency. These two sources of variation—measurement error and external variation—are independent. Therefore, the variance in the intervals recorded in memory (or more generally the state of the internal process when reinforcement is delivered) must be the sum of the internal variance (error variance) and the external variance (variation in the reinforcement latencies). In SET, the target response time on a subsequent trial is determined by a random draw from the store of remembered reinforcement latencies. Therefore, the variance in the remembered latencies translates directly into variance in the distribution of response latencies.

The above considerations might seem to imply that the variance in the distribution of timed responses *must* reflect the variability in the timed intervals. How could variation in the results of a timing operation not reflect variation in the intervals being timed? And, how could such variation in the results of timing not be manifest in the timing of responses based on those results? An answer to the latter question is that variation in the timing of responses need not reflect variation in the timing results if the computations conducted on recorded latencies distinguish between internal variance and external variance. In our second experiment, we show that large random variability in the short and long feed latencies does not increase the variability in the timing of the switch response; if anything, it reduces it. We argue that this result implies that the subjects distinguish measurement error (endogenous variability) from exogenous variability. That in turn implies that more than one timer times each individual feed latency, because only that would enable an estimate of measurement error.

How Rich is the Record?

Lurking behind the questions posed in the first two experiments is the broader question of whether a theory of timing must assume that the animal has a record in memory of its individual experiences (record-based theories) or whether response timing can be explained without such an assumption (record-free theories).

The poles of this contrast between record-based and record-free theories were staked out by SET and BeT. SET is a rich-record theory: it assumes that the animal remembers many individual reinforcement latencies and conducts computations on them. Skinner (1974, 1977, 1990) famously argued that memory had no place in a scientific psychology. The

original version of BeT was put forward in that behaviorist spirit. In the original BeT (Killeen & Fetterman, 1988), there is no timer, just a progression through a fixed sequence of behavioral states. Nor is there record of average reinforcement latency, let alone a trial-by-trial record; rather, the terminal behavioral state (the state when reinforcement is delivered) becomes a discriminative stimulus for the reinforced response.

Timer-free and record-free associative theories of timing continue to find favor in some neurobiological circles (see for example Karmarkar & Buonomano, 2007). Many neurally oriented timing theories are not entirely record-free, but none of them follows SET in assuming a readable record of the sequence of experienced durations. The third experiment asks whether the animal has the trial-by-trial record of past intervals postulated by SET and, more recently, by Wilkes and Gallistel's Analytic Theory of Associative Learning (TATAL, Gallistel & Wilkes, 2016; Wilkes & Gallistel, 2017).

Records of the feeding latencies on individual trials, if they exist, would constitute episodic memories. It was long thought that episodic memory was unique to humans, but there is now accumulating evidence for it in rodents (Crystal & Smith, 2014; Panoz-Brown et al., 2016; Wilson, Mattell, & Crystal, 2015; Zhou, Hohmann, & Crystal, 2012)

In previously published work (Balci et al., 2009), we showed that the mean of the distribution of switch latencies varies in the intuitively expected way with the relative frequency of the short and long trials: When short trials are much more frequent than long trials, the expected cost of a premature switch is greater than the expected cost of a late switch. When long trials are much more frequent than short trials, the reverse is true. Therefore, one might expect that subjects would move their switch-latency distribution away from the short feed latency toward the long feed latency when short trials are more frequent and away from the long feed latency toward the short feed latency when long trials are more frequent. Both mouse and human subjects do make these distributional shifts (Balci et al., 2009). When there is a step change from one condition (short trials more frequent) to the other (long trials more frequent), the shift in the distribution of response latencies occurs abruptly—in a single step—after relatively few trials (Kheifets & Gallistel, 2012). Just as importantly, it often occurs before the subject has missed a single pellet (Kheifets & Gallistel, 2012). This last finding implies that the repositioning of the switch-latency distribution is not driven by differential reinforcement and nonreinforcement of the short and long latencies.

The third experiment pushes the rich-record-versus-record-free question in theories of timing further by asking whether the latency to the abrupt adjustment in the mean of the response-time distribution is sensitive to an information-theoretic difference that can, we believe, be appreciated only by a system that keeps a record of the sequence of trials.

Except for SET, the timing theories known to us are associative theories. Like most associative theories, they generally employ delta-rule updating. The target switch latency is not extracted from a trial-by-trial record of the reinforcement latencies, as in SET. Rather, it is a running average of those latencies, as is associative strength in the theory of Rescorla and Wagner (1972). The running average is readable by a computational process that

computes on each new trial the arithmetic difference between the current value of the average and the outcome on the new trial and updates the average by adding some fraction of that difference to the slightly discounted current value of the running average (delta-rule updating).

Let p_{new} be the new probability of a positive outcome, an outcome that increases the average. Let p_{old} be the expected probability of that outcome on the basis of a running average over past experience. Positive outcomes depart from expectation by $1 - p_{\text{old}}$ and they will occur with relative frequency p_{new} . Negative outcomes, outcomes that decrease the average, depart from expectation by $-p_{\text{old}}$ and they will occur with relative frequency of $1 - p_{\text{new}}$. Therefore, Δ , the expected increment/decrement, on any given trial is $\Delta = \alpha (p_{\text{new}}(1 - p_{\text{old}}) - p_{\text{old}}(1 - p_{\text{new}}))$. Ignoring the constant of proportionality, α , multiplying through and cancelling quadratic terms gives $p_{\text{new}} - p_{\text{old}}$. Thus, on any trial after a change in p , the expected magnitude of the change in the estimate of p is proportionate to $|p_{\text{new}} - p_{\text{old}}|$ regardless of the direction of the change in p .

Delta-rule updating algorithms are thus not sensitive to the informational asymmetry that arises when the changes are between a midlevel probability ($p = .5$) and a more extreme probability (say, $p = .9$). When recent experience leads you to expect a high probability of one binary outcome (say $p(\text{long}) = .9$) when in fact the probability has now dropped to $.5$, information to the effect that the old value for p is wrong comes in more rapidly than when the transition goes in the opposite direction, from $.5$ to $.9$.

The informational asymmetry is intuitively obvious when one considers an extreme value for p . Suppose you believe all swans are white, $p(\text{white}) = 1.0$ when in fact $p(\text{white}) = .99$. The first black swan you encounter gives unbounded evidence that your belief is wrong. The expected number of trials to the first encounter of a black swan is only 100. Because the strength of the evidence against your belief has an infinite expectation after a finite number of trials, the rate at which evidence of error accumulates is infinite. Now suppose the reverse: you believe one swan in 100 is black when in fact they're all white. Even after you've seen 200 white swans and no black ones, you still have only weak evidence that your belief in $p(\text{white}) = .99$ is wrong. The average rate at which evidence of error accumulates when the direction of error is reversed is very slow.

The Kullback-Leibler divergence measures the informational distance between two probability distributions with the same support, for example, two Bernoulli distributions with different values for p . It gives the average information (information about the error in one's belief) gained per observation (Cover & Thomas, 1991). Intuitively, the more one's belief diverges from the truth, the more rapidly experience makes that apparent. As already explained, the divergence of a Bernoulli distribution with $p = 1.0$ from a Bernoulli distribution with $p = .99$ is infinite, whereas the divergence going the other way is very small (.01 bits per experience).

Optimal algorithms for detecting changes in stochastic parameters take advantage of all the available information, so they are sensitive to informational asymmetries. They detect the change from $p = .9$ down to $p = .5$ more rapidly than the change from $p = .5$ to $p = .9$.

However, the optimal algorithms known to us (for example, Adams & MacKay, 2006) operate on a record of the sequence of outcomes—a record that is not available according to associative theories of timing.

Our third experiment asks whether the mouse is sensitive to this informational asymmetry. In asking this question, we distinguish strongly between theories like SET, which assume a rich record of the sequence of past outcomes, and the vast majority of timing theories, which assume at most a running average (minimal-record or no-record theories). Our argument depends on the assumption that only algorithms operating on the remembered sequence of outcomes can show this asymmetry.

General Method

The experimental environment for the experiments with mice is schematized in Figure 1. It consisted of a polypropylene nesting tub connected by an acrylic tube to a Med Associates mouse test chamber. Three hoppers were ranged along the far wall of the test chamber. The interior of each hopper could be illuminated by an LED. The entrance to each was monitored by an infrared beam. Pellet dispensers attached to the flanking hoppers could deliver 20 mg food pellets into those hoppers. The middle hopper served only to initiate a trial. The mice lived in the environment 24/7. A bright LED illuminated the test chamber from 22:00 (artificial dawn) to 10:00 (artificial dusk), providing a reverse 12:12 day-night cycle. There were two feeding phases every 24 hr, during which the mice obtained pellets by performing the tasks described below. These feeding phases straddled the artificial dusk (test chamber light off) and dawn (test chamber light on), one from 9:00 to 13:00 and the other from 20:00 to 24:00.

In the switch paradigm, the illumination of the middle hopper signals that the subject may start a trial whenever it chooses by poking into the illuminated Trial-Initiation hopper. The interruption of the Trial-Initiation hopper's infrared beam extinguishes the illumination in that (the middle) hopper and illuminates the flanking hoppers. On any one trial, the controlling computer silently chooses at random one of the flanking hoppers to deliver a pellet. The subject may obtain a pellet from the chosen hopper by appropriately timing its visit(s) to the two illuminated hoppers.

When the computer chooses the short-latency hopper, then the subject's first poke into that hopper at or after the expiration of a 4s latency releases a pellet, unless the subject makes one of two possible errors: 1) If it goes directly to the long-latency hopper after trial initiation, the trial terminates immediately without a pellet release. 2) If it leaves the short hopper for the long hopper before the short latency has elapsed, then when the short latency elapses, the first poke into the long hopper terminates the trial without a pellet release. On rare occasions—less than 0.1% of the trials—the mouse may leave the short hopper after a brief visit, visit the long hopper briefly, but return to the short hopper so quickly that the first poke made when the short latency elapses is into the short hopper; in those rare cases, it gets its pellet.

If the computer chooses the long-latency hopper, the subject earns a pellet by going first to the short hopper and then leaving it for the long hopper before the long latency has elapsed. Again, there are two possible errors: 1) Going first to the long hopper; in this case, the trial terminates immediately without a pellet release. 2) Continuing to poke into the short hopper until after the long latency has elapsed; in this case, when the subject finally does switch to the long hopper, the first poke there terminates the trial without a pellet release. What matters to the outcome on a long trial is only that the mouse goes first to the short hopper and that the first poke at or following the elapse of the long latency is into the long hopper. If the mouse leaves the short hopper for the long hopper, then returns briefly to the short hopper before revisiting the long, it gets its pellet provided its first poke at or after the long latency elapses is into the long hopper. (Such trials occur, but they are rare.)

A trial ends with the turning off of the lights in the flanking hoppers. The turning off of these two hopper lights is usually coincident with the release of a pellet into one or the other hopper, except on error trials, when they turn off without a pellet release. Following trial termination, the middle (Trial-Initiation) hopper re-illuminates after an interval drawn from an exponential distribution with an expectation several times longer than the within-trial feed latencies. For more details on equipment and method, see (Gallistel, Balci, Freestone, Kheifets, & King, 2014).

In summary, the subject gets a pellet on every trial provided it always goes first to the short hopper and, on long trials, its first poke at or after the elapse of the long latency is into the long hopper. Switch latencies are calculated only on the long trials. The calculated latency is the latency at which the subject stops poking into the short hopper, that is, the latency to the termination of the last poke into the short hopper prior to the first poke into the long hopper.

Experiment 1. Narrowing and Widening the Temporal Goalposts

The subjects were eight CD-1 male mice (Harlan) approximately 50 days old at the start of the experiment. The intertrial intervals in the switch conditions were exponentially distributed with an expectation of 80 s. The short feed latency in the switch conditions was fixed at 4 s, measured from the onset of the first poke into the short hopper. The long feed latency, which was also measured from the onset of the first poke into the short hopper, was initially 12 s. It was automatically shortened to 8 s by the quasi real-time data-analyzing software when that software detected at least 800 measured switch latencies in the initial switch condition. It was advanced from the 4–8 condition to the final 4–12 condition when the software detected at least 600 measured switch latencies in the 4–8 condition. Because the automated quasi real-time data analysis occurred only twice per day, the numbers of recorded switch latencies when the program actually adjusted the long latency were substantially greater than these minima. The number of trials in each condition for each subject may be read off the vertical solid lines in Figure 4. The number of recorded switches for a given condition is in each case approximately half the number of trials.

To accustom the mice to the test environment, we began with a matching protocol, in which pokes into the unilluminated flanking hoppers were rewarded on concurrent variable schedules of reinforcement (VI 160 s VI 480 s). When the mouse had completed at least 200

cycles between the two hoppers, the VI schedules were reversed. Each mouse completed at least 200 cycles in this reversed VI condition before passing into the two-hopper autoshaping phase.

The purpose of the autoshaping phase was to teach the mouse the 4-s and 12-s feeding latencies. In this phase, the illumination of the Trial-Initiation hopper signaled that the mouse could initiate a trial by poking into it. The interruption of that hopper's infrared beam immediately extinguished the light in that hopper and illuminated a randomly chosen one of the two flanking hoppers. On trials when the left hopper was illuminated, a pellet was released into it at the end of 4 s, whether the mouse poked or not; when the right hopper was illuminated, a pellet was released into it at the end of 12 s whether the mouse poked or not. As in the ensuing switch phase, the intertrial intervals were exponentially distributed with a mean of 80 s. This autoshaping phase continued until the mouse reliably poked more rapidly into each hopper when it was illuminated (in learned anticipation of imminent food delivery) than during the intertrial intervals when it was not illuminated (and food delivery was not imminent). Because the trials to the acquisition of anticipatory poking differed from mouse to mouse, the number of autoshaping trials varied from slightly less than 100 to slightly more than 500, before the mouse advanced to the first of the switch conditions.

Statistical Analyses

Fitting different distributions to the switch-latency data—Because the form of the mathematical distribution that best describes the empirical distribution of these latencies is of theoretical interest, we fitted three distributional forms to the switch latencies in each condition: the normal, the gamma and the Wald (also known as the inverse Gaussian). In making these fits, setting an outlier criterion was essential, because there were extreme outliers that grossly distorted the estimates of, for example, the mean and that rendered all fits obviously poor.

Because the three distributional forms differ in their ability to accommodate a long right tail, the choice of an outlier criterion might have affected the conclusions regarding which forms provide a better description. To determine if that was the case, we used three different outlier criteria: 12, 15 and 20 s. We used only the latencies from the second half of the first condition, because it was apparent that the latencies were stabilizing over the first half for some subjects.

The comparisons of principal interest are for the between-condition effects on the coefficient of variation (cv) of the switch latency distributions. Of secondary interest are the effects on the means; these were assessed by *t* tests.

Computing the strength of the evidence for changes in the CVs—The cv is the ratio of the estimates for two distribution parameters, the standard deviation and the mean. How to compute *p* values for the differences in cvs is unclear; whereas the computation of Bayes Factors for the cv comparisons is straightforward. Like a *p* value, a Bayes Factor gives the strength of the evidence in favor of a between-condition difference when the data suggest it. Unlike a *p* value, it also gives the strength of the evidence in favor of no difference when the data suggest no difference (Gallistel, 2009). A Bayes Factor of 3 is

commonly given the same interpretation as a p value of .05; it indicates substantial evidence. A Bayes Factor of 10 corresponds roughly in its interpretive implications to a p value of .01 (strong evidence), while a Bayes Factor greater than 100 is regarded as “decisive.” Bayes Factors less than 1 support the null hypothesis (no difference between the conditions being compared). The reciprocal of a Bayes Factor that is smaller than 1 is the Bayes Factor (odds ratio) in favor of the conclusion that there is no difference.

To obtain a Bayes Factor for each cross-condition comparison in each mouse, we computed the likelihood functions for normal distributions parameterized by their means and cvs, given the switch latencies from two conditions, after discarding outliers (lower panel of Fig. 2). We then summed over the values for the mean to obtain the marginal likelihood functions for the cvs (upper panel of Fig. 2). Normalizing the marginal likelihood function for the cv from the earlier of the two conditions (scaling it so that it integrates to 1), gives the posterior probability distribution for that cv. This distribution is the null prior on the marginal likelihood function for the later condition in the comparison; it predicts where the later likelihood function should fall if the change in condition does not change the cv.

A two-tailed Bayes Factor is computed by constructing an alternative prior distribution to represent the hypothesis that the cv changed. The alternative prior specifies the range on either side of the mean prechange cv within which the postchange likelihood function could plausibly fall. We get this range of plausible effect sizes from Figure 7, which shows that the observed between-condition differences in the cvs ranged from -0.1 to $+0.1$. The two priors and the likelihood function are each represented by vectors specifying the probability densities (in the case of the priors) or likelihoods (in the case of the likelihood function) corresponding to each element in the vector of the plausible cv values that constitutes the support for these functions. The posterior marginal likelihoods of the competing hypotheses (the competing priors) are obtained by weighting each point in a given prior by the corresponding value in the likelihood function and summing across the results. The Bayes Factor is the ratio of the two posterior marginal likelihoods (Fig. 3, see Gallistel, 2009, for more details).

For reasons already mentioned, we excluded from these analyses latencies greater than 12 s (except for S3, for which the exclusion threshold was 18 s), and we used only the latencies from the second half of the first condition.

Results

Independent effects on mean and precision—Figure 4 is a scatter plot of each subject’s switch latencies throughout the three conditions. Figure 5 gives the condition-by-condition cumulative distributions of the switch latencies. Figure 6 gives the condition-by-condition cumulative distributions of the censored switch latencies used to compute and compare the cvs.

The responses to the changes in the long feed latency varied strikingly between subjects. However, in seven of the eight subjects, one or both of these changes—from the initial 4–12 condition to the 4–8 condition and then back to the 4–12 condition—elicited an unequivocal change in the cv. In what follows, we call attention to the various ways in which the subjects

responded to the changes in the long feed latency. A theoretically important feature of these responses is their immediacy and their abruptness. The changes are apparent in the points immediately following the vertical lines across the scatter plots in Figure 4, which indicate the points at which the long feed latency changed.

In response to the lowering of the long feed latency, S1 abruptly lowered its mean switch latency and tightened the distribution (reduced the scatter), thereby keeping the great majority of its switches between the temporal goalposts. When the long feed latency was increased back to 12s, it lowered its mean still further and tightened the distribution still further.

The scatter plot of successive switch latencies shows the abruptness and immediacy with which S1 responded to the shortening of the long feed latency. However, it does not enable us to judge whether the lowering of the mean and the reducing of the scatter altered the cv. For that, we turn to the semilog cumulative distributions in Figure 6. When cumulative distributions are plotted against a logarithmic x-axis, a change in the cv is usually manifest as a change in the steepness. Put another way, if scalar variability holds—if the cv does not change because the standard deviation and the mean change by the same factor—then the two empirical cumulative distributions will be parallel when plotted over a common logarithmic x-axis.

The shifts in the means are apparent in the cumulative distributions for S1 in Figures 5 and 6: the dashed plot for the 4–8 condition lies well to the left of the thin solid curve for the initial 4–12 condition, and the heavy solid curve for the final 4–12 condition lies still further to the left. On Figure 6, we read that the mean in the 4–8 condition was 85% of the mean in the initial 4–12 condition and the mean in the final 4–12 condition was 92% of that reduced mean. These shifts in the mean are highly significant ($p < .001$ by the 2-sample Kolmogorov-Smirnov test).

The standard deviations of the latter two distributions are also smaller, as is evident in the scatter plot in Figure 4. That is to be expected given the well-established scalar variability in response timing distributions. The question is whether the factors by which the standard deviations diminish are the same as the factors by which the means diminish, as predicted by scalar variability. If so, then the ratios between the cvs for the various conditions would be 1. However, we see that the cumulative distributions in the 4–8 and second 4–12 condition are steeper than in the first 4–12 condition. This implies that the cvs were reduced, which indeed they were: The cvs in the 4–8 condition and second 4–12 condition were 82% and 83%, respectively, of the cv in (the second half of) the initial 4–12 condition. The Bayes Factors on the top of the scatter plot for S1 in Figure 4 show that the evidence for these reductions in the cv is beyond decisive (Bayes Factors of >1000 and 660, respectively). Thus, S1 substantially reduced its cv—increased the precision of its response timing—in response to the narrowing of the goalposts. The effect was immediate and abrupt.

S2, by contrast, lowered its mean switch latency by 4% while *increasing* its cv by 11% in response to the reduction in the long feed latency. The reduction in the mean was highly significant ($p < .001$ by a two-tailed t test), but the increase in the cv was not. Particularly

striking in the scatter plot for S2 is the immediate and abrupt 24% increase in the mean and 18% increase in the cv when the long feed latency was lengthened back to 12 s. The Bayes Factor of 17 for the comparison between the first and final 4–12 condition tells us that there is strong evidence for an increased cv between the first and the final condition. This replicates the finding that subjects can control their cv and it establishes the important fact that a subject may change its cv in either direction in response to the same challenge.

S3 showed the most unexpected response to the shortening and lengthening of the long feed latency: instead of lowering its mean and tightening the dispersion, it immediately increased both. In consequence, the great majority of its switch latencies in the 4–8 condition were too long (> 8 s), with the result that it failed to get a pellet on most long trials throughout this condition. The standard deviation increased by more than the mean but not by a factor sufficient to suggest evidence of a change in the cv. The subsequent restoration of the long feed latency to 12 s, however, did produce strong evidence of independent control of the cv, because, paradoxically, in response to the lengthening of the long feed latency, the subject reduced the mean but not the dispersion, with the consequence that the cv in the final 4–12 condition was decisively greater than in the initial 4–12 condition (Bayes Factor > 1000). Thus, the cv may increase when the mean decreases. The Bayes Factor of greater than 1000, which unequivocally confirms increase in the cv, is surprising given that the two cumulative distribution plots in Figure 6 are so close. The difference in the cv arises because the second 4–12 distribution is leptokurtic (has long tails). The leptokurtosis is not readily apparent in Figure 6. However, in the scatter plot in Figure 4 one sees more unusually long switches in the second 4–12 condition than in the second half of the first 4–12 condition and more unusually short switches as well. These are what give the distribution its long tails.

S4 had a low mean during the second half of the initial 4–12 condition. Like S3, it unexpectedly increased substantially its mean switch latency in response to the shortening of the long feed latency (by 28%). Unlike S3 however, it radically reduced the dispersion, with the result that a 28% increase in the mean coincided with a decisive 23% decrease in the cv. Thus, a cv may decrease when a mean increases. Closer scrutiny of the scatter plot for S4 reveals a dramatic variation of the cv within the 4–8 condition: The immediate response to the shortening of the long feed latency was a clear *increase* in the dispersion, but this was followed by a very strong reduction in the dispersion during the last two thirds of the trials.

S5 responded to the shortening of the long feed latency with an immediate 14% reduction in its mean switch latency and a proportional reduction in the dispersion, hence no change in the cv. However, it responded to the subsequent increase in the long feed latency with a further reduction in the mean but a slight increase in the standard deviation, hence, a decisive *increase* in the cv. The increased cv persisted when the long feed latency was restored to 12 s. Notice that this restoration elicited several unusually long switches.

S6 responded to the shortening of the long feed latency with an immediate reduction in the mean and an almost commensurate reduction in the dispersion, giving, therefore, only weak evidence for a change in cv. Remarkably, however, it responded to the subsequent increase in the long feed latency with an immediate further *reduction* in its mean switch latency and disproportionately greater reduction in the dispersion, hence, a decisive *reduction* in its cv.

S7 was the only one of the eight subjects that did not change the precision of its timing across the three conditions. It changed both the means and the dispersions of its switch latencies by appropriate and proportional factors. The Bayes Factors for the three between-condition comparisons, that is, the reciprocals of the values given above the plot in Figure 4, tell us that the evidence for no between-condition changes in cv in this subject ranges from substantial to strong.

S8 showed an immediate but also continuing reduction in mean and dispersion in response to the reduction in the long feed latency. Taking the condition as a whole, the mean in the 4–8 condition was reduced by 23% while the cv increased by 25%. The subsequent lengthening of the long feed latency produced a slight further reduction in the mean and greatly reduced the dispersion, with consequently strong evidence of a decrease in the cv (by 12%). However, inspection of the scatter plot shows that these values are misleading because of the continuing change during the 4–8 condition. The mean of the switch latency distribution in the last third of the 4–8 condition is 6% *smaller* than the mean in the subsequent 4–12 condition, and the standard deviation is 4% *smaller*. When the last third of the 4–8 condition is compared to the subsequent 4–12 condition, the Bayes Factor of 0.14 (not shown in Fig. 4) favors the conclusion that the cv did not change. When this last third is compared to (the second half of) the first 4–12 condition, there is negligible evidence for a change in the cv (Bayes Factor = 1.1, also not shown in Fig. 4).

In summary, the cv results show that a decrease in the mean of a switch-latency distribution may accompany either a decrease or an increase in the cv, and likewise for an increase in the mean. Neither the direction nor the magnitude of the change in the mean predicts the direction or magnitude of the change in the cv (Fig. 7). For the most part, the effects of a change in condition on both the mean and the cv are striking for their immediacy and their abruptness, although there is also occasional evidence for a more or less continuous change in both parameters (see the 4–8 condition for S8 in Fig. 4).

The form of the switch-latency distribution—Figure 8 gives the cumulative distributions of the likelihood-maximizing parameters for the three distributional forms of theoretical interest—the normal, the gamma and the Wald—for the three conditions. Figure 9 gives the cumulative distributions of the maximum likelihoods; the greater the maximum likelihood, the better the fit.

The results in Figures 8 and 9 imply that the Wald distribution is not well suited to describing switch latencies. The values for the Wald lambda parameter range over more than two orders of magnitude (from 2 to 368, see Fig. 8), even though the empirical distributions look generally similar (Figs. 5 and 6). By contrast, the parameters of the normal and the gamma vary over only about a two-fold range. Secondly, the likelihood of the best Wald fit is often substantially worse than that for the other two forms (Fig. 9). The choice of an outlier criterion does not affect this conclusion as to how well the different forms describe the empirical distributions (compare across columns in Fig. 9).

The normal and the gamma distributions describe the latency distributions equally well (Fig. 9). However, this is because the values for the shape parameter of the gamma are generally

high (Fig. 8). The higher the value of this parameter, the more nearly the gamma approximates a normal distribution. Similarly, for the Wald distribution, the higher the value for lambda, the more nearly the distribution approximates a normal distribution. Only the high values for lambda gave reasonable approximations to the empirical distribution. The Wald fits with low values of lambda were grossly erroneous by simple inspection. In short, the empirical distributions of switch latencies are approximately normal and not well described by the Wald distribution.

Discussion

The results of Experiment 1 show that the variability in the switch latencies is not a product of noise processes over which the mouse has no control. Within limits, the precision with which switches are timed is adjusted in response to changing conditions. As is evident in Figure 4, the change in precision in response to a change in conditions is often both rapid—it occurs soon after the change in the long feed latency—and abrupt—it occurs over the span of very few switch trials.

The mean and precision of switch latencies vary independently: Decreases in the mean may occur together with either increases or decreases in the cv, and likewise for increases in the mean (Fig. 7). One can sometimes see the two parameters varying independently within a condition. Subject 4, for example, initially responded to the shortening of the long feed latency with an increase in the variability of its switch latencies, but then greatly reduced that variability, so that the net change in cv taking the conditions as a whole was unequivocally downward.

In the General Discussion we consider the implications of these results for contemporary theories of response timing. Here, we briefly mention two theoretically important findings regarding the form of the switch latency distribution. Current theories differ with regard to the mathematical form that they suggest should best describe an empirical distribution of response latencies. Some suggest that it should be normal (Gallistel & Wilkes, 2016), others that it should be gamma (Killeen & Fetterman, 1988), and still others that it should be Wald (Simen et al., 2011). Our results imply that the Wald distribution is not well suited to describing the empirical distribution of switch latencies. They also imply that the gamma is suited only if its shape parameter is approximately 20. Values for the shape parameter that high make the gamma a good approximation to the normal. Thus, these results imply that to describe the empirical distribution of switch latencies one must use a distribution that is a good approximation to the normal.

In the original version of the Behavioral Theory of Timing (BeT, Killeen & Fetterman, 1988), the shape parameter of the gamma distribution is n , the number of behavioral states the animal passes through before reaching the state that serves as a discriminative stimulus for the emission of a timed response. The implication of these results for BeT is that to explain the shape of the switch-latency distribution, one must assume that the mouse passes through many different behavioral states in a short amount of time. Taking a typical switch latency to be 6.5 s and the typical value of n to be 20, the mouse must pass through about three behavioral states per second.

Experiment 2. Jittering the Goalposts

The changes in the dispersion of switch latencies in response to changes in the long feed latency show that the variability in switch latencies must depend on the subject's timing of the feed latencies. All theories of timing known to us—including no-timer theories—have the property that the variability in the observed response latencies must depend in some appreciable measure on variability in the timed intervals themselves. As explained in the introduction, the objective durations of the experienced intervals determine the how long the internal process runs, hence the terminal state it attains. The terminal states attained by the internal process over successive trials determine somehow the subsequently observed response latencies.

As already noted, most associative models of timing assume delta-rule updating. A delta-rule updating model convolves a geometrically decaying kernel with the sequence of experienced intervals to obtain a running average of those intervals. Convolution is a linear operator; doubling the variability in the sequence of intervals doubles the variability in the trial-by-trial values of the running average. If the value of this running average on a given trial determines the subject's response latency on that trial, then increased trial-by-trial variation in the running average will appear as increased trial-by-trial variability in response latencies.

Perhaps most importantly, no experiment known to us addresses the question whether random trial-by-trial variation in reinforcement latencies about a nonzero mode is in fact manifest in increased variation in response timing that is centered on that mode. Our second experiment poses that question.

Experiment 2A: Mouse Subjects

Method—As in the previous experiment, we began with the matching and two-hopper autoshaping protocols, followed by the switch protocol, training the subjects initially with fixed short and long latencies of 3 s and 9 s, respectively. At the beginning of the sequence of sessions in which the location and magnitude of the trial-to-trial variation in feed latencies began to be changed, the nominal values of the two feed latencies were set to 4 s and 8 s. The relative frequencies of short and long trials were kept equal (50% short 50% long). The intertrial intervals were exponentially distributed with an expectation of 240 s. The subjects were 29 male mice of the C57BL/6 strain. They were approximately 50 days old at the start of training. Five of the initial subjects were excluded because they failed to reach clear proficiency on the switch task.

The switch protocol in this experiment differed from that in the previous experiment in the way in which overly long stays at the short hopper on long trials were treated. In this experiment, a short trial ended without reinforcement with the first poke made into the short hopper after the long feed latency had elapsed. No switch latency could be calculated on trials thus terminated. In the previous protocol, the trial ended without reinforcement only when the mouse eventually switched to the long hopper; a switch latency could therefore be calculated on those trials. Also, in the previous experiment, the feed latencies were timed

from the first poke into the short hopper. In this experiment, they were timed from the illumination of the two hoppers.

With the nominal values for the two feed latencies fixed at 4 s and 8 s, we varied the coefficient of trial-to-trial variation in one or both feed-latencies, using cv values of 0.1, 0.2 and 0.35. Figure 10 shows these distributions for each of the three possible loci of variation (only the short feed latency jittered, only the long, or both). Adding jitter to the feed latencies made it impossible for the mouse to obtain a pellet on every trial because the actual feed latencies on some nominally long trials were shorter than the actual feed latencies on some nominally short trials.

Crossing the three cv values with the three loci of variation yields nine conditions. Testing every possible ordering of the nine conditions was not feasible. However, some subjects experienced different orders for the location of the jitter while the cv remained fixed. For example, one such subject experienced cv = 0.1 jitter in first the short feed latency, then in the long feed latency, and then in both; whereas another such subject experienced first both jittering, then only the long, then only the short. For other subjects, the location of the jitter (short, long or both) was fixed, but they encountered different levels of jitter and in different orders. For example, for one such subject the jitter was always in the short feed latency, but the level of jitter increased from cv = 0.1 to 0.2 to 0.35; whereas another such subject experienced the same levels of jitter in the same short feed latency but in reverse order. On analyzing the data, we found no significant effect of order, so we collapse the different orders in presenting the results.

In analyzing these data, we included all recorded switch latencies. However, we stress that it is impossible to calculate a switch latency longer than whatever the computer-chosen long feed latency was on a given long trial with this variant of the switch protocol. When the mouse stayed at the short hopper longer than whatever that chosen long latency was on a given “long” trial, then the trial ended without reinforcement the moment that long chosen latency was exceeded (and no switch latency was calculated). This was true even when the computer-chosen “long” latency was shorter than the nominal (i.e., mean) feed latency on “short” trials.¹

Results—The jittering of the goalposts introduced large variability in the short feed latencies, the long feed latencies or both (Fig. 11). Contrary to our expectations, however, adding variability to the intervals being timed increased the precision with which the mice timed their switches (reduced the cvs) rather than decreased it. In the panels of Figure 12, the cumulative distributions for the conditions in which one or both feeding latencies were jittered lie significantly to the left of the cumulative distribution for the conditions in which the programmed feeding latencies had no extrinsic noise in them. (In every such comparison, $p < .01$ by the two-sample Kolmogorov Smirnov test.) The failure of the variability in switch latencies to increase in response to greatly increased trial-to-trial variability in the reinforcement latencies implies that mice distinguish between their endogenous interval-

¹‘Long’ and ‘short’ are in scare quotes because the jitter blurred the distinction between them in this experiment; the distinction was more in the mind of the experimenters than in the experience of the mouse.

measurement error and exogenous variability in the intervals timed. Doing so enables them to isolate the variability in their response timing from the variability in the timed intervals themselves.

Experiment 2B: Rat Subjects

Method—This version of the experiment was done by DF in the laboratory of Russell Church at Brown University. The subjects were 36 naïve male Sprague-Dawley rats (Taconic Laboratories, Germantown, NY). They were kept in a colony room on a 12:12 light-dark cycle (lights off at 8:30am). Dim red lights provided illumination in the colony room and testing rooms. Upon arrival, the rats were 8 weeks of age and weighed between 75 and 100 grams. During the first week, the rats were on a free-feeding schedule. After a week, their daily food (FormuLab 5008) was rationed to 16 grams per day. During the experimental session, the rats were fed 45-mg Noyes pellets (Improved Formula A) as a reward. Water was available ad libitum in both the home cage and the testing chamber.

Twenty-four experimental chambers (Med Associates, dimensions $25 \times 30 \times 30$ cm) were in two separate experiment rooms (12 in each room). Each chamber was contained in a sound-attenuating box (Med Associates, dimensions $74 \times 38 \times 60$ cm) with a fan for ventilation. Each experimental chamber was equipped with a pellet dispenser (Med Associates, ENV-203) on the front wall that delivered the reward into a food cup. A head entry into this cup interrupted a photo beam (Med Associates, ENV-254). On both sides of the food cup, there were two retractable levers. On the opposite wall, a water bottle protruded into the chamber allowing ad libitum access to water during the session. A lick on the spout of the water bottle completed an electric circuit. A Gateway OptiPlex 380 computer running Med-PC for Windows (version 4.1) controlled the experiments and recorded the data. The interruption of the photo beam and the completion of the lick and lever circuits were recorded in time-event format with 10-ms accuracy.

The rats were trained to press a lever in the first 2 days (FT-60 until the first press, then FR-1 for a total of 30 deliveries per lever). On days 3–8 we ran the switch task, with a 3 s short feed latency and a 6 s long feed latency and an exponentially distributed intertrial interval with an expectation of 90 s. However, many of the rats perseverated, so Days 9–15 included correction trials in which incorrect choices led to the same trial type on the next trial.

Days 16–30 were baseline sessions in which the external variability was not added to any part of the task. On days 31–40, we added $cv = 0.33$ variability. On days 41–50, we added $cv = 0.66$ variability. In a between-subject design, we added jitter to the short trial durations (12 rats), the long trial durations (12 rats), and the switch times (12 rats). Thus, the rat version of this experiment explored the effect of more extreme exogenous variability than tested in the mouse version. When the cv is 0.66, approximately 6% of the jittered times are less than 0. Whenever the draw from the jitter distribution yielded a reinforcement latency less than 0, the latency was set to 0.02 s, which is the smallest latency at which MedPC could arm a lever.

On a short trial, a pellet was released and the trial ended if and when the programmed short feed latency elapsed, *provided* that the short lever was already depressed at that moment or

was the first lever to be depressed after that moment. If, however, the short latency elapsed and the first lever depression at or after that moment was on the long lever, the press on the long lever ended the trial without a pellet being delivered. On a long trial, a pellet was released at or after the long feed latency elapsed, *provided* that the long lever was already depressed at that moment or was the first to be depressed after that moment. The recorded switch latency was the latency of the most recent release of the short lever. If, on a long trial, the last release of the short lever occurred after the long feed latency had elapsed, then when the rat eventually left the short lever and depressed the long lever, the trial ended without a pellet release. The recorded switch latency was the latency of the final release of the short lever.

This implementation of the switch paradigm made it possible to add scalar jitter to the switch latencies themselves in one of the three groups (rather than to the goalposts). The hope was to mimic the effect of an increase in the variability of the rat's internal timer. At each release of the short lever that was followed by a depression of the long lever—indicating that a switch had occurred—the computer drew a value from a Gaussian distribution whose mean was the trial time elapsed at the moment of the short lever was released and whose standard deviation was the programmed cv times that mean. This jitter value was added to the time at which the short lever was released. The jittered short-lever-release time thus obtained was compared to the short and long feed latencies. The principal effect of this jitter was felt on long trials. Suppose the rat made its final release of the short lever before the long feed latency elapsed, but the jittered value for that release time fell beyond the long feed latency, then when the rat depressed the long lever the trial ended without a pellet, even though the rat had switched (left the short lever) at an appropriate time. However, it was also sometimes felt on a short trial. Suppose there was a release of the short lever before the short feed latency had elapsed followed by a depression of the long lever. In the absence of jitter, this premature departure from the short lever would count as an error and no pellet would be received. However, if the jitter added to the premature release produced a jittered time longer than the short feed latency, it did not count as an error; the rat was reinforced on arriving at the long lever despite having left the short lever prematurely. The purpose of the complex jitter in this group was to mimic an increase in reinforcement variability that derived from behavior (switch times) instead of variability that derived from the trial durations themselves.

We analyzed all the switch latencies obtained. We stress that, like the protocol used with the mice in Experiment 1, but unlike the protocol used in Experiment 2A, the rat protocol allowed us to record switch latencies on trials where the rat stayed too long at the short hopper, that is, longer than the computer-chosen long feed latency. We stress again that, in the jitter conditions, this computer chosen “long” feed latency was often shorter than the nominal (i.e., mean) feed latency at the short hopper—and vice versa, the latency at which poking on the short hopper was reinforced was sometimes longer than the nominal (mean) feed latency at the long hopper.

Results—As in Experiment 2A, jittering either the short and or the long feed latency greatly spread out the feed latencies experienced when pressing the short or long levers, respectively (Fig. 13, panels B and C). The attempt to achieve similar effects by adding jitter

to the short-lever release times was not successful. Jitter was only applied to short-lever releases that occurred after 3 s had elapsed and that were followed by a depression of the long lever, because only at that point was it clear that a switch had been made. In consequence, there was never a pre-3-s release on short trials to which jitter could be applied. On most short trials, the rat received its pellet when the 3 s had elapsed. It sometimes got a pellet on a short trial (on arrival at the long lever) even though it left the short lever prematurely, but these “undeserved” reinforcements, while potentially confusing to the rat, are not evident in the feed-latency quantiles. On most long trials, the rat had depressed the long lever by the time 6 s had elapsed. In that case, it got its pellet provided that its release time on the short lever + the jitter was less than 6 s. It failed to get a pellet when this jittered release time was greater than 6 s. Thus, there were “undeserved” nonreinforcements on long trials. Like the underserved reinforcements on short trials, these are not evident in the feed-latency quantiles.

Jittering the goalposts and/or decreasing the probability of reinforcement on the long lever did not decrease the precision of the rat’s switch times (did not increase the cv). Indeed, the cv decreased slightly, but insignificantly over the 3 phases. Figure 14 is a representative plot of the switch latencies of one rat from each of the three groups over the three phases (no jitter, jitter with a cv of 0.33 and jitter with a cv of 0.66). In every phase, almost all its switches fell between the mean values for the goalposts despite the extreme trial-to-trial variability in the reinforcement latencies from the two levers. Figure 15 is a scatter plot of the cvs when there was jitter (Phases 2 and 3) versus the cvs in the initial phase when there was no jitter. The points scatter around the identity, which means that, generally speaking, the cv was no greater in the jitter conditions than in the initial no-jitter conditions despite the extreme variability in the latencies at which poking was reinforced at one or the other hopper.

Discussion

The results from both versions of this experiment imply that mice and rats do not confound endogenous errors in their measurements of elapsed or elapsing intervals with objective variability in the intervals themselves. As pointed out in both the general introduction and the introduction to this experiment, all timing theories known to us assume that the subject has a target latency that is determined by the results of timing processes of some kind. In all of these theories, even in no-timer theories, the result on any given trial of an internally unfolding “timing” process is jointly determined by the rate at which the process unfolds and by the duration of the external interval whose conclusion terminates the unfolding. Thus, the variance in the trial-by-trial results of this internal process must be the sum of the variance in the rates at which the internal process unfolds and the (scaled) variance in external durations. And, in all of these theories, the variability in the results of timing should be manifest in the variability in the timing of the subject’s responses.

One may imagine a theory in which the variability in response timing is decoupled from the variability in the results of timing operations. Obviously, such a theory must then address the source of the variability in response timing. In doing so, it must explain why the variability in the timing of responses scales with the target value. (Scalar variability holds in the switch

paradigm, just as in other timing paradigms, Fetterman & Killeen, 1995; Gallistel, King, & McDonald, 2004) And, the theory must explain how the subject may make the rapid and abrupt changes in both the mean and the precision of their response timing that were evident in Experiment 1 (see Fig. 4). Most timing theories assume that the timing of the response on the current trial is based on a running average of the reinforcement latencies experienced on preceding trials. As noted in our introduction, running averages are linear operators; increasing the trial-to-trial variability in the inputs by a given factor increases the trial-to-trial variability in the running average by that same factor. The only way to make input variability have a negligible effect on the variability in response timing in these theories is to make the decay rate in the running average very slow, but, in that case, one would not observe the rapid and abrupt adjustments in the mean and precision of response timing seen in Experiment 1.

A subtler problem for a theory that attempts to isolate variation in the target response time from variation in the results of timing processes is how it is possible for the target switch latency not to vary from trial to trial. The target could not be a running average, because increasing the trial-to-trial variability in the input intervals by some factor increases the trial-to-trial variability in a running average by that same factor. A fixed average—an average that does not vary as new data come in—may only be obtained by fixing the sample from which it is computed. Running averages are by definition not based on fixed samples. Fixing a sample would seem to imply that a sequence of results from previous trials was preserved in memory rather than being folded into a running average. Moreover, assuming that the target does not vary trial-by-trial raises the question of how the target may be changed when the mean of the timed intervals changes, as in our first experiment. Addressing that question would seem to require a change-detection algorithm operating on the sequence of timing results preserved in memory (Gallistel, Krishan, Liu, Miller, & Latham, 2014; Gallistel & Wilkes, 2016; Wilkes & Gallistel, 2017). Our third experiment addresses the question whether the subject's ability to change its target switch latency (hence the mean of its switch-latency distribution) depends on a remembered sequence of trial durations rather than only on a running average of those durations.

In the General Discussion, we then return to the question of how it is possible for a subject to isolate variability in its target latency from variability in the results of its timing, while preserving the scalar variability in the distributions of its response times.

Experiment 3. Sensitivity to Informational Asymmetry

The distribution of switch latencies is usually positioned approximately where it should be between the short and long temporal goalposts, given the subjects' uncertainty about the true value of the elapsed interval within a trial. However, the location of the distribution of switch latencies is also sensitive to an exogenous uncertainty of a different kind from the one manipulated in Experiment 2; it is sensitive to the uncertainty arising from the relative frequency (the complementary probabilities) of the randomly programmed short and long trials (Balci et al., 2009). When long trials are more probable than short, the risk of losing a pellet incurred by a tardy switch is greater than the risk incurred by a premature switch, and the distribution of switch latencies shifts toward shorter latencies. When short trials are more

probable than long, the relative risks reverse, and the distribution of switch latencies shifts in the opposite direction.

The shift in the distribution of switch latencies in response to a step change in the probabilities of short and long trials is itself a step shift, rather than a gradual shift spread over many trials. Moreover, it often occurs before the mouse has lost a single pellet (Kheifets & Gallistel, 2012), so adjustment to a change in the relative frequency of the two trial types cannot be based on differential reinforcement. The experienced probabilities of reinforcement are in all cases very close to the programmed probabilities.

These previous results imply that the mouse estimates discrete probabilities, that is, it forms the ratio of two numerical estimates, an estimate of the number of short trials experienced and an estimate of the total number of trials experienced (or possibly the ratio of the number of short trials to the number of long trials, that is, the odds ratio). These previous results further imply that it does not make these estimates by reinforcement learning using delta-rule updating or a particle filter, as has often been assumed for human subjects (Behrens, Woolrich, Walton, & Rushworth, 2007; Nassar, Wilson, Heasly, & Gold, 2010), because estimates made that way change gradually—unless the learning rate is set to 1, in which case, the estimates can only be 0 or 1, depending on the most recent trial duration (cf. Simen et al., 2011). Recent delta-rule models assume that the learning rate can itself be modulated by delta-rule updating of a running-average estimate of variability (Nassar et al., 2010; Simen et al., 2011).

In this experiment, we ask whether the latencies with which subjects adjust the distribution of their switch latencies to a change in the relative frequencies of the short and long trials are sensitive to the informational asymmetry between the case where the subject believes $p(\text{long})$ is .9 when in fact there has been an unsigned drop to .5, versus the case where the subject believes it is .5 when in fact there has been an unsigned rise to .9. In both cases what we call a subject's belief rests on its prior experience. Because the changes in probability are unsigned, there is a mismatch in the postchange trials between its prior experience and its current experience of the relative frequencies of the two kinds of trials. An ideal observer/agent will adjust to the transitions from .9 to .5 more quickly than to the transitions from .5 to .9, because the Kullback-Leibler (KL) divergence is asymmetrical. The KL divergence, D_{KL} , measures the informational divergence of one probability distribution, Q , from another probability distribution, P , with the same support:

$$D_{\text{KL}}(P||Q) = \sum_i P(i) \log_2 \frac{P(i)}{Q(i)} \quad (1)$$

It is called the *divergence of the Q distribution from the P distribution* rather than the *distance* between them because it is asymmetric: For any two unequal values of p , both greater than or equal to .5, the divergence of the higher probability from the lower probability is greater than the divergence of the lower from the higher. Thus, when adjusting downward from an erroneously high estimate, the rate of information gain is greater than when adjusting upward from an erroneously low estimate.

Therefore, when we drop the value of the hidden parameter $p(\text{long})$ from a high value (.9) to a middle value (.5), information about the new value of p accumulates more rapidly than when we raise the probability over the same interval (from .5 to .9). The divergence (in bits) when the assumed value of $p = .9$ and the true value is .5 is approximately .74, while the divergence going the other way is approximately .53. Thus, if a subject detects changes by processing the remembered sequence of short and long trials with an approximately optimal change-detecting algorithm, then its response to the change in the relative frequencies of the short and long trials should happen sooner when $p(\text{long})$ goes from .9 to .5 when it goes the other way, by a factor of about $.74/.53 = 1.4$.

As a control, we also switched $p(\text{long})$ from .85 down to .5, because although this is a downward transition and although .85 is not much lower than .9, the divergence of Bernoulli($p = .85$) from Bernoulli($p = .5$) is .49 bits, which is slightly less than the .53 bit divergence of Bernoulli($p = .5$) from Bernoulli($p = .9$). Therefore, the trials required for the shift to appear in response to this slightly smaller downward transition should be approximately the same as the trials required for the upward transition from Bernoulli($p = .5$) to Bernoulli($p = .9$).

The question is theoretically important because the learning-rule assumptions in most timing theories imply that subjects' adjustment latencies will not be sensitive to this informational asymmetry. The learning-rule assumptions in most timing theories predict that subjects' latencies to adjust to the two directions of transition will be the same.

Method

Subjects were seven male C57BL/6 mice from Jackson Laboratories, aged about 4 weeks on arrival in the lab and about 5 weeks at the beginning of training. The initial training was that described in the general methods for the mice experiments: matching, followed by two-hopper autoshaping, followed by the switch protocol with feed latencies initially set at 4 s and 12 s and subsequently narrowed to 4 s and 8 s. The intertrial intervals were exponentially distributed with an expectation of 240 s. The protocol was identical to that for the mouse subjects in Experiment 2A: if the mouse stayed too long at the short hopper on a long trial (a trial where it should switch when the short feed latency has elapsed), then the first poke into the short hopper at or after the elapse of the long feed latency terminated the trial without reinforcement, and no switch latency could be calculated. This is in contrast to the protocol used with mice in Experiment 1 and with rats in Experiment 2B, where these trials (stayed-too-long trials) terminated only when the subject finally switched to the long hopper, and a switch latency was calculated. When there were switches on a majority of long trials with feed latencies at 4 s and 8 s and $p(\text{long})$ at .5 and the switch latencies had stabilized, we began to run sessions in which subjects switched either from $p(\text{long}) = .5$ to $p(\text{long}) = .9$ or from $p(\text{long}) = .9$ to $p(\text{long}) = .5$. After each change in $p(\text{long})$, it remained constant for approximately 300 trials before the next change. The changes in $p(\text{long})$ were unsignaled and they occurred at randomly chosen points within the middle third of a session.

Results

We estimated the trial at which the change in distribution occurred using a tri-linear transition function to map from trials to the values of the parameter vector of the distribution fit to the switch latencies. The transition function for a given parameter vector, θ was flat with value θ_b up to t_s and flat with value θ_a after t_f , as seen in Figure 16. θ_b denotes the parameter vector before the transition; θ_a denotes it after the transition. The elements of θ_b (one for each parameter) were obtained by fitting the distribution function to the switch latencies before t_c , while the elements of θ_a were obtained by fitting it to the switch latencies after t_c . The transition function describes a linear transition from the vector θ_b to the vector θ_a between trial t_s and t_f . The values for the two parameters of the transition (t_s and t_f) were obtained by an iterative search for the values that maximized the likelihood (that is, the joint probability of the observed switch latencies). We took $t_s - t_c$ to be T , the estimated start of the transition, that is, the estimated latency to detect the change in relative probabilities of the short and long trials.

Figure 17 gives the cumulative distributions of the estimated shift latencies (T 's) for the three transitions of interest. The downward transitions from $p(\text{long}) = .9$ to $p(\text{long}) = .5$ (thick solid plot in Fig. 17) occurred after significantly fewer trials than the upward transitions over the same p interval (thick dashed plot), as predicted by the fact that the divergence of Bernoulli($p = .9$), the old distribution of long and short trials, from Bernoulli($p = .5$), the new one, is 40% greater than the divergence going the other way.

One sees also in Figure 17 that the latencies of the downward transitions from $p(\text{long}) = .85$ to $p(\text{long}) = .5$ (thin solid plot) had approximately the same distribution as the upward transitions from $p(\text{long}) = .5$ to $p(\text{long}) = .9$, as would be expected from the fact that the divergences of the old distributions from the new are approximately the same in these two cases. Thus, we conclude that the process that detects changes in hidden values of p is sensitive to the information-theoretic divergence of the old distribution from the new distribution. As previously noted, this is not true for processes that estimate values of p by delta-rule updating. It would seem to require a theory in which change points are detected by computations carried out on a record in memory of the experienced sequence of short and long trials.

General Discussion

Experiment 1 showed that the mouse has some degree of control over the precision of its response timing and that it can increase or decrease precision (the inverse of the cv) independently of the changes it makes in the mean. The adjustments in precision appear to be both rapid and abrupt. The second experiment showed that mice and rats distinguish between measurement error and exogenous variability in the measured reinforcement latencies. The third experiment showed that the number of trials preceding a step-shift in switch-latency distributions in response to changes in the relative frequencies of short and long trials is appropriately sensitive to the Kullback-Leibler divergence of $\text{Bern}(p_{\text{old}})$ from $\text{Bern}(p_{\text{new}})$.

Humans attempting to track changing probabilities behave like mice in that they make step adjustments to their probability estimates after observing several or even many trials (Gallistel, Krishan, et al., 2014), rather than adjusting trial by trial, as would be expected in delta-rule, particle-filter and Bayesian updating models (Brown & Steyvers, 2009; Nassar et al., 2012; Nassar et al., 2010; Steyvers & Brown 2006; Wilson, Nassar, & Gold, 2010). Step changes in response latencies in response to changed reinforcement latencies have been observed in humans (Simen et al., 2011) and in pigeons (Higa, Thaw, & Staddon, 1993). Short-latency step changes in the distributions of visit durations have also been observed in matching paradigms in response to step changes in the hidden parameters of the concurrent variable interval schedules of reinforcement (Gallistel et al., 2007; Gallistel, Mark, King, & Latham, 2001; Mark & Gallistel, 1994). In short, step-like adjustments to changes in hidden stochastic parameters are frequently observed.

Implications for Theories of Timing

As explained in the Introduction, theories that seek to explain the timing of conditioned responses fall into two broad classes, record-free theories and record-based theories. The original version of the Behavioral Theory of Timing (Killeen & Fetterman, 1988) and neurally oriented theories (Fiala et al., 1996; Grossberg & Schmajuk, 1989; Karmarkar & Buonomano, 2007; Yamazaki & Tanaka, 2009) are examples of record-free theories. Scalar Expectancy Theory (SET, Gibbon, 1977; Gibbon, Church, & Meck, 1984), The Analytic Theory of Associative Learning (TATAL, Gallistel & Wilkes, 2016; Wilkes & Gallistel, 2017) and a later version of BeT (Fetterman & Killeen, 1995) are examples of record-based theories. In a record-based theory, there is at least one referential record in memory, whereas in a record-free theory, there are no records in the brain of any objective fact about the animal's experience. A referential record in memory encodes a quantitative fact gleaned from experience, such as the duration of an individual reinforcement latency or an average reinforcement latency. In theories that assume sophisticated models of experience, these records are, in essence, numbers: They are intrinsically orderable and the brain computes with them—it adds, subtracts, multiplies and divides them.

Record-based theories have two sub-classes: 1) minimal-record theories and 2) rich-record theories. Minimal-record theories keep the number of learned, referential quantities in memory to a minimum, usually just a running average of some experiential statistic (e.g., reinforcement latencies and/or reinforcement probabilities). Rich-record theories assume that the subject records the sequences of experienced intervals and their outcomes; in other words, it has an episodic memory that encodes a variety of facts about each trial/episode (see Crystal, Alford, Zhou, & Hohmann, 2013; Crystal & Smith, 2014; Panoz-Brown et al., 2016; Wilson et al., 2015). Thus, rich-record theories assume extensive referential memory, while minimal-record theories assume very little.

Rich-Record Theories

The most influential rich-record theory of timing is SET (Gibbon, 1977). Despite the richness of the record it assumes, it cannot account for the fact that subjects independently vary the mean and cv of their switch-latency distribution, nor for their ability to distinguish their measurement error from variability in the intervals being timed. In SET, scalar noise is

a basic postulate. It is assumed that when reinforcement latencies are fixed, the reinforcement latencies in memory have a dispersion that is proportional to the fixed reinforcement latency. The dispersion of the conditioned response latencies is determined by the dispersion of the latencies in memory, and the variance of remembered intervals is the sum of the external variance and the endogenous variance. Therefore, increasing the external variability must increase the variability of the response latencies. The same point applies to many models other than SET; we begin with SET because the sparseness and clarity of its assumptions make it easy to see the problem.

Distinguishing measurement error from external variability—Measurement error is estimated by measuring the same thing repeatedly, thereby eliminating variability in the thing measured. The ever-onward flow of time prevents an interval-measuring process from measuring the identical interval on successive occasions; therefore, the only way to estimate the measurement error in interval timing is to apply multiple independent interval-measuring processes to each interval measured. Recent results on the timing of reinforcement latencies by individual cerebellar Purkinje cells make it plausible that the brain mechanisms make multiple measurements of single intervals. Johansson, Jirenhed, Rasmussen, Zucc, and Hesslow (2014) show that individual Purkinje cells time and remember the durations of the CS-US intervals in eyeblink conditioning in decerebrate ferrets. This gives neurobiological insight into how the brain could distinguish between the contributions of measurement error and external variability to the variability in the results of its temporal measurements. If several cells independently time and remember each experienced interval, then the within-trial variance in the results of those measurements estimates the measurement error. The estimate of the external variance is then the variance in the between-trial results minus the within-trial variance.

The source of scalar variability—The source has always been assumed to be noise. We suggest that it is not noise; rather, it is a limit on the precision in a discrete (digital), noiseless representation of quantities of all kinds, including durations. The precision of an analog record is noise limited; a digital record is noiseless, but its precision is limited by the amount of physical resources allocated to the representation of a quantity, for example, the number of base-pairs allocated to the representation of each number in a DNA-based memory (Extance, 2016).

The representation of quantity is precision limited in the fixed-point numbers with which a maximally efficient computer most often operates and in scientific notation. In both cases, the representation takes the form mb^p . In scientific notation, b is 10; in a fixed-point binary number, it is 2. In a fixed-point number, the number of bits in the coefficient is fixed. When scientific notation is used, the number digits in m is commonly fixed by the percent accuracy with which the quantities are assumed to be known. The representations of quantity in scientific notation and in a computer using a fixed-point representation obey Weber's law because the interval between any two representable quantities is proportional to the scale factor, b^p . Killeen and Taylor (2001) show that stochastic counters have a similar property. When a quantity in memory that is the target for a behavior (e.g., a decision threshold) has some uncertainty due to the limited number of bits in its representation of that quantity

(quantization error), it is functionally appropriate for the response-generating process to dither the output (the timing of the behavior) so as to span the uncertainty inherent in the limited number of bits in the representation of the target.

In short, at least one neurally plausible way to explain the results of Experiments 2 A and B within the context of a rich-model theory is to assume that the duration of each trial is measured by multiple timers. This makes it possible to distinguish between measurement error and endogenous variability. At least one way to explain the results of Experiment 1 is to assume that target intervals (decision criteria) are specified with a limited number of bits. A threshold represented by 4 bits would be specified to within $\pm 12.5\%$, which is roughly the lower end of the range of estimates for the Weber fraction in interval timing. A threshold represented by 3 bits would be specified to within $\pm 25\%$, which is roughly the upper end of the Weber fraction estimates for interval timing. Varying the number of bits with which the threshold is represented and dithering the response within the resulting range of uncertainty would explain the results of Experiment 1, in which the mice varied the precision of their response timing in response to varying task demands.

Minimal Record Theories

In a later development of BeT, Fetterman and Killeen (1995) abandoned the original commitment to intrinsically unordered behavioral states as the causes of timed responding (see Record-Free Theories below for a discussion of the original BeT). They adopted the SET assumption that the brain contained a pacemaker and a counter and that behavior was initiated on any given trial when the count of pacemaker pulses reached a criterial value drawn from memory for past experience (or a running average over those experiences). They also incorporated several additional sources of variation. Counts are, of course, both referential and intrinsically ordered.

With this development, BeT went from being a record-free theory to a minimal-record theory. In its new incarnation, it, like SET, and like the quasilogarithmic theory of Josefowicz, Staddon, and Cerutti (2009), offers no means of explaining how the precision of response timing may be adjusted independently of the mean response latency and independently of objective variation in the reinforcement latencies.

Neurally oriented theories of timing treat the timing of intervals as the product of complex and extensive neural circuitry (Balci et al., 2011; Durstewitz, 2003; Finnerty, Shadlen, Jazayeri, Nobre, & Buonomano, 2015; Grossberg & Schmajuk, 1991; Matell & Meck, 2004; Oprisan & Buhusi, 2011; Simen et al., 2011; Simen, Rivest, Ludvig, & Killeen, 2013; Yamazaki & Tanaka, 2009). The record-free versions of these theories do not posit a memory for previously experienced intervals. Even the record-based versions (Durstewitz, 2003; Matell & Meck, 2004; Oprisan & Buhusi, 2011; Simen et al., 2011) do not posit a generally readable episodic memory that stores the results of each interval in a sequence of separately experienced intervals, and makes the whole sequence accessible to diverse and complex later computations. To posit that would be to propose a read–write memory mechanism, which theorists concerned with neural plausibility are reluctant to do (Gallistel & King, 2010).

The Simen et al.'s (2011) neurally oriented model is an amalgam of the record-based version of BeT and drift diffusion theories of decision making (Ratcliff & McKoon, 2008). Like BeT and like drift diffusion theory, it assumes a time signal that rises linearly to a fixed output threshold. Like BeT, it also assumes that the system varies the output latency by varying the rate of rise in the time signal. Unlike record-based BeT, the Simen et al. theory assumes Gaussian noise in the time signal. As they explain, this assumption entails that the empirical distribution of response latencies be best described by the Wald distribution, the distribution that describes the level-crossing times for a drift-diffusion process.

A problem our results pose for drift-diffusion theories of timing in general is that the Wald distribution poorly describes the distribution of switch latencies (Experiment 1). However, we believe that the results of Experiment 2 pose a more serious problem for the Simen et al. theory. Their learning rule rapidly adjusts the rate of rise in response to changes in the experienced reinforcement latency. Thus, their theory predicts that in our Experiment 2, where the experienced reinforcement latencies vary widely and randomly over a large range, the response latencies should show greatly increased variability, which they do not; the variability in response latencies is unaffected or even diminished by variability in the timed durations.

Record-Free Theories

The essential idea behind record-free theories is most simply grasped by considering the original version of BeT (Killeen & Fetterman, 1988). In this theory, a Poisson (random rate) process generates pulses that move the subject through a sequence of behavioral states called adjunctive behaviors. Examples of behavioral states in a pigeon are preening the upper body, preening the lower body, fluttering the feathers, scratching the substrate, and so on. It is assumed that behavioral states can serve as discriminative stimuli for operant responses. If a subject emits an operant response while in a given behavioral state and that response is followed immediately by reinforcement, then that behavioral state acquires discriminative control over that response. The reinforcement strengthens the subject's tendency to emit that response when in that behavioral state. It is natural to imagine that the strengthening of the tendency of a stimulus to elicit a response reflects the strengthening of synaptic conductances in a neural pathway connecting that activity evoked by the stimulus to the activity that produces the response, which is why the idea of reinforcement (strengthening) appears in all neurally oriented theories of timing. It is the same idea as in behaviorist theorizing but it is applied to notional synapses rather than to response tendencies or to associative bonds.

Behavioral states have no intrinsic order, unlike counts. Fluffing the feathers is not intrinsically a later or larger behavior than preening the upper body. The order in which the states occur in a timing task may vary between subjects for a given task and within subjects across tasks. It is important to bear this in mind when we refer to successive behavioral states by index numbers. The index numbers have an intrinsic order but the states they index do not. The theory does assume, however, that the states occur in a fixed sequence for a given subject and given timing task. There is, however, no counter, only a pacemaker. Because durations are not measured, a fortiori they are not recorded in a memory.

Because stepping through the behavioral states is driven by pulses from a Poisson process, the time taken to arrive at the n^{th} state is gamma distributed, with parameters n and τ , the shape and scale parameters of the gamma distribution. We estimated n for the switch latencies in Experiment 1 when we fitted gamma distributions to them. The cumulative distributions of these estimates are in the fourth row of Figure 8.

Because the mean of a gamma distribution is $n\tau$, we have $\hat{\tau} = \hat{\mu} / \hat{n}$, where the hat denotes estimate and $\hat{\mu}$ is the mean of the observed latencies. If the rate of the Poisson process were constant across timing tasks, then tasks that reinforced after longer latencies would reinforce responses to states with higher index numbers. In that case, the cv of those response-latency distributions would be smaller because the cv of the gamma distribution is $\sqrt{n}/n$. However, the second key assumption in BeT is that τ , which is conceived of as a state of arousal, is proportionate to T , which is the average period between reinforcements. In the situations for which BeT was originally devised, most simply responding on a fixed interval schedule, doubling the fixed interval doubles T , thereby halving the subject's state of arousal, thereby doubling τ , the average interval between pulses from the pacemaker, with the result that it takes the subject on average twice as long to reach the n^{th} behavioral state. Because the number of pacemaker pulses required to reach the reinforced state does not change as we change reinforcement latency, neither does the cv. Thus, the original version of BeT derives scalar variability rather than postulating it, which is part of its attraction as a theory.

The median cvs in Experiment 1 implied that the state the subject was in when it switched was approximately the 20th state it had passed through since the start of the state sequence (fourth row of Fig. 8). The median mean of the switch latencies was about 6.5 s, so when BeT is applied to our data, we must assume that the mouse passes through roughly three behavioral states per s. This does some violence to our own notion of what could plausibly be called a behavioral state, but that is not the most serious problem that our data pose for BeT.

The first question in applying BeT to the switch paradigm is how to characterize the reinforced response or responses. It would seem that the switch response itself—leaving the short hopper for the long—cannot be regarded as the reinforced response, because it occurs between the two feed latencies, at trial times that are almost never reinforced. The only responses that are reinforced at the time they are made are poking into the short hopper and poking into the long hopper. Thus, in applying BeT we assume that these are the reinforced responses. Therefore, the behavioral states at the latencies when these reinforcements are delivered are the states that control the conditioned behavior. In short, we assume that the state that gains discriminative control of the switch response is the first state at which the tendency to poke into the long hopper becomes stronger than the tendency to poke into the short hopper.

There is also a question as to what constitutes the contextual rate of reinforcement, because this determines the value of τ . At least four contexts suggest themselves on first consideration: 1) T is the average interval between reinforcements when the subject is at a given hopper; 2) T is the average interval between reinforcement during a trial; 3) T is the

average interval between reinforcements during one of the 4-hr-long feeding periods; 4) T is the average interval between reinforcements in the test chamber.

Under Context Definition 1 (the hopper is the context), the context changes depending on whether the mouse is at the short- or long-latency hopper. It would seem that T cannot readily be computed under this first definition. It would have to have one value early in a trial, when the mouse is at the short-latency hopper, a different value later in the trial, when the mouse is at the long-latency hopper, and no definable value when the mouse is in transition between hoppers or before it has arrived at the short-latency hopper. For this reason, we do not entertain this definition.

Either Context Definition 3 (feeding session is the context) or Context Definition 4 (test environment is the context) is the original definition. Under either of these definitions, T is essentially fixed in our protocols, hence also τ , because the animal is in the chamber 24/7, with two feeding phases per day, each lasting 4 hr. When T is fixed, τ has the same value throughout a trial. Thus, the cv must decrease from the first feed latency to the second as the inverse of the square root of the ratio between the long and short feed latencies.

We assume that τ is determined by the average rate of reinforcement during a trial (Context Definition 2). Therefore, it is the same on long and short trials. It does vary somewhat with changes in condition, because $T = p_s L_s + (1 - p_s) L_l$, where p_s is the probability of reinforcement at the short-latency hopper and L_s and L_l are the short and long latencies. However, it does not scale inversely with the changes in n induced by those changes.

Because τ must be assumed to be constant throughout a trial, the fact that there are two different responses reinforced at substantially different latencies when the subject is at substantially different points in the sequence of behavioral states poses a problem for BeT. It implies that the cvs for the poking at the short and long hoppers must differ by large amounts. In our 4–12 conditions, the long feed latency is three times longer than the short one, so the cv for the long-latency state must be smaller than the cv for the short-latency state by a factor of $1/\sqrt{3} = 0.58$. This violates the empirically well-established scalar variability principle. More importantly, it has been shown to be false for mice in the double-switch paradigm.

Gallistel et al. (2004) ran a two-switch version of the switch protocol. There were three hoppers, where poking paid off at three latencies, 5, 15 and 45 s, depending on the hopper. As in the 1-switch paradigm, a computer chose one of the three to pay off on any given trial. The mice had no way of knowing at the start of a trial which hopper the computer had chosen. They learned to go first to the 5 s hopper, then, if that did not pay off, to the 15 s hopper, and then, if that also did not pay off, on to the 45 s hopper. This protocol yields the means and cvs for departures (stop times) from the short-latency hopper (between 5 and 15 s) and for departures from the middle-latency hopper (between 15 and 45 s). The mean latencies differed for these two switches by a factor of 3, while the cvs, which had values similar to those we report here, did not differ (Gallistel et al., 2004, p. 11, Table 2). Moreover, the cvs for the two switches in the double-switch protocol were the same as for the stop latencies obtained from the same subjects when run in the peak procedure with

blocked target intervals of 5, 15 and 45 s, prior to their being run in the double-switch protocol. Thus, in the switch paradigm, contrary to the predictions of the original version of BeT, within-trial cvs do not get smaller in proportion to the $\sqrt{L}/L$, where L stands for reinforcement latency.

Entirely record-free theories have been all but abandoned by theorists not concerned with neural plausibility. However, they remain popular with theorists for whom this is a principal concern. An example are dynamic state models, such as that put forward by Karmarkar and Buonomano (2007). Their model is conceptually similar to the original BeT, except that it uses states of complex neural circuits instead of behavioral states, and it has no pacemaker. If a large object is dropped into a swimming pool with an irregular shape, a complex pattern of waves reflecting off the walls and reinforcing and canceling each other plays out over time. The same is true for the spatial pattern of activity in a perturbed complex neural circuit with appropriately chosen parameters for the interactions between its neural elements. The state of the surface of the water differs at each different moment in time, and so does the simulated activity pattern of such a circuit. Thus, if the brain forges associations (altered synaptic conductances) between those neurons in the circuit that are active at the time a reinforced response occurs, that response will occur only at the time when that pattern appears in the network, the next time activity in the network is triggered by an appropriate perturbation.

As the authors themselves stress, the different spatial activity patterns marking different points in time are not intrinsically ordered. A fortiori, this model makes no attempt to account for the scalar variability that is such a characteristic feature of interval timing. The cv is a metric concept. As the authors themselves stress, there is no temporal metric in their theory; the different brain states, like the different behavioral states in the original BeT, do not measure anything, and the associative connections that form do not record anything.

Spectral timing models (Grossberg & Schmajuk, 1991; Yamazaki & Tanaka, 2009), including beat-frequency models (Matell & Meck, 2004; Matell, Meck, & Nicolelis, 2003; Oprisan & Buhusi, 2011) are cousins to this no-pacemaker, no-ordered-time-markers model, even when they include some assumptions that move them technically into the class of record-based theories. Spectral theories posit a spectrum of dynamics in neurons that provide the inputs to response-generating neurons. In the cerebellum, for example, the input neurons to the Purkinje cell are the parallel fibers that originate from the tiny, extraordinarily numerous granular cells, which constitute much more than 80% of the neurons in a mammalian brain. In the original spectral timing theory, the neurons whose intrinsic dynamics are assumed to be such that they peak at the appropriate time are assumed to become selectively associated with the response. (There is no evidence that granule cells have the assumed dynamics.) Beat frequency models posit a spectrum of damped oscillations at different frequencies. Different subsets peak simultaneously at different times after the population is set in motion. The subset that peaks at the reinforcement latency becomes associated with a conditioned response.

Neither the postulated presynaptic neurons with different rise and fall times in a spectral model nor the different oscillatory subsets in a beat-frequency model constitute intrinsically

ordered time markers, nor do the associative connections between them and the response. One would have to read both the associative strengths and the dynamics of each associated input neuron to infer the duration of the learned interval. Thus, the strengthened associations in spectral timing theories do not constitute a referential record.

For all of these theories, varying the feed latencies should spread the spectrum of associative strengths out over a broader range of presynaptic neurons (a broader portion of the timing spectrum). That should either eliminate effective timing altogether, or, at the very least, increase the variability of the output. Moreover, none of these models offers a way for the animal to adjust the precision of its timed behavior. As with SET and the record-based version of BeT, the endogenous noise is not under the subject's control in these models.

It may also be remarked that none of the record-free theories of interval timing do justice to what is known about the sophisticated use animals make of remembered durations. Salient aspects of conditioned behavior depend on arithmetic relations between interval durations learned from different experiences. The number of trials to the acquisition of the conditioned response depends on the ratio of the CS–US interval to the average US–US interval (Balsam & Gallistel, 2009; Gallistel & Gibbon, 2000; Gottlieb, 2008; Ward, Gallistel, & Balsam, 2013; Ward et al., 2012). The choices pigeons make can be made to depend on arithmetic differences in separately experienced interval durations (Gibbon & Church, 1981) or on separately experienced numbers of pecks (Brannon, Wusthoff, Gallistel, & Gibbon, 2001). The conditioned responses in sensory preconditioning and secondary reinforcement paradigms can be made to depend on the differences between separately learned intervals (Denniston, Blaisdell, & Miller, 1998; Denniston, Blaisdell, & Miller, 2004; Matzel, Held, & Miller, 1988). These results seem to require the assumption of explicitly represented, readable representations of the durations of experienced intervals, that is, rich records of experience.

Change Detection

That mice attach probabilities to short and long trials is remarkable. The trials are separately experienced higher-level events in the flow of the subject's experience, with lengthy and highly variable intervals between them. They are what Crystal and his collaborators call episodes (Crystal et al., 2013; Crystal & Smith, 2014; Panoz-Brown et al., 2016; Wilson et al., 2015). These episodes are constituted of different lower-level point events:

[(light-on-in-left-hopper)→(poke in left hopper at or up to short feed latency)-
→(pellet-drop-in-left-hopper)]

and

[(light-on-in-right-hopper)→(poke in right hopper at, after or up to long feed
latency)→(pellet-drop-in-right-hopper)].

Unlike the point events of which they are constituted, these episodes have duration. The fact that the mouse can attach probabilities to episodes distinguished by their durations would seem to imply that the representations of these events are countable entities in a hierarchical data structure in memory. The results from probability shifting experiments using the switch protocol, such as Experiment 3, seem to imply that the mouse can count the number of

remembered short episodes and the number of long episodes and compute the ratio of the two counts (the odds ratio).

The data from human trial-by-trial probability estimating give an important further result: In trial-by-trial updating models, estimates of a Bernoulli probability should vary trial by trial, because Bernoulli outcomes are always at the extremes (1 or 0), while the current estimate of p usually lies in between. Thus, the observed outcome is far from the current estimate of p on most trials. In fact, however, the Bernoulli probability estimates from human subjects remain constant over long sequences of trials, then change abruptly (Gallistel, Krishan, et al., 2014; Ricci & Gallistel, 2017). These and other results from human subjects have led us to a model in which the subject is assumed to compute the strength of the evidence that the most recent observations are inconsistent with the current estimate (Gallistel, Krishan, et al., 2014; Robinson, 1964). The Kullback-Leibler divergence comes naturally into this computation because it is a factor in the measure of the extent to which the current estimate differs from the estimate implied by recent observations.

We have not been able to devise a method of obtaining trial-by-trial estimates of a hidden Bernoulli probability parameter from nonverbal subjects. We do know, however, that nonverbal subjects adjust as abruptly as human subjects. Mice adjust the location of their switch-latency distribution within the span of a single trial, often before they have missed a single reinforcement (Kheifets & Gallistel, 2012). The step-like abruptness of the adjustment is not what one would expect from any trial-by-trial updating process. Trial-by-trial updating in response to a change in input implies a gradual adjustment—unless the learning rate is set to 1, in which case the behavior on trial $n + 1$ will be entirely determined by the outcome on trial n . Similarly abrupt adjustments occur in response to changes in the hidden parameters of the exponential distributions of interreinforcement intervals in concurrent variable-interval matching experiments (Gallistel et al., 2007; Gallistel et al., 2001; Mark & Gallistel, 1994).

In Experiment 3, we have now shown that the latency of the abrupt adjustment to a change in the hidden probability parameter is appropriately sensitive to the asymmetry in the Kullback-Leibler divergence of a midlevel probability (.5) from a high probability (.9). The adjustment latency is shorter when the probability changes from .9 to .5 than when it changes from .5 to .9. This result further constrains theories of the processes that mediate learned adjustments to nonstationary stochastic processes.

Normative stochastic models for the real-time detection of changes in the parameters of stochastic processes (e.g., Adams & MacKay, 2006), which are sensitive to the KL asymmetry, operate on a record of the sequence of outcomes. They need this record in order to compute the relative likelihoods of stochastic models for the sequences that differ in where the change is assumed to have occurred. Theories of stochastic change-point detection based on normative assumptions (Gallistel, Krishan, et al., 2014; Gallistel et al., 2001) predict step adjustments, because once a change has been detected and its locus within the sequence of past experiences estimated, the new estimate of the hidden stochastic parameter is based only on the postchange portion of the sequence.

Minimal-record theories for the estimation of hidden stochastic parameters employ delta-rule updating to obtain a running average estimate of the hidden parameter. They face the following challenge: In order to obtain an accurate and stable estimate of the current value of the parameter, the learning rate must be slow. That is, the coefficient in the delta rule has to be small, so that individual experiences produce little trial-by-trial change in the running average. However, to explain abrupt behavioral adjustments to changes in the input, the learning rate has to be high. To explain large changes within the span of a single trial, the learning rate has to be set to 1 (Nassar et al., 2010; Simen et al., 2011).

The challenge is especially strong when subjects estimate the Bernoulli p parameter, because the outcomes of a Bernoulli process are always at one extreme or the other (either a “success” or a “failure”, a 1 or a 0), whereas the value of the parameter is often well away from the extremes. With the learning rate set to one, the behavior on trial $n + 1$ is entirely determined by the outcome on trial n . In the switch paradigm, this would mean that a feed at the short hopper should lead to a long string of trials in which the subject pokes only at that hopper, and likewise for the long hopper. That is not observed. Nonetheless, when the hidden Bernoulli parameter changes, subjects often adjust to the change after relatively few further trials, and, when they do adjust, they make the full adjustment within the span of a single trial. Thus, earlier results showing that adjustments are maximally abrupt in both human and mouse subjects (Gallistel, Krishan, et al., 2014; Kheifets & Gallistel, 2012) are evidence in favor of rich-record theories as opposed to minimal-record theories.

In conclusion, the results of three experiments seem to imply that rodents retain in memory a trial-by-trial (episode-by-episode) record of recent events, not just running averages. They make multiple measurements of each experienced latency, which enable them to distinguish between their measurement error (within-sample variance) and objective variation in the latencies (between-sample variance). That is how the animal may know its measurement error, that is, its uncertainty. Knowing its uncertainty enables it to choose an appropriate precision with which to represent its estimate of an interval. By varying the precision with which they represent the results of interval measurements and dithering their responses to span the resulting uncertainty, subjects can control the precision of their response timing. Because they preserve in memory the sequence of episodes, they are able to do approximately normative detection of changes in the values of the hidden parameters of the stochastic processes generating those episodes. That is why their adjustments to these changes are sensitive to the informational asymmetry inherent in them. And that is why the changes are abrupt: Having detected the change and estimated the locus in the experienced sequence at which it occurred, they base their new behavior on only the postchange portion of the remembered sequence.

Acknowledgments

This work was supported by NIMH Grant R01 MH077027 to CRG and by Russell Church. The authors declare no competing financial interests.

References

- Adams, RP., MacKay, DJC. Bayesian online changepoint detection. 2006. Retrieved from <http://arxiv.org/abs/0710.3742>
- Balci F, Freestone D, Gallistel CR. Risk assessment in man and mouse. *Proceedings of the National Academy of Science USA*. 2009; 106(7):2459–2463.
- Balci F, Freestone D, Simen P, deSouza L, Cohen JD, Holmes P. Optimal temporal risk assessment. *Frontiers in Integrative Neuroscience*. 2011; doi: 10.3389/fnint.2011.00056
- Balci F, Papachristos EB, Gallistel CR, Brunner D, Gibson J, Shumyatsky GP. Interval timing in the genetically modified mouse: A simple paradigm. *Genes, Brains & Behavior*. 2008; 7:373–384.
- Balsam P, Gallistel CR. Temporal maps and informativeness in associative learning. *Trends in Neurosciences*. 2009; 32(2):73–78. <http://dx.doi.org/10.1016/j.tins.2008.10.004>. [PubMed: 19136158]
- Behrens TE, Woolrich MW, Walton ME, Rushworth MF. Learning the value of information in an uncertain world. *Nature Neuroscience*. 2007; 10(9):1214–1221. [PubMed: 17676057]
- Brannon EM, Wusthoff CJ, Gallistel CR, Gibbon J. Numerical subtraction in the pigeon: Evidence for a linear subjective number scale. *Psychological Science*. 2001; 12(3):238–243. [PubMed: 11437307]
- Brown SD, Steyvers M. Predicting and detecting changes. *Cognitive Psychology*. 2009; 58:49–67. [PubMed: 18976988]
- Cover, TM., Thomas, JA. *Elements of information theory*. 2nd. New York: Wiley Interscience; 1991.
- Crystal JD, Alford WT, Zhou W, Hohmann AG. Source memory in the rat. *Current Biology*. 2013; 23:387–391. [PubMed: 23394830]
- Crystal JD, Smith AE. Binding of episodic memories in the rat. *Current Biology*. 2014; 24:2957–2961. DOI: 10.1016/j.cub.2014.10.074 [PubMed: 25466681]
- Denniston JC, Blaisdell AP, Miller RR. Temporal coding affects transfer of serial and simultaneous inhibitors. *Animal Learning and Behavior*. 1998; 26(3):336–350.
- Denniston JC, Blaisdell AP, Miller RR. Temporal coding in conditioned inhibition: analysis of associative structure of inhibition. *Journal of Experimental Psychology: Animal Behavior Processes*. 2004; 30:190–202. [PubMed: 15279510]
- Durstewitz D. Self-organizing neural integrator predicts interval times through climbing activity. *Journal of Neuroscience*. 2003; 23:5342–5353. [PubMed: 12832560]
- Extance A. How DNA could store all the world's data. *Nature*. 2016; 537(7618):22–24. [PubMed: 27582204]
- Fetterman JG, Killeen PR. Categorical scaling of time: Implications for clock-counter models. *Journal of Experimental Psychology: Animal Behavior Processes*. 1995; 21:43–63. [PubMed: 7844506]
- Fiala JC, Grossberg S, Bullock D. Metabotropic glutamate receptor activation in cerebellar Purkinje cells as substrate for adaptive timing of the classically conditioned eye-blink response. *Journal of Neuroscience*. 1996; 16(11):3760–3774. [PubMed: 8642419]
- Finnerty G, Shadlen M, Jazayeri M, Nobre A, Buonomano D. Time in cortical circuits. *Journal of Neuroscience*. 2015; 35:13912–13916. [PubMed: 26468192]
- Gallistel CR. The importance of proving the null. *Psychological Review*. 2009; 116(2):439–453. [PubMed: 19348549]
- Gallistel CR, Balci F, Freestone D, Kheifets A, King AP. Automated, quantitative cognitive/behavioral screening of mice: For genetics, pharmacology, animal cognition and undergraduate instruction. *Journal of Visualized Experiments (JoVE)*. 2014; (84) (2013). doi: 10.3791/51047
- Gallistel CR, Gibbon J. Time, rate, and conditioning. *Psychological Review*. 2000; 107(2):289–344. [PubMed: 10789198]
- Gallistel, CR., King, AP. *Memory and the computational brain: Why cognitive science will transform neuroscience*. New York: Wiley/Blackwell; 2010.
- Gallistel CR, King AP, Gottlieb D, Balci F, Papachristos EB, Szalecki M, Carbone KS. Is matching innate? *Journal of the Experimental Analysis of Behavior*. 2007; 87(2):161–199. [PubMed: 17465311]

- Gallistel CR, King A, McDonald RJ. Sources of variability and systematic error in mouse timing behavior. *Journal of Experimental Psychology: Animal Behavior Processes*. 2004; 30(1):3–16. [PubMed: 14709111]
- Gallistel CR, Krishan M, Liu Y, Miller RR, Latham PE. The perception of probability. *Psychological Review*. 2014; 121:96–123. [PubMed: 24490790]
- Gallistel CR, Mark TA, King AP, Latham PE. The rat approximates an ideal detector of changes in rates of reward: implications for the law of effect. *Journal of Experimental Psychology: Animal Behavior Processes*. 2001; 27:354–372. [PubMed: 11676086]
- Gallistel CR, Wilkes JT. Minimum description length model selection in associative learning. *Current Opinion in Behavioral Science*. 2016; 11:8–13.
- Gibbon J. Scalar expectancy theory and Weber's Law in animal timing. *Psychological Review*. 1977; 84:279–335.
- Gibbon J, Church RM. Time left: linear versus logarithmic subjective time. *Journal of Experimental Psychology: Animal Behavior Processes*. 1981; 7(2):87–107. [PubMed: 7241054]
- Gibbon, J., Church, RM., Meck, WH. Scalar timing in memory. In: Gibbon, J., Allan, L., editors. *Timing and time perception*. Vol. 423. New York: New York Academy of Sciences; 1984. p. 52–77.
- Gottlieb DA. Is the number of trials a primary determinant of conditioned responding? *Journal of Experimental Psychology: Animal Behavior Processes*. 2008; 34(2):185–201. [PubMed: 18426303]
- Grossberg S, Schmajuk NA. Neural dynamics of adaptive timing and temporal discrimination during associative learning. *Neural Networks*. 1989; 2:79–102.
- Grossberg, S., Schmajuk, NA. Neural dynamics of adaptive timing and temporal discrimination during associative learning. In: Carpenter, GA., Grossberg, S., editors. *Pattern recognition by self-organizing neural networks*. Cambridge, MA: MIT Press; 1991. p. 637–674.
- Higa JJ, Thaw JM, Staddon JER. Pigeons' wait-time responses to transitions in interfood-interval duration: Another look a cyclic schedule performance. *Journal of the Experimental Analysis of Behavior*. 1993; 59:529–541. [PubMed: 16812693]
- Johansson F, Jirenhed DA, Rasmussen A, Zucc R, Hesslow G. Memory trace and timing mechanism localized to cerebellar Purkinje cells. *Proceedings of the National Academy of Science*. 2014; 111(41):14930–14934. DOI: 10.1073/pnas.14153711
- Jozefowicz J, Staddon JER, Cerutti DT. The behavioral economics of choice and interval timing. *Psychological Review*. 2009; 116(3):519–539. [PubMed: 19618985]
- Karmarkar UR, Buonomano DV. Timing in the absence of clocks: encoding time in neural network states. *Neuron*. 2007; 53:427–438. [PubMed: 17270738]
- Kheifets A, Gallistel CR. Mice take *calculated* risks. *Proceedings of the National Academy of Science*. 2012; 109:8776–8779.
- Killeen PR, Fetterman JG. A behavioral theory of timing. *Psychological Review*. 1988; 94:455–468.
- Killeen PR, Taylor T. How the propagation of error through stochastic counters affects time discrimination and other psychophysical judgments. *Psychological Review*. 2001; 107:430–459.
- Mark TA, Gallistel CR. Kinetics of matching. *Journal of Experimental Psychology: Animal Behavior Processes*. 1994; 20(1):79–95. [PubMed: 8308494]
- Matell MS, Meck WH. Cortico-striatal circuits and interval timing: coincidence detection of oscillatory processes. *Cognitive Brain Research*. 2004; 21:139–170. [PubMed: 15464348]
- Matell MS, Meck WH, Nicolelis MA. Interval timing and the encoding of signal duration by ensembles of cortical and striatal neurons. *Behavioral Neuroscience*. 2003; 117(4):760–773. [PubMed: 12931961]
- Matzel LD, Held FP, Miller RR. Information and expression of simultaneous and backward associations: Implications for contiguity theory. *Learning and Motivation*. 1988; 19:317–344.
- Nassar MR, Rumsey KM, Wilson RC, Parikh K, Heasley B, Gold JI. Rational regulation of learning dynamics by pupil-linked arousal systems. *Nature Neuroscience*. 2012; doi: 10.1038/nn.3130
- Nassar MR, Wilson RC, Heasley B, Gold J. An approximately Bayesian delta-rule model explains the dynamics of belief updating in changing environment. *Journal of Neuroscience*. 2010; 30(37): 12366–12378. [PubMed: 20844132]

- Oprisan SA, Buhusi CV. Modeling pharmacological clock and memory patterns of interval timing in a striatal beat-frequency model with realistic, noisy neurons. *Frontiers in Integrative Neuroscience*. 2011; doi: 10.3389/fnint.2011.00052
- Panoz-Brown D, Corbin HE, Dalecki SJ, Sluk CM, Wu JE, Crystal JD. Rats remember items in context using episodic memory. *Current Biology*. 2016; doi: 10.1016/j.cub.2016.08.023
- Ratcliff R, McKoon G. The diffusion decision model: theory and data for two-choice decision tasks. *Neural Computation*. 2008; 20(4):873–922. [PubMed: 18085991]
- Rescorla, RA., Wagner, AR. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In: Black, AH., Prokasy, WF., editors. *Classical conditioning II*. New York: Appleton-Century-Crofts; 1972. p. 64-99.
- Ricci M, Gallistel CR. Accurate step-hold tracking of smoothly varying periodic and aperiodic probability. *Attention, Perception and Psychophysics*. 2017; doi: 10.3758/s13414-017-1310-0
- Robinson GH. Continuous estimation of time-varying probability. *Ergonomics*. 1964; 7:7–11.
- Simen P, Balci F, deSouza L, Cohen JD, Holmes P. A model of interval timing by neural integration. *The Journal of Neuroscience*. 2011; 31:9238–9253. [PubMed: 21697374]
- Simen P, Rivest F, Ludvig EA, Killeen P. Timescale invariance in the pacemaker-accumulator family of timing models. *Timing and Time Perception*. 2013; 1:159–188. DOI: 10.1163/22134468-00002018
- Skinner, BF. *About behaviorism*. New York: Random House; 1974.
- Skinner BF. Why I am not a cognitive psychologist. *Behaviorism*. 1977; 5(2):1–10.
- Skinner BF. Can psychology be a science of mind? *American Psychologist*. 1990; 54(11):1206–1210.
- Steyvers M, Brown S. Prediction and change detection. *Advances in Neural Information Processing Systems*. 2006; 18:1281–1288.
- Ward RD, Gallistel CR, Balsam PD. It's the information! *Behavioral Processes*. 2013; 95:3–7.
- Ward RD, Gallistel CR, Jensen G, Richards VL, Fairhurst S, Balsam PD. Conditional stimulus informativeness governs conditioned stimulus—unconditioned stimulus associability. *Journal of Experimental Psychology: Animal Behavior Processes*. 2012; 38(1):217–232. DOI: 10.1037/a0027621 [PubMed: 22468633]
- Wilkes, JT., Gallistel, CR. Information theory, memory, prediction, and timing in associative learning. In: Moustafa, A., editor. *Computational models of brain and behavior*. New York: Wiley/Blackwell; 2017.
- Wilson G, Mattell MS, Crystal JD. The influence of multiple temporal memories in the peak-interval procedure. *Learning & Behavior*. 2015; 43(2):163–178. [PubMed: 25788176]
- Wilson RC, Nassar MR, Gold JJ. Bayesian online learning of the hazard rate in change-point problems. *Neural Computation*. 2010; 22:2452–2476. [PubMed: 20569174]
- Yamazaki T, Tanaka S. Computational models of timing mechanisms in the cerebellar granular layer. *The Cerebellum*. 2009; 8:423–432. [PubMed: 19495900]
- Zhou WG, Hohmann AG, Crystal JD. Rats answer an unexpected question after incidental encoding. *Current Biology*. 2012; 22:1149–1153. [PubMed: 22633809]

Fig. 1.

Test environment and switch protocol. State 1: the Trial-Initiation Hopper is illuminated and the flanking hoppers are not. A poke into the illuminated Trial-Initiation Hopper initiates State 2, in which the flanking hoppers are illuminated and the Trial-Initiation Hopper is not. The mouse learns to go first from the Trial-Initiation hopper to the Short-Latency Feeding Hopper (Arrow 3) and then on to the Long-Latency Feeding Hopper (Arrow 4) if and when poking in the short-latency hopper fails to release a pellet when the short feed latency has elapsed.

Fig. 2.

The lower panel shows contour plots of the likelihood functions for normal distributions parameterized by their means and cvs and given the second half of the latency data from Condition 1 of Subject 1 (upper right set of ovals) and from Condition 2 (lower left set of ovals). The contours are for the first, third and fifth natural log unit decreases in likelihood. Thus, the innermost contour delimits the parameter values that fall within 37% of the maximum likelihood, while points outside the outermost contour represent parameter combinations that are more than 150 times less likely than the likelihood-maximizing combination. Summing along the Mean axis gives the marginal likelihood functions for the cvs of the two distributions (upper panel). The dashed vertical lines indicate the maximum likelihood values. The fact that the marginal likelihood functions for these two cvs do not overlap indicates that the evidence for a difference in cv is “decisive.”

Fig. 3.

Graphic illustration of the computation of the Bayes Factors for the three cross-condition cv comparisons for Subject 1. The solid curve is the null prior; it represents our expectation if the change in condition does not produce a change in cv. The low-lying, sparsely dashed curve is the alternative prior; it represents our expectations if the change in conditions does change the cv in one direction or the other. It spreads the unit mass of prior probability from $\mu_n - 0.1$ to $\mu_n + 0.1$, where μ_n is the mean of the null prior. The prior distributions are plotted against the left axis, as indicated by the arrows in the top plot. The densely dashed curve is the marginal likelihood function for the cv given the postchange data. It is plotted against the right axis, as indicated by the arrow in the top plot. Intuitively, the question is, which prior distribution better predicts this likelihood function. The question is answered by taking the dot products between the likelihood function and each of the prior distributions. The Bayes Factor is the ratio of the two dot products. Intuitively, the extent to which one conclusion (one hypothesis) is favored by the evidence is measured by how much more prior probability it puts under the likelihood function. In the middle panel, the null prior clearly puts more prior probability under the likelihood function, which is why these data favor the no-change conclusion by almost 10:1 ($1/0.11$). In the top and bottom panels, the alternative prior puts nontrivial prior probability under the likelihood function, while the null prior puts essential none, which is why these data favor the change conclusion by “decisive” Bayes Factors (1108 and 660).

Fig. 4.

Switch latencies for each subject across the three conditions. The heavy horizontal lines mark the temporal goalposts. The vertical black lines mark the transitions between conditions. In any condition, a switch at a latency above the upper line (a switch too long delayed on a long trial) was not reinforced; at any latency below that line, it was reinforced. On short trials, a switch at a latency below the lower line (a premature switch) was not reinforced. Above the plots are the results of a statistical analysis for the three between-condition differences in the cv (first 4–12 vs. 4–8, 4–8 vs. second 4–12, and first 4–12 vs. second 4–12). The condition-bracketing lines terminating in open down-pointing arrows indicate a particular between-condition comparison: The number in the middle of each condition-bracketing line is the Bayes Factor for the hypothesis that the cvs in that comparison differed. When the Bayes Factor is greater than 3, it indicates substantial evidence for a difference in cv. In that case, the heavy up- or down-pointing arrow immediately in front of the Bayes Factor indicates the direction of the difference in cv. When the Bayes Factor is less than 1, it indicates that the data support the conclusion that there is no difference in the cvs in that comparison. In that case, there is no directional (up or

down) heavy arrow in front of the Bayes Factor. Only the second half of the switch latencies in the initial condition were included in the statistical analyses. Latencies longer than 12 s were excluded from the Bayes Factor computations, except for S3, where the exclusion criterion was >18 s.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Fig. 5.

The condition-by-condition empirical cumulative distributions of the switch latencies. All the latencies are included in these plots, but the upper limit on the x axis is set to 15. The deviation between 1 (the top of a panel) and the height of a cumulative distribution at a given latency (e.g., the 12 s latency) gives the fraction of the latencies excluded from statistical analysis by using the given latency as the criterion for excluding outliers.

Fig. 6.

The cumulative distribution of only those latencies used for the statistical analyses, that is, after the exclusion of outliers. Semilog plots. On each plot are the factors by which the means and cvs differ between the conditions. The effects on the means were significant by a t test with p values $< .005$ in every comparison but two: In S3, the difference between the two 4–12 conditions was not significant ($p = .16$), while in S4, the difference between the 4–8 mean and the second 4–12 mean was $p = .04$. The strength of the evidence for the effects on the cvs was assessed by Bayes Factors, which are given at the tops of the plots in Figure 4.

Fig. 7.

Between-condition differences in the cvs versus differences in the means. There is no correlation; that is, the difference in the means does not predict either the direction or the magnitude of the difference in the cvs.

Fig. 8.

Cumulative distributions of the likelihood-maximizing parameters for the normal, the gamma and the Wald (inverse Gaussian) distributional forms. Experimental condition varies across columns; the distributional parameter varies across rows. The mean is a natural parameter for the normal and the Wald distribution. Notice the extreme range of likelihood-maximizing values for the Wald lambda parameter (third row). Notice the high values for the gamma n parameter (also known as the shape parameter, fourth row).

Fig. 9.

The cumulative distributions of maximum likelihoods for the normal, gamma and Wald distributions. Row 1: first 4–12 condition; Row 2: 4–8 condition; Row 3: second 4–12 condition. The outlier cutoff criterion varies by column. The lower portions of the dashed plots lie to the left of the other two, toward lower maximum likelihoods, indicating that Wald fits are often much poorer than the other two. The plots for the Normal (thin black lines) and the gamma (broad gray lines) are hard to distinguish in many panels because they approximately superimpose, indicating that the gamma and the normal yield equally good fits.

Fig. 10.

Feed-latency probability density functions that differ in the location of the variation (short, long or both) and the magnitude of variation. At the highest level of variation ($cv = 0.35$), the actually programmed long feed latency on a given trial was sometimes shorter than the nominal value of the short feed latency and the actually programmed values for the nominally short feed latency ranged over the entire interval from 0 to the nominal value of the long feed latency.

Fig. 11.

Quantiles (5%, 25%, 50%, 75% and 95%) of the short feed latency distributions (dotted verticals) and of the long feed latency distributions (solid verticals). **A.** With no jitter. The quantiles collapse together because the feed latencies are strongly concentrated at the programmed latencies. **B.** With a cv of 0.35 in the programmed short feed latency and no jitter in the long. **C.** With a cv of 0.35 in the programmed long feed latency and no jitter in the short. **D.** With a cv of 0.35 in both programmed latencies. The shortest 5% of the experienced long feed latencies overlapped with the longest 5% of the experienced short feed latencies in this condition.

Fig. 12.

Cumulative distributions of the switch latency cvs under varying levels of trial-to-trial jitter in one or both programmed feed latencies. Jittering the goalposts increased the precision in timing of switches (moved the cumulative distributions to the left toward lower cvs) rather than decreased it. The other cumulative distributions all lie well to the left of the solid curve. The top panel is for the condition where only the long feed latency was jittered, the middle panel for when only the short feed latency was jittered, and the bottom panel for when both latencies were jittered.

Fig. 13.

Quantiles of the distributions of short feed latencies (dotted verticals) and of the long feed latencies (solid verticals) in Experiment 2B. **A.** With no jitter. **B.** Jitter added to the short feed latencies. **C.** Jitter added to the long feed latencies. **D.** Jitter added to the short-lever release latencies that exceeded 3s.

Fig. 14.

Representative sequences of switch latencies (small open black circles) across the three switch phases from one rat in each group. The solid horizontal lines are the nominal (mean) positions of the temporal goalposts.

Fig. 15.

Scatter plot of the cvs in Phases 2 (filled symbols) and 3 (open symbols) when jitter was added, plotted against the cv from the same subject in Phase 1 when there was no jitter. The dashed diagonal is the identity. Points below the identity indicate a lower cv in a jitter condition than in a no-jitter condition.

Fig. 16.

The tri-linear transition function used to estimate T , the transition latency in trials. The estimated values for t_s , the start of the transition, and t_f , the finish of the transition, were those that maximized the probability of the switch latency data given the previously computed estimates for the parameter vectors of the distributions fitted to the switch-latency data before and after the transition, θ_b and θ_a ; t_c denotes the trial on which we changed $p(\text{long})$.

Fig. 17.

The cumulative distributions of transition latencies following changes with different Kullback-Leibler divergences. The thick solid plot lies substantially to the left of the thin solid plot and the thick dashed plot, as is expected if the process that detects changes in the parameter of a hidden Bernoulli process is sensitive to the fact that the Kullback-Leiber divergence of $\text{Bern}(p = .9)$ from $\text{Bern}(p = .5)$ is greater than the divergence going the other way; hence, in principle, detectable in fewer trials.