Behavioral/Systems/Cognitive

A Model of Interval Timing by Neural Integration

Patrick Simen, Fuat Balci, Laura deSouza, Jonathan D. Cohen, and Philip Holmes

Princeton Neuroscience Institute, Princeton University, Princeton, New Jersey 08544

We show that simple assumptions about neural processing lead to a model of interval timing as a temporal integration process, in which a noisy firing-rate representation of time rises linearly on average toward a response threshold over the course of an interval. Our assumptions include: that neural spike trains are approximately independent Poisson processes, that correlations among them can be largely cancelled by balancing excitation and inhibition, that neural populations can act as integrators, and that the objective of timed behavior is maximal accuracy and minimal variance. The model accounts for a variety of physiological and behavioral findings in rodents, monkeys, and humans, including ramping firing rates between the onset of reward-predicting cues and the receipt of delayed rewards, and universally scale-invariant response time distributions in interval timing tasks. It furthermore makes specific, well-supported predictions about the skewness of these distributions, a feature of timing data that is usually ignored. The model also incorporates a rapid (potentially one-shot) duration-learning procedure. Human behavioral data support the learning rule's predictions regarding learning speed in sequences of timed responses. These results suggest that simple, integration-based models should play as prominent a role in interval timing theory as they do in theories of perceptual decision making, and that a common neural mechanism may underlie both types of behavior.

Introduction

Diffusion models can approximate neural population activity whenever the spike trains of individual neurons in a population are approximately Poisson and independent. Diffusion models—random walks with infinitesimal time steps—have been used to account both for interspike times in individual neurons (Gerstein and Mandelbrot, 1964) and for response times and choice probabilities in psychological research on decision making (Ratcliff and Rouder, 1998). They have not, however, been widely used to model interval timing abilities.

Here we show that a specific form of diffusion model arises from simple assumptions about neural integration, noise, and cortical wiring; that this model accounts for the full shape of the time-estimate distribution in classic timing tasks; that it is well suited to learning new durations instantly; and that behavioral data from a novel timing task rule out any model unable to learn new durations from a single exposure. The model posits that to time an interval, average firing rates in neural integrator populations ramp up linearly over time in response to stochastic inputs with a constant mean. This mean, which determines

the ramp rate, is tuned so that integrator activity hits a fixed threshold level at the right time. Thresholds, starting triggers, and variable spike rates can easily be implemented within the same, simplified neural modeling framework that produces neural integrators.

We present the model in terms of a hierarchy of modeling levels: (1) at the lowest level, the individual neuron is modeled as a nonhomogeneous Poisson spike generator, with a rate parameter determined at each moment by a leaky integral of spikes previously received from other neurons; (2) at the next level, populations of these units use balanced excitation and inhibition to act either as integrators or bistable switches; (3) circuits of switch and integrator populations implement a timer; and finally, (4) overall circuit behavior ultimately reduces to a one-dimensional drift-diffusion approximation with adjustable drift, a single, constant threshold, and a diffusion coefficient (i.e., noise level) proportional to the square root of the drift.

The diffusion model at the top of this hierarchy is intimately related to existing psychological theories of timing, such as scalar expectancy theory (SET) (Gibbon, 1977) and the behavioral theory of timing (BeT) (Killeen and Fetterman, 1988). Mathematically, it is only a slight variation on BeT. However, as we show, this variation allows the model to be interpreted in terms of spiking neural activity in a way that is hard to reconcile with BeT. The diffusion model furthermore produces inverse Gaussian distributions of time estimates that fit a large body of data better than BeT's predicted gamma distributions (and far better than normal distributions).

Finally, diffusion models give excellent accounts of decision-making behavior and can also account for electrophysiological data in decision-making tasks (Roitman and Shadlen, 2002). They may therefore serve to unify theories of decision making and timing in terms of a common process of stochastic neural integration.

Materials and Methods

Here we describe and formalize our basic assumptions about processing in individual neurons (level 1 of the modeling hierarchy) and in cortical populations (level 2). We couple these assumptions with a simple spike-counting argument to derive a drift-diffusion model of interval timing—the "opponent Poisson" diffusion model (Eq. 4)—which occupies the top level of the hierarchy (level 4). The simplicity of this model facilitates mathematical analysis, provides an elegant description of the model's behavioral predictions, and highlights the relationship of the model to existing psychological timing models. However, it assumes a number of computational functions that require justification in terms of neural mechanisms. These include the basic functions of temporal integration and switch toggling that are implicit in spike counting.

We therefore justify the spike-counting argument, and the leap from level 2 to level 4 in the modeling hierarchy, by describing a timing circuit at an intermediate modeling level (level 3). This timing circuit—the stochastic ramp and trigger (SRT) model—implements temporal integration, a mechanistic response threshold, and an adjustable clock speed. It also accounts for the operations of resetting and learning that are needed during intertrial intervals between timed stimuli.

We end the Materials and Methods section by deriving learning rules that allow the model to learn durations from a single exposure and by defining a behavioral task called "beat the clock" that tests human learning speed. In Results, we compare the model's predictions to data in the literature and to data from the beat-the-clock experiment.

Overview of the modeling hierarchy: neural processing assumptions underlying a time-scale-invariant drift-diffusion model of interval timing

Despite a wealth of behavioral and physiological data in humans and nonhuman species, there is much disagreement about the neural basis of timing abilities. Models of timed behaviors abound, ranging from models of the conditioned eyeblink response (with response times in the subsecond range) to models of fixed-interval conditioning (with tens of seconds between peak responding) to models of circadian rhythms. These models rely on a variety of dynamic phenomena, such as oscillations (Matell and Meck, 2004; Miall, 1989); ramping quantities, both linear and logarithmic (Ivry and Richardson, 2002; Dragoi et al., 2003; Durstewitz, 2003; Reutimann et al., 2004; Shea-Brown et al., 2006; Wackermann and Ehm, 2006); synfire chains (Haß et al., 2008); weighted sums of basis functions (Grossberg and Schmajuk, 1989; Machado, 1997; Ludvig et al., 2008); and repeatable stochastic processes (Karmarkar and Buonomano, 2007; Ahrens and Sahani, 2008). For durations in the seconds to minutes range, however, data have not clearly favored any particular model.

What is arguably the simplest approach—counting up the ticks of a noisy clock (Treisman, 1963)—works well in accounting for behavioral data, and as we show, it can also account qualitatively for physiological data. Linear accumulation that arises specifically by counting the ticks of a noisy pacemaker is a basic feature of BeT (Killeen and Fetterman, 1988), and it is also consistent with SET (Gibbon, 1992). SET identifies time scale invariance (or "scalar invariance" as it is called in SET) as the primary behavioral phenomenon to be explained—a phenomenon in which response time distributions are approximately Gaussian, with a mean equal to the timed duration, and an SD in constant proportion to the mean such that rescaling the time axis causes response time distributions for different intervals to superimpose. Although normality is often violated in empirical response time data, rescaling the time axis does typically cause response time distributions for different intervals to superimpose (e.g., Buhusi et al., 2009).

Under the assumption of a single Poisson process for a clock signal, the number of ticks counted by a given time has a Poisson distribution, and the time of first passage of this summed value above a fixed threshold has a gamma distribution that approximates a Gaussian when the threshold is large (Gibbon, 1992). Different durations can be accurately timed by changing either the threshold (as in process model implementations of SET) or the tick rate (as in BeT) (Killeen and Fetterman, 1988). As the timed durations grow, however, the threshold-adjustment approach reduces the coefficient of variation (CV) of first-passage times (i.e., the SD

of these times divided by their mean); in SET, scale invariance and approximately Gaussian response time distributions arise instead from noise in a memory comparison process (Gibbon, 1992). BeT, in contrast, produces scale-invariant gamma distributions using a fixed threshold, but as we demonstrate later, the tick rate must be exceedingly slow to account for typical behavioral CVs in the range of, e.g., 0.1 to 0.3.

The proliferation of alternatives to this type of tick-counting model derives in part from difficulties in envisioning a neural implementation of simple linear accumulation (cf. Church and Broadbent, 1992). In process implementations of SET, linear accumulation is often described in abstract computational terms of buffer storage and logical/arithmetic operations (Gallistel and Gibbon, 2000; Gibbon et al., 1984), and in BeT, ticks are equivalent to state transitions in a sequence of abstract behavioral states.

We propose a ramping firing rate model that addresses these difficulties while preserving dependence on the linear accumulation of Poisson inputs over time. This model depends on a number of assumptions about neural processing, most of which have been used by other modelers in various combinations. These assumptions, and the ways in which they determine how a given modeling level emerges from the level below, are as follows.

(1) Neurons (model level 1) (Fig. 1a) can represent slowly varying quantities with a nonhomogenous Poisson firing rate code (cf. Shadlen and Newsome, 1998). This Poisson process's rate parameter $\lambda(t)$ is equivalent to a shot-noise/leaky integrator/Ornstein–Uhlenbeck process with sigmoidal squashing (f), applied to a weighted sum of input spikes from other units $I(t)$. The sequence of input spikes at times s_i is essentially a sequence of brief rectangular pulses, each of area 1, but is technically a sum of Dirac-delta functions [i.e., $\delta(t - s_i) = 0$ for $t \neq s_i$, and $\delta(0)$ equals the limit of a sequence of rectangular pulses of area 1 as the base of the rectangles shrinks to 0; thus, $\int x(t)\delta(t - s)dt = x(s)$ for all continuous functions $x(t)$]. The spike-time sequence s_i is also assumed to be Poisson.

(2) Population activity (model level 2) (Fig. 1b) can be reduced (by the assumption of independence) to a single population firing rate variable $[V(t)]$, which is also a shot-noise process. Self-excitation (k) determines the time constant of leaky integration in this shot-noise process. Longer time constants (corresponding to slower leak) and even perfect integration (zero leak) can be achieved via recurrent excitation of a population by itself. However, the tuning of this recurrent excitation must be extremely precise to achieve perfect integration and to avoid instability (cf. Seung et al., 2000). Connection strengths w_{ij} between the model's population units determine the magnitude of each spike in a Poisson spike train of rate V_j being sent from the jth population to the ith population.

(3) Feedback control mechanisms can achieve the necessary level of tuning precision. We propose one such mechanism in the supplemental material (available at www.jneurosci.org) based on computing temporal derivatives in a neurally plausible way, but other approaches also exist (e.g., Koulakov et al., 2002; Goldman et al., 2003; Goldman, 2009).

(4) To the extent that a population's inputs are not independent spike trains, the inputs can be decorrelated by a proportional level of inhibition: coincident spikes in a densely connected population can trigger compensating, inhibitory spikes that cancel the coincident excitation (cf. van Vreeswijk and Sompolinsky, 1996; Shadlen and Newsome, 1998; Renart et al., 2010). Without such a mechanism for decorrelation, strong correlations between the firing rates of different neurons in A would tend to destroy the ability of population B to act as a temporal integrator of a rate code, producing spreading synchronization of firing and subsequent resetting of individual neuronal membrane potentials that would wipe out the temporal integral of their inputs. For simplicity, we therefore assume that any excitatory drive I is balanced by an inhibitory drive γI, with γ a constant with a value <1.

(5) Uncorrelated excitatory and inhibitory spikes drive an integrator population (model level 3) (Fig. 1c) to compute a running difference of the two spike counts to time intervals. At any point in time, this difference has a Skellam (1946) distribution, which quickly converges to a normal distribution as spikes accumulate (and which equals the Poisson distribution for $\gamma = 0$). We refer to such a spike-counting and differencing process as an "opponent Poisson process."
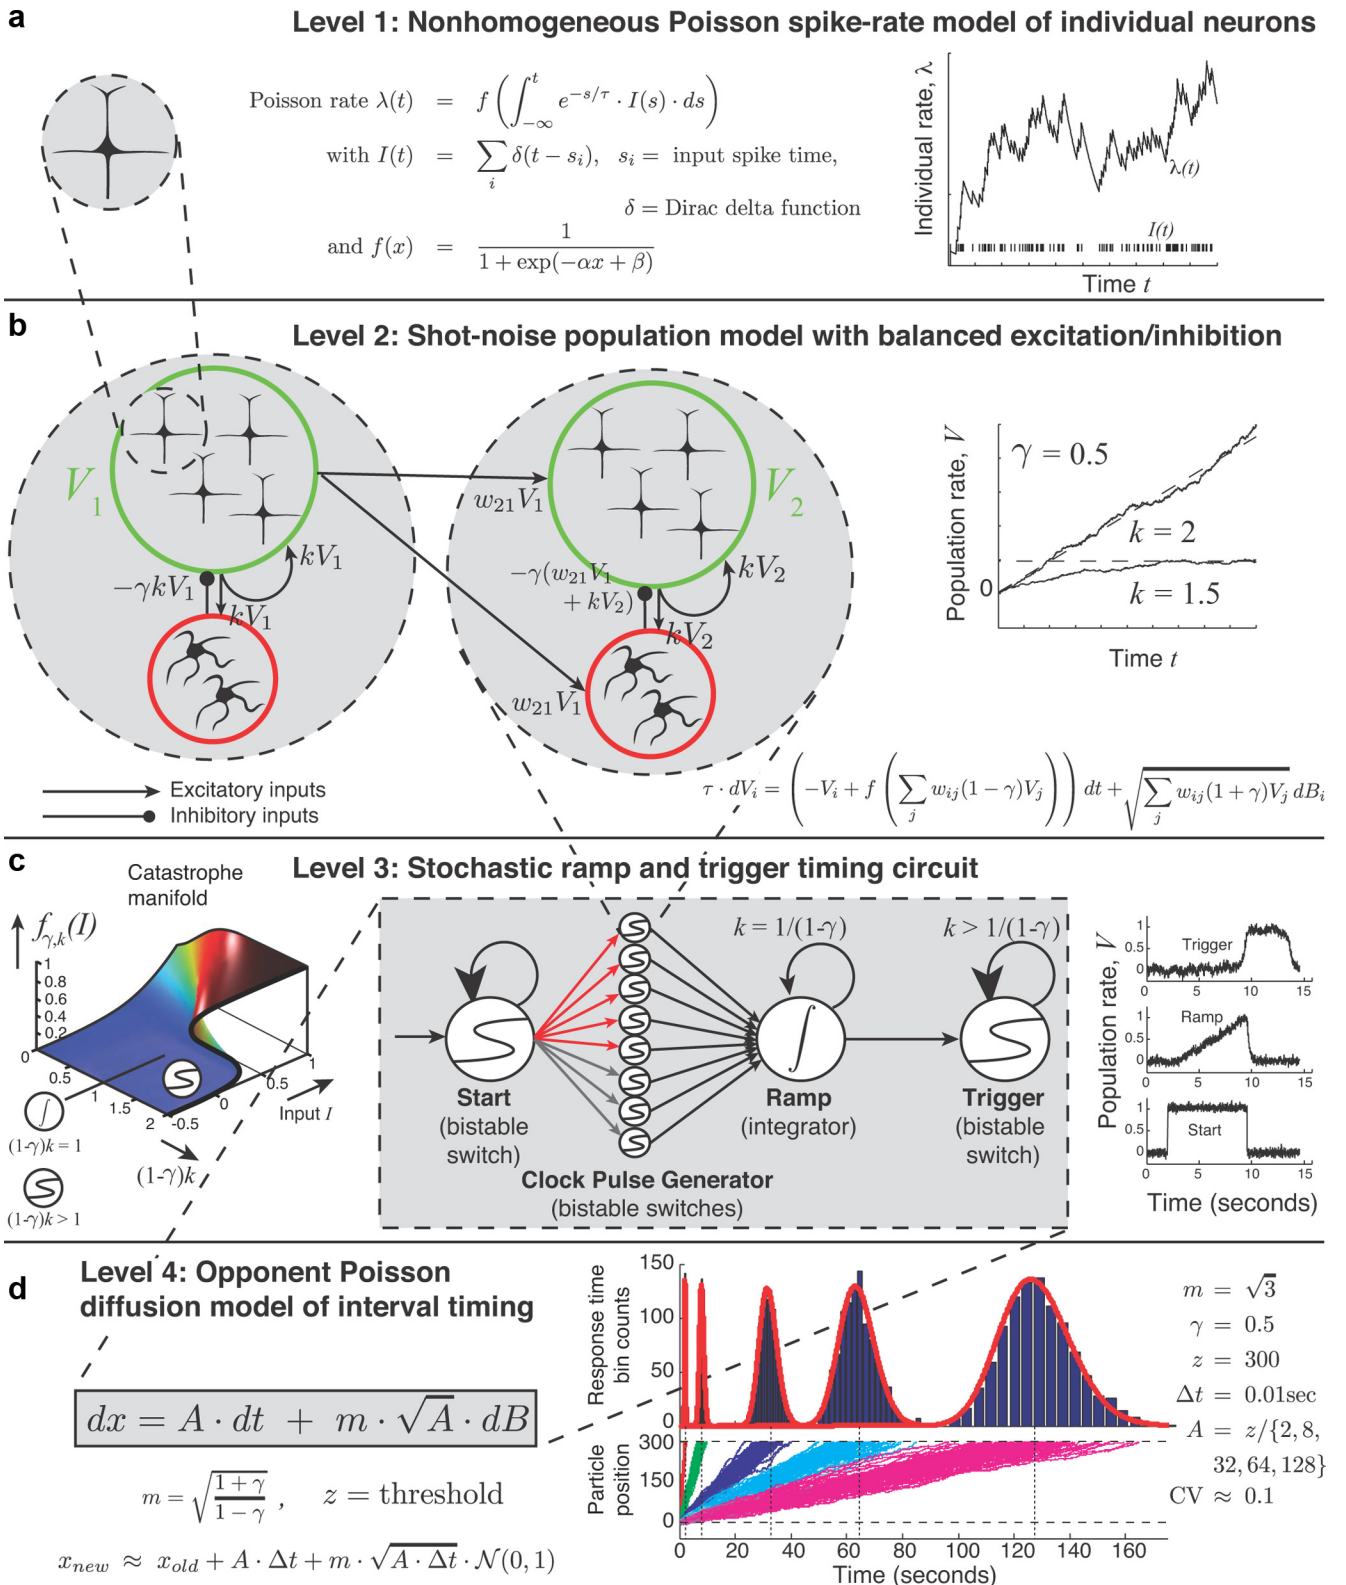
Figure 1. **a**, Individual neurons modeled as nonhomogenous Poisson spike generators with rate $\lambda(t)$. Right, Simulated response $\lambda(t)$ to input $I(t)$ over time. **b**, Subpopulations of inhibitory (red circle) and excitatory (green circle) neurons, with population firing rates equal to the sum of individual unit rates. Excitatory connections of weight w in the system also excite inhibitory interneurons to produce inhibition with a weight of γw. Self-excitation ($k \equiv W_{ii}$) determines the time constant of leaky integration. Simulated response of a self-exciting, shot-noise process to a weakly weighted, high-rate spike train input is shown at right for $k = 1/(1-\gamma) = 2$ (corresponding to perfect integration for $\gamma = 0.5$) and $k < 1/(1-\gamma) = 1.5$ (corresponding to leaky integration). **c**, The SRT model. Balanced self-excitation [$k = 1/(1-\gamma)$] produces an integrator (ramp), which accumulates clock pulses. Strong self-excitation [$k > 1/(1-\gamma)$] produces bistable switches, whose activation functions $f_{\gamma,k}$ are defined by a slice of the colored "catastrophe manifold" on the left. Switches initiate the timer circuit (start); produce stable, adjustable firing rates/clock speeds (clock-pulse generator); and produce a punctate behavioral response, or decision threshold (trigger). Duration is encoded by the proportion S of active start-to-clock weights (red arrows). On the right, time courses show start, ramp, and trigger activations during a timed interval. **d**, Opponent Poisson diffusion model. Simulations (red, green, blue, and magenta traces) are shown for 1000 trials each with threshold z at 300 and drifts A of $z/2$, $z/8$, $z/32$, $z/64$, and $z/128$, respectively ($m = \sqrt{3}$; $\gamma = 0.5$; $\Delta t = 0.01$ s). Corresponding response time distributions are shown above, as blue response time histograms (20 bins each) fit by scaled inverse Gaussians (red curves). Distributions are scale invariant (CV ≈ 0.1).

(6) When spike rates are high, the resulting difference of spike counts is approximated by a drift-diffusion process (cf. Gerstein and Mandelbrot, 1964), with drift (i.e., average tendency to rise) equal to the difference between the net excitatory and net inhibitory weighted spike rates ($I - \gamma I$). This model constitutes level 4 of the modeling hierarchy (Fig. 1d). Such a process is formally defined by the following stochastic differential equation (SDE), in which x represents the population firing rate of the integrator, t is time, A represents the net excitatory input, B (for Brownian motion) represents the white noise approximation to input variability, and $c\sqrt{t}$ is the SD of $x(t)$, assuming $x(0) = 0$ (cf. Gardiner, 2004):

$$dx = A \cdot dt + c \cdot dB. \tag{1}$$

Equation 1 can be simulated with arbitrary precision (as in Fig. 1d) by the Euler–Maruyama method (Higham, 2001). This method defines the following random walk in discrete time, with small time steps of duration Δt and random noise sampled from a standard normal distribution $\mathcal{N}(0, 1)$:

$$x_{\text{new}} = x_{\text{old}} + A \cdot \Delta t + c\sqrt{\Delta t} \cdot \mathcal{N}(0, 1). \tag{2}$$

(In all SDE simulations in the paper, approximation error is held to negligible levels by setting Δt small enough so that at least 200 time steps are typically simulated for each of the model's timed responses.) Because the variance of a difference between independent Poisson spike counts equals the sum of their variances, each of which is proportional to t, we take c to equal the square root of summed excitatory and inhibitory inputs, and we rewrite A as $(1 - \gamma)I$:

$$dx = (1 - \gamma)I \cdot dt + \sqrt{(1 + \gamma)I} \cdot dB. \tag{3}$$

Other neural models of decision making make similar assumptions (e.g., Gold and Shadlen, 2000; Eckhoff et al., 2008). For convenience, we define a constant of proportionality, $m \equiv \sqrt{(1 + \gamma)/(1 - \gamma)}$. Substituting m in Equation 3 gives a particular form of drift-diffusion model—the opponent Poisson diffusion model—in which m is constant across durations encoded by different levels of A:

$$dx = A \cdot dt + m\sqrt{A} \cdot dB. \tag{4}$$

The dynamical and statistical properties of such models are well understood, and their response time predictions have a simple and exact mathematical form as an inverse Gaussian (or Wald) distribution corresponding to first passages through a single, absorbing boundary, or threshold, $z > 0$ (Luce, 1986). Figure 1d shows simulations of such a process with noise and threshold set to give a plausible CV for humans of 0.1. [Because these simulation traces only briefly dip below 0, they show that the absence of a lower bound is unnecessary for producing a close approximation to an inverse Gaussian response time distribution when CVs are in a psychologically plausible range. As we demonstrate explicitly in Section 3 of the supplemental material (available at www.jneurosci.org), including a neurally plausible reflecting boundary at 0 would have very little impact on the shape of the response time distribution even for larger CVs.]

(7) Timed responses are emitted when the integral $x(t)$ reaches the threshold z. In the SRT model (level 3), the crossing of this threshold is detected by a "trigger" population, whose recurrent excitation is stronger than that of an integrator population, causing bistability in its firing rates (cf. Cragg and Temperley, 1955; Wilson and Cowan, 1973). Transition of the trigger population from a low to high rate triggers a behavioral response and rapid resetting of the integrator to 0 (cf. Lo and Wang, 2006). Optimal time estimation (i.e., unbiased estimation with minimal variance) requires that we take z to be fixed and deterministic, and integration to be perfect (see Results).

If our assumptions about the independence of spike trains are violated (for example, because spike-time correlations cannot be effectively cancelled), then it is reasonable to assume that the population firing rate dynamics will be better described by a diffusion process with a larger diffusion coefficient (more noise per unit firing rate), or an altogether different functional relationship between firing rates and noise levels. In

the supplemental material (available at www.jneurosci.org), we consider a spectrum of models with power-law relationships between drift and noise (with c proportional to A^r in Eq. 1) that bracket the current model: from a model in which noise does not grow at all as firing rates increase ($r = 0$) to a model in which noise grows in direct proportion to summed-input firing rates ($r = 1$). With a constant threshold, the CV of the constant-noise model increases with increasing duration, and that of the proportional-noise model decreases. In contrast, the CV of the current model ($r = 1/2$) (Eq. 4) remains constant. In the constant-noise case, an optimization argument detailed in the supplemental material (available at www.jneurosci.org) nevertheless recovers approximate time scale invariance over a limited, but potentially large, range of durations. For greater values of r ranging up to, but not including, the biologically implausible proportional-noise case ($r = 1$), approximate time scale invariance can similarly be achieved, but at a greater cost in neural resources. Thus the model is robust to variations in some of its structural assumptions.

Model level 2: neural population model

The preceding discussion motivates the use of a diffusion equation in which the noise coefficient is proportional to the square root of the drift (Eq. 4). We now complete a simplified description of neural population dynamics by including leaky integration and saturating nonlinearities (model level 2). These features take the model a step closer to biophysical plausibility, and they are critical for the real-time operation and response triggering of the complete SRT model at level 3. However, the predictions of the diffusion timing model at level 4 can be understood without reference to lower modeling levels.

Population firing rates are modeled as outputs $V_i(t)$ of leaky integrators with sigmoidal (specifically, logistic) activation functions applied to a weighted sum of inputs, $I(t)$—i.e., a standard, artificial neural network unit (Cohen and Grossberg, 1983; Hertz et al., 1991; cf. Lapique, 1907). In addition, to model the shot noise of Poisson spike train inputs in an analytically tractable form, we assume that Gaussian white noise is added to the population's inputs, which requires terms dB_i (Smith, 2010):

$$\tau \cdot dV_i = \left(-V_i + f\left(\sum_j w_{ij}(1 - \gamma)V_j \right) \right) \cdot dt$$
$$+ \sqrt{\sum_j w_{ij}(1 + \gamma)V_j} \cdot dB_i. \tag{5}$$

with

$$f(y) = \frac{1}{1 + \exp(-\alpha \cdot (y - \beta))}. \tag{6}$$

Here, $V_i(t)$ represents the average firing rate of population i; τ is the time constant; α determines the slope of the sigmoidal activation function f; and β represents the "offset voltage," the value of the input x such that $f(x) = 0.5$. Connection strengths from population j to population i are denoted by w_{ij}. These weight both the activations of afferent units and the white noise contributed by the connection. For convenience, we define k_i as the recurrent weight from a population to itself: $k_i \equiv w_{ii}$. [In supplemental material 2 (Simplified Equation; available at www.jneurosci.org), we show that Equations 5 and 6 approximate a system that more consistently treats sigmoidal squashing of both input activations V_j and interconnection noise dB_i; similar behavior of this alternative system suggests that the specific functional form of Equations 5 and 6 is not critical, as long as the population model is approximately a leaky integrator with sigmoidal squashing.]

Model level 3: SRT timing model

The neural plausibility of the opponent Poisson diffusion model of interval timing (Eq. 4) depends on whether the following critical assumptions hold: (1) spikes can be losslessly added and subtracted over time, which amounts to perfect temporal integration; (2) spike trains from different neurons are independent; and (3) fixed, deterministic thresholds can be applied to the accumulating spike count. In this section and in the supplemental material (available at www.jneurosci.org), we propose and analyze neural mechanisms that allow us to relax these assumptions

without losing the explanatory power of timing by ramping in the opponent Poisson diffusion model. The complete SRT model also includes a mechanism for integrator reset at the end of an interval, so that an account can be given of the full time course of neural activity both within and across multiple trials of a timing task. We show in Figure 2 that the SRT model explains salient features of electrophysiological data from rats and monkeys performing timing tasks.

Ramping in the linear regime. We first address assumption 1: perfect temporal integration by neural populations. This topic has received a great deal of attention in theoretical neuroscience, both because of its computational usefulness [it plays a central role, for example, in influential models of decision making (Shadlen and Newsome, 2001) and eye movement control (Seung, 1996)] and because of the difficulty that noisy neural systems ostensibly have in implementing temporal integration. The simplest hypothetical mechanism for temporal integration in neural systems parallels the conventional method for integrating in analog electronic circuits (Jackson, 1960), namely, via the cancellation of leak in a leaky integrator, or capacitor, by precisely tuned positive feedback (Seung et al., 2000). Given current knowledge about the time constants and noise characteristics of neural processing, the level of synaptic tuning precision required to cancel leak by this method is thought by many to be implausible (cf. Wang, 2001; Major and Tank, 2004).

We nevertheless use this conventional, recurrent feedback approach to temporal integration, because in the specific case of linearly ramping integrators, the necessary precision can be achieved by a relatively simple method. This feedback control method, which corrects for nonzero second derivatives of a ramp with respect to time, is discussed in the supplemental material (available at www.jneurosci.org). Other methods for achieving ramping and temporal integration (e.g., Koulakov et al., 2002; Goldman et al., 2003; Goldman, 2009) may also be suitable for implementing the linear firing-rate growth and rapid learning rules (to be discussed shortly) on which this timing model depends. The current model uses only two connection strength parameters, however, making mathematical analysis of it particularly straightforward.

Recall that V_i in Equation 5 represents the firing rate of a population, and $I_i = \Sigma_{j \neq i} w_{ij}V_j$ represents the weighted sum of excitatory inputs from other populations, minus interconnection noise. When a recurrent feedback connection strength $k_i = w_{ii}$ is added to Equation 5, the result is the following:

$$\tau \cdot dV_i = [-V_i + f((1 - \gamma) \cdot$$

$$(k_iV_i + I_i))]\, dt + \tilde{c}dB_i. \quad (7)$$

For convenience, we use $\tilde{c}$ to represent the weighted sum of the noise terms contributed by inputs and by recurrent excitation.

To simplify the analysis of ramping behavior, and without loss of generality, we can set the time constant τ to 1, and also α to 4, so that $f'(\beta) = 1$. The linearization of f around β, call it $\tilde{f}$, is then an affine function with slope equal to the identity function: $\tilde{f}(y) = y - \beta + 0.5$. Setting the feedback strength $k = 1/(1 - \gamma)$ in Equation 7 then cancels the leak defined by the $-V_i$ term. The result is the following approximation for Equation 7, in which $\tilde{A} = (1 - \gamma) \cdot I_i - \beta + 0.5$:

$$dV_i = \tilde{A}\, dt + \tilde{c}\, dB_i. \quad (8)$$
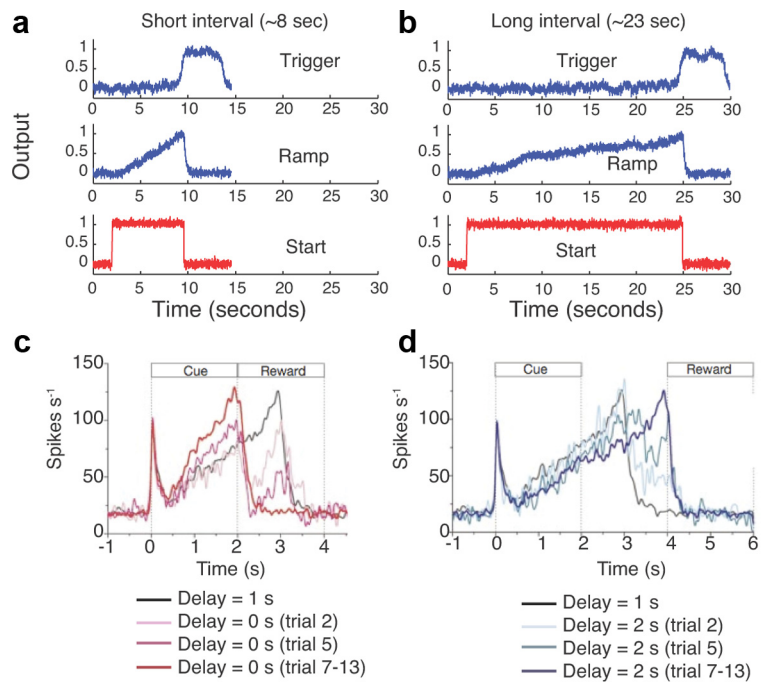
Figure 2. *a*, Time courses of start, ramp, and trigger unit activations in the SRT model (level 3) when used to time a short interval. Blue time courses are governed by system dynamics. Red time courses are input signals consisting of a 1 or 0 with superimposed white noise; inputs shut off when thresholds are crossed in these simulations (but can be parameterized to remain on even after output-triggering to support learning from errors). Hysteresis in the trigger unit leads to sustained activation following threshold crossing (the duration of this sustained activation can be parameterized as desired). *b*, A longer interval, timed using a lower clock speed (a smaller proportion S of active start-to-clock weights, colored red in Fig. 1c). *c, d*, Binned spike statistics of single, sensory thalamus neurons encoding reward delays from Komura et al. (2001), showing initial transients followed by ramps. Ramp slopes rapidly update within five to seven trials after decreases in reward delay from 1 to 0 s after cue offset (*c*) and increases from 1 to 2 s (*d*) (reprint permission pending).

If $\beta = 0.5$, then $\tilde{A}$ is identical to A in Equation 4. We will generally assume, however, that β is substantially larger than 0.5 in our ramping units, so that ramp activation will rapidly decay to 0 if the input to an integrator population suddenly goes silent. Thus, V_i is a perfect integral of $\tilde{A}$ if noise $\tilde{c} = 0$, as long as I_i remains in the range for which f is approximately linear.

Noise $\tilde{c}$, of course, cannot be assumed to be 0. In fact, we must assume under the balanced self-excitation model of integration that the spikes transmitted from an integrator population to itself should contribute to the shot noise in its own inputs. The following, modified drift-diffusion equation approximately implements this assumption; it closely mimics the behavior of a simulated, nonhomogeneous Poisson firing rate model that implements our assumptions exactly (Fig. 1*b*):

$$dV_i = (1 - \gamma)I_i \cdot dt + \sqrt{(1 + \gamma) \cdot (I_i + k_iV_i)} \cdot dB_i. \quad (9)$$

Furthermore, since $k_i = 1/(1 - \gamma)$, Equation 9 reduces to the following, with $\tilde{A} = (1 - \gamma)I_i$:

$$dV_i = \tilde{A} \cdot dt + m \cdot \sqrt{(\tilde{A} + V_i)} \cdot dB_i. \quad (10)$$

Importantly, this alteration of the opponent Poisson timing model (the inclusion of the V_i term in the noise coefficient of Eqs. 9 and 10) produces only negligible violations of the predicted properties of the simpler model (Eq. 4) when the integrator population is large. As the integrator population size goes to infinity, the contribution of V_i to the noise coefficient goes to 0. This results because both the mean and variance of the spike count in a homogeneous Poisson process with rate parameter λ over a duration t are equal to λt (Ross, 1998). If we weight each spike in a spike train by W, the expected sum of spikes over t is $W\lambda t$. If we then increase the rate parameter to $\lambda' = \lambda/W$ (with $W < 1$), the expected sum becomes λt once again, but the variance becomes $W\lambda t < \lambda t$ [since for a random variable X, i.e., the spike count, $\text{Var}(W \cdot X) = W^2\text{Var}(X) =$

$W^2\lambda't = W^2(\lambda/W)t = W\lambda t$]. Noise reduction can therefore be achieved in an integrator population by increasing the input spike rate (e.g., by increasing the integrator population size) and proportionally reducing each spike's weight (e.g., by decreasing k_i without otherwise changing Eqs. 9 and 10). In simulations, approximately scale-invariant inverse Gaussian first-passage time distributions arise for a wide range of parameter values (i.e., γ; λ, the spike rate sent to the integrator; and w_{ij}, the weight applied to these spikes before integration).

Regarding assumption 2, that correlated noise does not invalidate the model, the supplemental material (available at www.jneurosci.org) details simulations of a drift-diffusion model of timing with a large number of moderately correlated units. An argument based on response-time variance minimization suggests that even when such correlations are present, approximate time scale invariance can be achieved by averaging over large populations.

Thresholds: a latching/bistable trigger mechanism. Linear systems are unsuitable for performing a critical element of decision making. Specifically, they cannot easily be made to binarize their outputs into distinct categories that can be labeled "0" and "1." Yet, like the action potential of the individual neuron, the readout operation that triggers responses in the SRT model should render a yes/no decision about whether the value of the ramp activation is greater than a threshold value or not. The model of a latch in digital electronics offers exactly the desired properties: it binarizes its input voltage and can store the result indefinitely.

When positive feedback is sufficiently strong [$k_i > 1/(1 - \gamma)$] in a linearized version of our neural population model (Eq. 5), its activation level V_i becomes an exponentially increasing or decreasing function of time (depending on input strengths). When the squashing effect of its nonlinear activation function is taken into account, this means that activation is eventually pinned against a ceiling of nearly 1, or a floor of nearly 0 (Figs. 1c, 2a,b, "trigger" time course). Thus the system is bistable, with only transient intermediate levels of activation (cf. Cragg and Temperley, 1955; Harth et al., 1970; Wilson and Cowan, 1973). We refer to such bistable populations as "switches." Among other roles in the SRT model, switches implement the decision to respond by effectively applying a threshold to the ramping time variable in a ramp population. The "catastrophe manifold" (Jordan and Smith, 1999) in Figure 1c shows what happens to the sigmoid activation function $f_{\gamma,k}$ of a population as self-excitation k grows from the leaky range [$k < 1/(1 - \gamma)$] to the level of perfect integration [$k = 1/(1 - \gamma)$] and beyond, into the switch range [$k > 1/(1 - \gamma)$].

With a neurally implemented, bistable decision mechanism, the assumption of a fixed, deterministic threshold across trials may no longer be tenable. Instead, noisy processing in the trigger population produces decisions at different levels of ramp activation on different trials (although possibly with a constant mean value across durations). If this threshold noise is large relative to ramping noise, then the distribution of threshold values must contribute to the response time predictions. In the supplemental material (Section 4.4; available at www.jneurosci.org), we derive a response-time distribution for this more realistic model, which has approximately normally distributed thresholds across trials (cf. Lo and Wang, 2006).

The complete SRT timing circuit. We can now assemble a complete timing circuit based on the neural population model of Equation 5 (Fig. 1c, gray box). A "start" signal, modeled for simplicity as a step function of time equal to 0 or 1 in Figure 2, *a* and *b*, triggers a set of switches whose summed activations constitute an adjustable clock-pulse generator. This generator layer drives a ramping unit ("ramp") parameterized to act as much like a perfect integrator as possible [$k \approx 1/(1 - \gamma)$]. The strength of this drive determines the rate at which the ramp unit's output activation increases linearly over time (Fig. 2a,b). The ramp unit in turn excites the trigger [a bistable switch unit with $k > 1/(1 - \gamma)$] through a fixed connection strength. Below a critical level of input (Eq. 5, $\beta - 0.5$), the ramp unit decays back to 0, as shown in Figure 2a, around time equal 10 s, and at ~25 s in Figure 2b. This decay is equivalent to the integration of a strong negative input (with magnitude $0.5 - \beta < 0$) until zero activation is reached.

Similar, nearly linear ramping of firing rates followed by rapid decay has been observed in the rat sensory thalamus (Komura et al., 2001) as

well as in monkey presupplementary motor cortex (Mita et al., 2009) in interval timing tasks. Mita et al. (2009) in fact achieved a better fit to averaged firing rate data with an exponential curve than with a linear function, but their untransformed data appear reasonably linear. [A nearly exponential ramp can in any case be achieved by the current model with imperfect tuning of $k > 1$ (Shea-Brown et al., 2006).] A slowly ramping pattern is similarly observed in the contingent negative variation in human scalp potentials in timing tasks (Macar and Vidal, 2003).

The data of Komura et al. (2001) furthermore support a theory of rapid (five to seven trials) duration-learning, as shown in Figure 2, *c* and *d*. There, expected rewards following a cue are unexpectedly delivered early by 1 s (Fig. 2c) or are delayed by 1 s (Fig. 2d). Brody et al. (2003) found similarly rapid updating of firing rate ramps in monkey prefrontal cortical neurons in a vibrotactile frequency discrimination task with predictable interstimulus intervals. We address the SRT model's learning speed in the next section.

Adjustable clock speeds are achieved in the SRT model's clock-pulse layer through weights from the start unit to the clock-pulse units (Fig. 1c, red connections) that are strong enough to cause latching in clock-pulse layer units; start-to-clock pulse weights of this strength or greater will be referred to as "active." The subset of active clock-pulse populations sends inputs at a fixed rate to the ramp unit to generate a ramping activation of adjustable slope (and as we show in the next section, two simple learning rules can adapt the start-to-clock weights so as to encode new durations after a single trial). A transient burst of firing, like those seen in the rat data in Figure 2, *c* and *d*, at the cue onset, naturally occurs in more biophysically detailed spiking models of neural populations when a quiescent population receives a step voltage input (e.g., Gerstner, 1999). In contrast, drift-diffusion reductions of spiking models, such as Equation 9, are too simplified to capture this phenomenon.

Before a critical level of ramp activation is reached in the SRT model, bistability in the trigger unit prevents propagation of this ramping signal to the rest of a larger neural system in which the timer is assumed to be embedded. [Evidence for such bistability in membrane potentials exists for neurons in both the prefrontal cortex and striatum (Onn and Grace, 2000).] After the timed interval has elapsed—a duration that is assumed to be long relative to the duration of a 0–1 switch in the trigger unit—the trigger unit suddenly transitions, thereby approximating the step-function behavior of a true threshold.

The start signal can itself be implemented by a trigger population whose function is to detect cues that indicate the onset of an interval. Cue categorization is a perceptual decision-making process that can be modeled successfully by drift-diffusion models implemented with nearly identical modeling assumptions (cf. Bogacz et al., 2006). When the output trigger indicates that an interval has elapsed, it inhibits the start switch and stops the ramping process, which thereafter resets rapidly to 0. To enable the detection of responding that is too early, the start switch can also be parameterized to remain on even after the trigger switches on, and to turn off only when a cue from the environment is received.

This SRT timing scheme can naturally be extended to chains of timers, with the output trigger of one timing element serving as the start trigger of the next, or to chains of alternating SRT timers and cue-categorization processes.

Learning rules for adaptive timing

We show in Results that adapting the rate of spike accumulation and using a fixed threshold is the optimal time estimation policy for the opponent Poisson diffusion model (i.e., the policy that sets the model's expected first-passage time equal to the timed duration and minimizes the variance of the first-passage time distribution). We further show that this model fits behavioral data better than plausible alternatives. We have not yet specified, however, how this rate adaptation is to occur.

We now describe two simple learning rules that update the weights between the SRT model's start switch and its clock pulse generator in response to error feedback from the environment. Each such weight is either active, so that the corresponding clock pulse unit transitions to the "on" state when the start switch activates (in which case we assume for simplicity that the weight is 1), or it is inactive, so that the corresponding clock pulse unit remains in the "off" state at all times (in which case we

assume it is 0). Durations are encoded in the model by the proportion of active start-to-clock pulse weights (Fig. 1c, number of red weights out of the total number of start-to-clock pulse weights), which we denote by S. By the assumption of independent Poisson outputs of the clock pulse switch units, the total rate of spikes coming from the clock pulse generator is proportional to S (for simplicity, we assume that the weighted spike output of the clock pulse generator equals S). As noted previously, in the SRT timer's ramp unit, the effective drift, or "clock speed," $\tilde{A}$ resulting from a given value of S is $\tilde{A} = S - \beta + 0.5$.

With these rules, the model can learn a duration after a single exposure. The rules can be made less sensitive to individual training examples by changing weights by only a proportion of what the rules specify, resulting in asymptotic learning of durations over multiple trials. In Results, we discuss evidence for human and nonhuman animal behavior that is consistent with the rapid duration-learning that these rules predict.

Late-timer update rule. A simple rule for updating the start-to-clock pulse weights S and clock speed $\tilde{A}$ involves detecting when the timer has failed to cross threshold by the time (T) at which the target interval elapses on training trials; we refer to these as "late-timer trials." Such trials can be detected by simultaneously detecting activation in the start switch, the presence of the environmental cue signaling the end of the interval, and a lack of activation in the trigger. The information used by the late-timer update rule comprises the fixed threshold value, z, which is assumed to remain constant across trials, the current ramp level at the end of the training interval, V_T, and the clock speed, $\tilde{A}$.

From these three quantities, a weight increment ΔS can be computed instantaneously (Fig. 3a). We denote the active weight proportion S at the end of trial n as S_n. Interval durations begin at time 0 and end at T. The weight-change rule is then as follows:

$$S_{n+1} = S_n + \Delta S$$

$$\Delta S = (S_n - \beta + 0.5) \cdot \frac{z - V_T}{V_T}. \tag{11}$$

Here, β refers to the midpoint of a ramp unit's activation function in Equation 5 (an input stronger than $\beta - 0.5$ is required to cause ramp-up in the ramp unit).

The result of this rule (accurate learning in one trial) can be derived by computing the average first-passage time on trial $n + 1$, denoted t_{n+1}, assuming a late-timer trial on trial n. We denote the true duration in both trials as T (here, $\tilde{A}_n = S_n - \beta + 0.5$ indicates the slope of the ramp on trial n):

$$
\begin{aligned}
t_{n+1} &= \frac{z}{\tilde{A}_{n+1}} + \frac{z}{S_{n+1} - \beta + 0.5} \\
&= \frac{z}{S_n + (S_n - \beta + 0.5) \cdot \dfrac{z - V_T}{V_T} - \beta + 0.5} \\
&= \frac{V_T}{S_n - \beta + 0.5} = \frac{V_T}{\tilde{A}_n} = T.
\end{aligned}
\tag{12}
$$

Thus, if the timer's ramp slope is too shallow on trial n, the slope will increase so that the ramp hits the threshold at exactly T on trial $n + 1$.

Early-timer update rule. Another simple rule for updating S and $\tilde{A}$ applies to the case of early threshold crossings. Unlike the case of the late timer, the model cannot instantly correct timing errors when the ramp crosses the threshold early. Instead, it must wait for the environmental cue signaling the end of the interval before it has the information needed to correct the error. We assume that this signal will always occur at some point within the current trial—a more complex update rule would be needed to handle omissions of the end cue. Early-timer trials can be detected by simultaneously detecting activation in the start switch, the absence of the environmental cue, and activation in the trigger (the start switch can be parameterized to remain active even after the trigger activates if the environmental cue is not present, and to shut off once the cue

arrives). The information used by the early-timer update rule comprises the fixed threshold z and the current clock speed, $\tilde{A}$.

The early-timer rule adapts the active start-to-clock pulse weight proportion S continuously throughout the interval's elapsing, at time t_{early}, and the end-cue arrival at time T (Fig. 3a, third stimulus, during which the weight proportion S is continuously updated without affecting the timer's behavior until the next stimulus onset). For simplicity, we first define an update rule for the ramp slope, $\tilde{A}$, and then convert it into a weight-change rule for changing S. The update rule for $\tilde{A}$ is the following differential equation, which specifies that $\tilde{A}$ decreases continuously over the interval between t_{early} and T, starting at an initial value of $\tilde{A}(t_{early})$:

$$\frac{d\tilde{A}}{dt} = -\frac{\tilde{A}^2}{z}, \quad \text{with} \quad \tilde{A}(t_{early}) = \frac{z}{t_{early}}. \tag{13}$$

The solution of this differential equation, $\tilde{A}(t)$, converges on the correct slope, $\tilde{A}(T) = z/T$, by the end of the longer-than-expected duration. Moving the $-\tilde{A}^2$ term to the left side of Equation 13 and integrating both sides with respect to t gives the following:

$$\int_{\tilde{A}(t_{early})}^{\tilde{A}(T)} -\tilde{A}^{-2} d\tilde{A} = \int_{t_{early}}^{T} \frac{1}{z} dt$$

$$\Rightarrow \frac{z}{\tilde{A}(T)} = T. \tag{14}$$

At time T, the slope $\tilde{A}(t)$ stops changing and preserves its now correct value, $\tilde{A}(T) = z/T$, for the next timing trial. Finally, since $\tilde{A}(t) = S(t) - \beta + 0.5$, the slope-change rule (Eq. 13) becomes a weight-change rule:

$$\frac{dS}{dt} = \frac{-(S - \beta + 0.5)^2}{z}. \tag{15}$$

The performance of both learning rules is illustrated in Figure 3, b and c. Training example durations are plotted against trial number as squares. The update rules were applied to generate the predicted time on the subsequent trial; predictions are plotted as circles. Predictions of this deterministic-ramping version of the model (which captures the true model's average dynamics) had time-scale-invariant Gaussian noise added to ensure that the learning rules worked in the face of realistic levels of noise. The first 50 trials used a training duration of 10 s on average, with a small amount of trial-to-trial variability; the next 25 trials used an average 2 s duration; and the last 25 used an average 5 s duration. At the first trial of a new average example duration, the timer severely misestimates the actual duration; by the second trial, its estimate is as good as it will be at any other point during the block of constant durations (on trial 0, the prediction was arbitrarily chosen to be 100 s and is not plotted).

Although these update rules are deterministic, they still work in the face of noise in both training examples and predictions. The timer elapses early as often as it does late, so that the mean prediction time equals the mean observation time. However, substantial sequential effects should be produced under this model. Less pronounced sequential effects and more gradual learning of duration examples can be achieved simply by multiplying the right sides of Equations 11 and 15 by a learning rate constant $l < 1$ (Fig. 3c, where $l = 0.1$).

The biological plausibility of these rules is debatable. In favor of plausibility, we note that the parameters of both learning rules represent quantities such as firing rate and synaptic strength that are local to the weights being learned, or that can be transmitted by diffuse, lateral interactions among clock pulse generator units. Less parsimoniously, these rules' nonlocal parameter (the threshold value z) must be estimated at the synaptic site of learning through a slow process, e.g., by reward-guided hill-climbing over many trials. Furthermore, two different learning processes occur depending on the type of error (early or late), possibly requiring more cellular machinery than a single process would. The importance of timing to survival may nonetheless have supplied enough
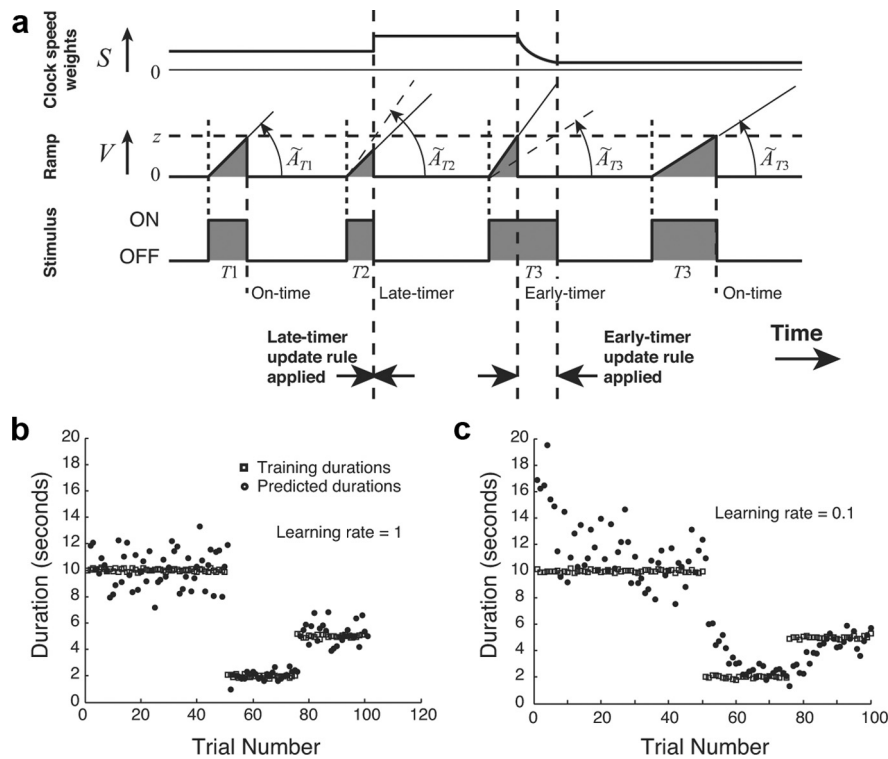
Figure 3. ***a***, Example of a learning sequence. After an accurately timed duration *T1*, a late response for duration *T2* triggers the late-timer rule and a punctate increase in the active start-to-clock pulse weight proportion *S*. A subsequent early response for duration *T3* triggers the early-timer rule and a continuous decrease in *S* until the end of the stimulus. The final repetition of *T3* is timed correctly. Top row, Ramp slope as a function of time. Note that the slope is constant whenever the timer is running. Middle row, Ramping timer. Bottom row, Stimulus. ***b***, Performance of the weight-learning algorithm illustrating one-shot learning, with a learning rate of 1. ***c***, Slower learning performance with a learning rate of 0.1.

selective pressure for such mechanisms to evolve, since these rules are the only rules by which a ramping timing model can learn accurately from single exposures. As we will show, if animals use neural integration to time intervals, then behavioral evidence strongly favors the neural implementation of rules similar to these.

Beat the clock: a timing task that tests learning speed
We will evaluate the SRT model primarily by comparing its predictions to existing behavioral data in the literature [obtained from male Sprague Dawley rats in the data sets of Church et al., (1998) and Yi (2007)]. Testing one feature of the model, however—its learning speed—requires a type of behavioral task that is less widely used, in which participants issue a single response at a time that is as close as possible to some target duration signaled by a salient cue.

For learning to take place in such a task without abstract forms of feedback (e.g., verbal or graphical indications of timing accuracy following a response), target durations must be experienced by participants. Thus, salient cues are needed to signal their onsets and offsets. Yet allowing responding to occur after a cue signaling the end of a duration makes it easy to perform such tasks without timing, simply by doing signal detection. We therefore developed a novel task, the beat-the-clock task, in which participants must respond before (but as close as possible to) the offset of the duration stimulus. At that time, reward feedback is presented whether participants have responded or not. To test learning speed, target durations are held constant within miniblocks of trials, but they change frequently and unpredictably many times per session. This task requires the anticipation and planning of a timed response to maximize monetary rewards.

The instantiation of this task that we gave to human participants involves a deadline for responding to changes in a persistent visual stimulus, a square in the center of a computer screen. At trial onset, a green square appears. If the deadline is not missed, a red border briefly appears around the square when the participant responds. At the deadline, the square

disappears and reward feedback is displayed until the next trial. Participants respond by key press to receive a monetary reward. Responses that occur after the deadline earn no money. Responses that occur before the deadline earn a reward that is an exponential function of time ($\$0.25 \cdot e^{-6 \cdot (T - t)/T}$, with t denoting current time in the trial and T the target duration), increasing from nearly $0 at the start of the interval up to $0.25 at the time the stimulus disappears. Thus, optimal (reward-maximizing) performance in this task requires responding, on average, at a time just before the deadline.

Participants were explicitly instructed not to count, tap, or adopt any rhythmicity to time the intervals. Counting strategies dramatically increase timing precision in humans (Rakitin et al., 1998), and counting could also be used to achieve rapid changes in time estimates. To prevent covert counting, we simultaneously present a distractor task. A sequence of digits is shown continuously in the middle of the viewing screen, except during reward feedback, with the digits changing rapidly and at irregular intervals (digits are visible for an exponentially distributed duration, with a mean of 500 ms, lower limit of 100 ms, and upper limit of 5 s; interdigit intervals are exponentially distributed with a mean of 50 ms and upper limit of 100 ms). Any "3" that appears requires pressing a button other than the timing response button (target "3" durations are always 2 s, to allow for sensorimotor latency). The total reward earned at the end of the task is reduced by the proportion of target stimuli not detected.

We collected data from 17 human participants (eight males, nine females) recruited from the Princeton University Psychology Department's paid experiments website. Each participant performed a 1-h-long session of the beat-the-clock task divided into 5 min blocks. Each block of trials was divided into miniblocks, within which the duration of the stimulus being timed was fixed. Intertrial interval durations were temporally unpredictable, drawn from a truncated exponential distribution (location parameter, 2 s; lower bound, 1 s; upper bound, 5 s; mean, ~2.4 s).

To prevent anticipation of a change in the target stimulus duration, miniblock lengths were gamma distributed with a mean of 7, shape parameter of 21, and scale parameter of 1/3 (additionally, miniblock lengths were truncated, with a lower limit of 4 and an upper limit of 20). Across miniblocks, the stimulus duration T performed a random walk, generated by the following algorithm for determining the *i*th block's stimulus duration:

$$T(i + 1) = T(i) + 0.8 \cdot (8 - T(i)) + 0.5 \cdot t(i) \cdot \mathcal{N}(0, 1).$$

(16)

The random walk was thus attracted to a value of 8 s, but any given miniblock's stimulus duration was likely to be 50% larger or smaller than the duration in the preceding miniblock. A single set of intertrial intervals, stimulus durations, and miniblock lengths was used for all participants, so that group performance could be examined with the same set of transitions among learned durations.

Results
We first derive analytical results for the opponent Poisson diffusion model: (1) it requires a fixed threshold, set to the maximum possible value, to minimize the variance of its time estimates; (2)
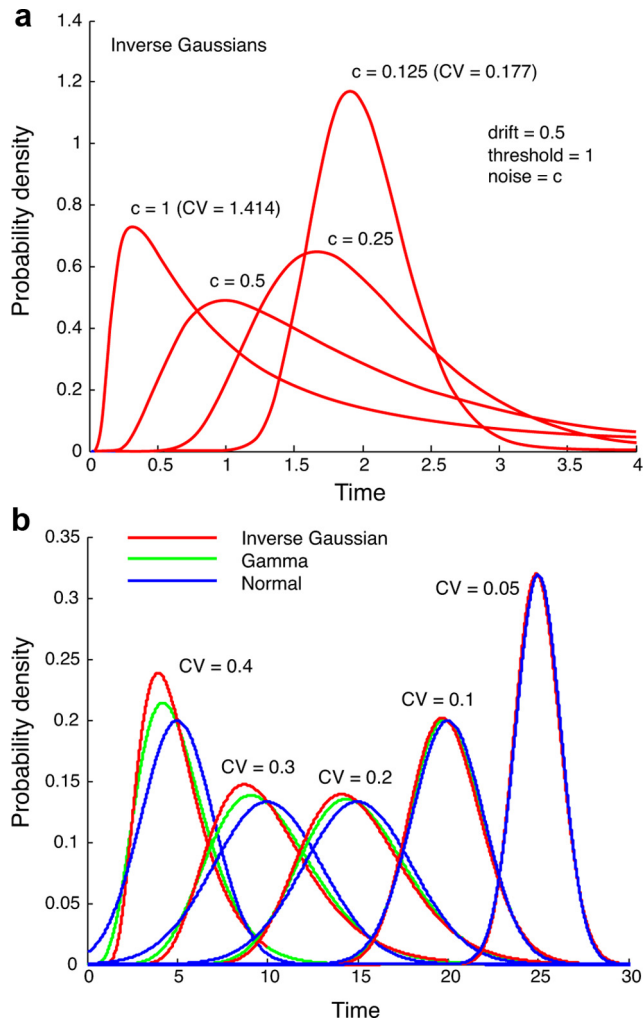
Figure 4. *a*, Inverse Gaussian probability densities generated by a one-dimensional drift-diffusion process ($dx = A \cdot dt + c \cdot dB$) starting at $x = 0$, with drift $A = 0.5$, a single threshold $z = 1$, and various levels of noise c (here c is not proportional to $\sqrt{A}$). All distributions shown have a mean of 2. As c goes to 0, this density converges on a delta function with all of its mass located at exactly 2. The CV in all cases is $c\sqrt{2}$. *b*, A comparison of inverse Gaussian (red), gamma (green), and normal (blue) distributions with identical CVs, shown for a decreasing sequence of CVs from 0.4 to 0.05 as the mean increases from 5 to 25.

parameterizing the model in this way leads to time scale invariance; and (3) the ratio of skewness to CV of the model's time-estimate distributions must always equal 3. Next we compare the model's predictions regarding steady-state response time distributions with existing rat behavioral data and other data from nonhuman species. Finally, we discuss behavioral evidence for the rapid duration learning made possible by the learning rules of Equations 11 and 15, including results from the beat-the-clock experiment.

Optimal temporal estimation by the opponent Poisson diffusion model
For a drift-diffusion model (Eq. 1) with drift A, diffusion coefficient c, and a single threshold z, the response-time density is the Wald or inverse Gaussian density (Wald, 1947; Luce, 1986). This density, defined over first-passage times t, can be expressed most conveniently in terms of the threshold-to-noise ratio, $\eta = (z/c)^2$, and the threshold-to-drift ratio, $\mu = z/A$, as follows:

$$p(t, \mu, \eta) = \left(\frac{\eta}{2\pi t^3}\right)^{1/2} \cdot \exp\left(\frac{-\eta(t - \mu)^2}{2\mu^2 t}\right). \quad (17)$$

Examples of inverse Gaussian densities corresponding to varying levels of noise but with equal means are shown in Figure 4.

The first three moments of the inverse Gaussian distribution are as follows:

$$E(t) = \mu; \quad (18)$$

$$\sigma^2 \equiv \mathrm{Var}(t) = \frac{\mu^3}{\eta}; \quad (19)$$

$$\mathrm{Skewness}(t) = E\left[\left(\frac{t - \mu}{\sigma}\right)^3\right] = 3 \cdot \sqrt{\frac{\mu}{\eta}}. \quad (20)$$

The CV of first-passage times is therefore as follows:

$$\sigma/\mu = \sqrt{\frac{\mu^3}{\eta} \cdot \frac{1}{\mu}} = \sqrt{\frac{\mu}{\eta}}. \quad (21)$$

As we now show, Equation 21 implies an important property of the opponent Poisson timing model: that optimal time estimation requires keeping the threshold fixed and adapting the drift to time different durations, which in turn implies time scale invariance.

Translating back into the model's basic parameters, A, z, and c, gives the following expression for σ:

$$\sigma = \sqrt{\frac{\mu^3}{\eta}} = \frac{z^{3/2}/A^{3/2}}{z/c} = \frac{cz^{1/2}}{A^{3/2}},$$

$$\Rightarrow \mathrm{CV} = \frac{cz^{1/2}}{A^{3/2}} \cdot \frac{A}{z} = \frac{c}{\sqrt{Az}}. \quad (22)$$

Substituting $c = m\sqrt{A}$ from Equation 4 into Equation 22 gives the following:

$$\mathrm{CV} = \frac{m\sqrt{A}}{\sqrt{Az}} = \frac{m}{\sqrt{z}}. \quad (23)$$

Optimal time estimation requires minimizing both $E[(t - T)^2]$ and $\mathrm{Var}(t)$, where T is the duration of the interval being timed. It is easy to minimize the first quantity by setting $\mu = z/A \equiv T$, making t an unbiased estimate of T. Minimizing $\mathrm{Var}(t)$ under the constraint that $z/A = T$ is equivalent to minimizing the CV, $m/\sqrt{z}$. The drift A no longer factors into this quantity, so maximizing z minimizes the CV for all durations. Plausibility requires that z is bounded above by some value $z_{\max}$. Thus, we set z to its maximum for all values of T and use the learning rules of Equations 11 and 13 to set $A = z_{\max}/T$. We thereby obtain a constant CV (an extension of Weber's law to the domain of interval timing) and superimposition of mean-normalized response-time distributions—i.e., time scale (or "scalar") invariance (Gibbon, 1977).

Whether or not the threshold is fixed at $z_{\max}$, the expressions for mean, variance, and skewness in Equations 18–20 imply that a specific relationship must hold among the first three central moments of the response time distribution for the drift-diffusion model, namely, that the skewness equals three times the CV. This strong prediction of the model might seem to conflict with a general consensus in the timing literature that timed response time distributions are approximately Gaussian (implying a skewness of nearly 0). For a small enough CV, however, a skewness-to-CV ratio of 3 implies only a very small positive skewness, and in any case, it is well known that a heavier right tail often arises in empirical response time distributions in timing tasks (Roberts,

1981; Cheng and Westwood, 1993; Cheng and Miceli, 1996; Rakitin et al., 1998). In the next section, we show that data from rats in the peak-interval task are clearly consistent with this prediction.

Before considering fits to empirical data, it is worth noting that BeT also predicts a skewed response time distribution (Killeen and Fetterman, 1988). The absence of inhibitory inputs to the accumulator in BeT and SET and the comparison of its excitatory spike count to a deterministic, integer-valued threshold results in a gamma distribution of response times (specifically an Erlang distribution) (Ross, 1998) with shape parameter κ and scale parameter θ. The rate of excitatory spikes equals $1/\theta$, and the threshold is κ. For this model, the mean is $\kappa\theta$, the SD is $\theta\sqrt{\kappa}$, the CV is $1/\sqrt{\kappa}$, and the skewness is $2/\sqrt{\kappa}$. In contrast to the drift-diffusion model, these expressions imply that skewness equals two times the CV, rather than three. [These skewness-to-CV ratio values of 2 and 3 happen to equal the respective indices of the gamma and inverse Gaussian distributions in the family of "Tweedie distributions," a subset of the exponential family of distributions that also includes the Poisson, with an index/skewness-to-CV ratio of 1, and the normal, with an index/skewness-to-CV ratio of 0 (Jorgensen, 1987).] As their respective threshold parameters go to infinity, the CV and skewness of both the gamma and inverse Gaussian distributions go to 0, and both converge to a normal distribution (Fig. 4b). However, the gamma and inverse Gaussian also approximate each other closely even for thresholds of, e.g., 100 (CV = 0.1), at which level noticeable skewness and deviation from normality are still present.

Importantly, though, it is difficult to interpret the ticks counted in BeT as individual action potentials. Since a CV of 0.1 implies $\kappa = 100$, timing a 100 s interval would require a clock spike rate of 1 Hz ($1/\theta = 1$), which seems implausibly low [indeed, BeT is presented by Killeen and Fetterman (1988) in terms of transitions between abstract behavioral states, not in terms of firing rates]. If high-rate excitatory spike trains could simply be counted up to a large threshold without any background noise present, then CVs for this model could easily be reduced by orders of magnitude below what is observed in behavior. Integrating an inhibitory spike train along with the excitatory spike train (setting $\gamma > 0$ in Eq. 3), however, gives the model enough flexibility to fit behavioral data while accommodating realistically higher spike rates.

Opponent Poisson diffusion model fits to behavior

We now show that maximum likelihood fits of inverse Gaussian distributions to timing data from rats are better than fits of gamma and normal distributions. We further show that skewness-to-CV ratios are closer to the predicted 3 of the inverse Gaussian distribution than the predicted 2 of the gamma distribution and the predicted 0 of the normal distribution. Previously published behavioral data from rats (Church et al., 1998) were obtained using the peak-interval procedure (described below). We also examine findings from other species that show evidence of timing by drift diffusion.

Peak-interval task performance in rats

In the peak procedure (Catania, 1970), subjects are presented with a randomly interleaved sequence of reinforced, discrete fixed interval (FI) trials and nonreinforced "peak" trials that last much longer than the FI trials. In the FI trials, subjects are reinforced for their first response after the fixed interval elapses relative to the onset of a cue, whereas no reinforcement is delivered in peak trials. When responses are averaged across many peak trials, the rate of responding (e.g., lever pressing) increases smoothly as a function of time into the trial, reaches its peak at around the

time of reinforcement availability, and then smoothly decreases thereafter. The resulting response rate distribution, or peak response curve, is typically approximately Gaussian, often with a slight positive skewness (e.g., Roberts, 1981).

Church et al. (1994) showed, however, that the response rate in individual peak trials is not a smooth function of time. In individual peak trials, responding instead follows a "break-run-break" pattern, in which subjects abruptly increase the rate of responding midway through the usual interval (the start time) and abruptly decrease their response rates after the interval elapses without reward (the stop time). "Middle times" are defined as the times halfway between start and stop times; these approximate the times at which response rate curves reach their peaks. Analyses of middle time distributions are less vulnerable to averaging artifacts than analyses based on raw response rate curves, which can be contaminated by numerous sources of noise. Across trials, the shapes of start, stop, and middle time distributions are similar to the response rate curves themselves, but are more positively skewed.

To evaluate how well the opponent Poisson model's predicted inverse Gaussian distribution fits rat behavioral data, we analyzed the last 10 sessions of the previously published peak-interval data of Church et al. (1998), in which separate groups of five rats received extensive training (50 sessions) with one of three different intervals: 30, 45, and 60 s. We estimated the trial-based expected time of reinforcement using the middle times (Church et al., 1994). Start and stop times were estimated using a relative likelihood change point detection algorithm, which decides at each point in time whether an increase in the response rate has occurred. We used one liberal (Bayes factor, 10) and one conservative (Bayes factor, 100) decision criterion for change point detection (for details, see Balci et al., 2009b); the midpoint estimates were essentially the same regardless of the criterion. Similar results were obtained with the ordinary least squares–cumulative sum procedure used by Church et al. (1994). In calculating these values, we adopted exclusion criteria commonly used in the timing literature; specifically, we excluded trials in which the start time exceeded the FI, the stop time was less than the FI, or the stop time exceeded three times the FI (cf. Cheng and Miceli, 1996; Balci et al., 2008a). We compared fits of Gaussian, gamma, and inverse Gaussian distributions to the data, and we further estimated the CV and skewness of these middle times (computed from the empirical data) to test whether the skewness-to-CV ratio equaled the predicted value of 3.

Figure 5a shows the distribution of middle times for 30, 45, and 60 s intervals, with middle-time CVs approximately equal to 0.2 across all durations. Fitted inverse Gaussians are superimposed on the empirical histograms. Inverse Gaussians fit the data better than the gamma and normal distributions (regardless of the change point algorithm's decision criterion), as demonstrated by the larger log likelihood values for the inverse Gaussian in Table 1 (all three models have the same number of free parameters, so they are comparable in terms of model complexity). Given these log likelihoods, the inverse Gaussian is more likely than the gamma, which is much more likely than the Gaussian. Unfortunately, because $\mu/\eta = (1 + \gamma)/(1 - \gamma)/z$, any increase in z can be compensated by a suitable increase in γ. Thus, fits of the inverse Gaussian do not uniquely specify values of γ and z in Equations 3 and 4.

Figure 5b shows the superimposition of the fitted inverse Gaussian distributions after rescaling the time axis by the average middle time. The near-perfect match indicates the time scale in-

variance of the middle times (Church et al., 1994, their Fig. 9, and Church et al., 1998, their Fig. 4*b*, which plot middle times that we also find are better fit by inverse Gaussians than by normal or gamma distributions). The superiority of the inverse Gaussian fits coupled with clear evidence of time scale invariance is also consistent with an opponent Poisson diffusion model of timing in which the threshold is constant and drift is adapted across conditions, since the CV of such a model is inversely proportional to the square root of the threshold (Eq. 23).

These results are sufficient for selecting the inverse Gaussian over the normal and gamma. Statistical analysis is needed, however, to determine whether skewness is the feature of the data that gives the inverse Gaussian its advantage. We found that this was the case. When averaged across different durations and change point decision criteria, the estimated skewness-to-CV ratio of the data was close to 3 (mean, 2.57; SEM, 0.64). This ratio was significantly different from 0 (which holds for the normal distribution; $t_{(14)} = 3.99$; $p = 0.0013$) but was not significantly different from the predicted ratio of 2 for the gamma ($t_{(14)} = 0.887$; $p = 0.39$) or from the ratio of 3 for the inverse Gaussian ($t_{(14)} = 0.67$; $p = 0.52$). The same results held when the analysis was conducted separately with different change point decision criteria.

Our analysis of a separate published data set (Yi, 2007), in which 12 rats were tested on the peak procedure with 60 and 120 s intervals corroborated these findings. When averaged across different durations and decision criteria, the skewness-to-CV ratio was 3.07 (SEM, 0.53) for this data set. This value was significantly different from 0 ($t_{(11)} = 5.82$; $p < 0.001$), trended toward a significant difference from 2 ($t_{(11)} = 2.23$; $p = 0.067$), and was not significantly different from 3 ($t_{(11)} = 0.13$; $p = 0.90$). Similar results have also been observed in other data sets and using different start- and stop-point algorithms (Church et al.,1994) on the same data set (D. Freestone, personal communication). Thus, skewness appears to be the feature of the data that distinguishes the fit quality of these distributions: the gamma can accommodate a larger skewness for a given CV than the Gaussian, and the inverse Gaussian in turn can accommodate a larger skewness-to-CV ratio than the gamma, and this ordering echoes that of the model-fit qualities.

Inverse Gaussian response time distributions across nonhuman species
Across species, when individual, timed responses are analyzed (i.e., middle points rather than the entire peak response curve), we find that the literature contains many more examples of inverse Gaussian-distributed response times across species. For instance, in a task used by Brunner et al. (1992), starlings invested foraging time on a patch that delivered reward with fixed inter-reward times. In an unpredictable manner, the patch was fully depleted (i.e., rewards were no longer available), and in those cases subjects were observed to "give in," i.e., to stop responding at a patch, at some point after the fixed interreward interval had elapsed without reward delivery. The decision to give in could only be made by tracking the elapsed time. In line with our analysis of peak-interval middle times, the distributions of such giving-in times were found to be inverse Gaussian (Gibbon and
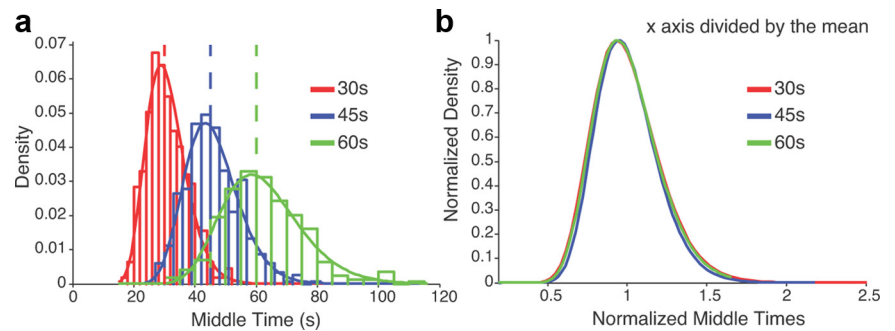
Figure 5. Inverse Gaussian fits of middle times of rat response time distributions in the peak-interval task. Middle times (midpoints of distinct periods of high-rate responding in individual peak trials) are correlated with overall mean response times but may provide a better measure of an animal's underlying reward-time estimate. *a*, Empirical distributions (histogram) and fitted inverse Gaussian densities for behavior in 30, 45, and 60 s schedules. *b*, Time scale invariance is demonstrated in plots of fitted densities against a normalized time axis, with time divided by the mean middle time.

Table 1. Fits of Gaussian, gamma, and inverse Gaussian distributions to middle times of rat behavioral data from the peak-interval procedure, indicating clear superiority of the inverse Gaussian fits

	Gaussian	Gamma	Inverse Gaussian
30 s	log likhd., −1803.46	log likhd., −1787.62	log likhd., −1784.44
	$\mu_N = 30.91$	$k = 22.98$	$\mu = 30.91$
	$\sigma = 6.55$	$\theta = 1.35$	$\eta = 692.70$
45 s	log likhd., −1812.96	log likhd., −1801.06	log likhd., −1799.42
	$\mu_N = 45.90$	$k = 27.65$	$\mu = 45.90$
	$\sigma = 8.84$	$\theta = 1.66$	$\eta = 1239.40$
60 s	log likhd., −1381.05	log likhd., −1371.13	log likhd., −1370.27
	$\mu_N = 62.47$	$k = 23.40$	$\mu = 62.47$
	$\sigma = 13.12$	$\theta = 2.67$	$\eta = 1417.48$

Larger log likelihood (likhd.) values indicate better fits; all models considered have equal complexity (two parameters). Parameters are as follows: Gaussian, mean μ_N, standard deviation σ; gamma, shape k, scale θ; inverse Gaussian, mean μ, shape η.

Church, 1990) and time-scale-invariant. Approximately inverse Gaussian start, stop, and middle times have also been observed in other studies (e.g., Church et al., 1998; Matell and Portugal, 2007), and the peak response curve itself appears inverse Gaussian, with a large CV and extreme positive skewness, in lower organisms such as goldfish (cf. Drew et al., 2005).

Despite these consistent patterns, however, the shape of behavioral response distributions has usually been dismissed in modeling efforts, since positive skewness in empirical data can arise from multiple sources of response time variability in a model (for discussion that recommends ignoring skewness for these reasons, see Gibbon and Church, 1990). Both the Poisson counter and opponent Poisson diffusion models of timing, however, make strong a priori predictions about skewness that can be tested unambiguously. Thus, distributional shapes clearly have the power to distinguish among distinct but closely related hypotheses about the neural basis of timing, at least when behavioral CVs are 0.2 or greater.

We have now seen that a single, inverse Gaussian distribution with different parameter values can account for timing behavior across species. Relatively smaller values of η—corresponding to larger values of c in Equation 1—are found for rodents compared to humans, and response rate curves for goldfish require still smaller η values. These differences in precision may have resulted from differences in the task performed, from innate differences in the level of noise c perturbing neural processing in each species, or from a combination of both factors.
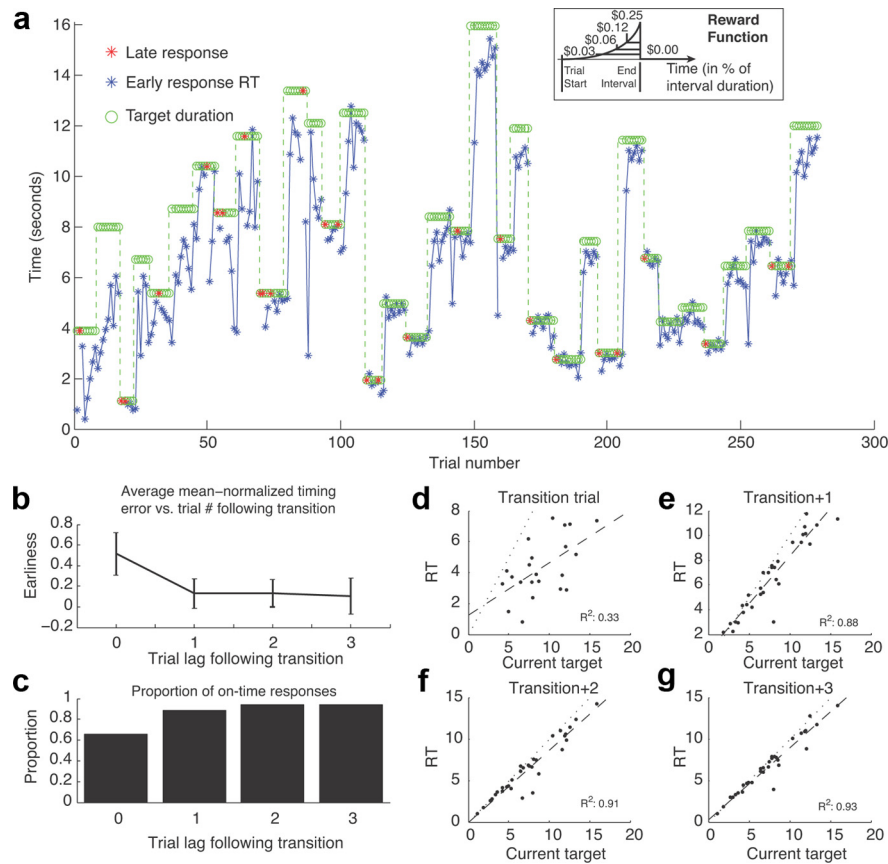
Figure 6. Individual human performance in the beat-the-clock task. **a**, Stimulus durations (green circles), timed responses (blue stars), and late responses (red stars). **b**, Average, mean-normalized timing error on trials 1, 2, 3, and 4 of each miniblock. Error bars indicate the SD of the earliness error. **c**, Proportion on the same trials of failures to respond on time. **d**, Scatterplot of response times as a function of stimulus duration on the first trial of each miniblock. The identity line is dotted; the best linear fit ($R^2 = 0.33$) is dashed (on this trial, linearity is an artifact; all points must be below the identity line since participants cannot make responses later than 200 ms after deadline). **e–g**, Response times versus stimulus duration on the second, third, and fourth trials of each miniblock, respectively. R^2 asymptotes at around 0.9 one trial after miniblock transition.

Evidence for rapid learning of duration in rodents and pigeons

Evidence for rapid learning of temporal intervals has been demonstrated in rats under aversive conditioning paradigms after a single delivery of a shock (Davis et al., 1989; Bevins and Ayres, 1995). For instance, Bevins and Ayres (1995) varied the latency of a single shock for different rats in contextual fear conditioning and observed that the peak and spread of the distribution of freezing responses increased as a function of experienced shock latency.

Experiments using positive reinforcement have demonstrated abrupt, if not immediate, adjustments to manipulations of temporal task parameters. Meck et al. (1984), for example, showed that rats adjusted the time of maximum-rate responding abruptly in two steps rather than gradually when a fixed interval to reward was changed from 10 to 20 s and from 20 to 10 s in the peak-interval procedure (Catania, 1970; Roberts, 1981), a pattern observed only in individual subjects' data.

Another form of data suggesting rapidly updated, temporally controlled behavior involves wait time: the time spent waiting before resumption of responding after consuming a reward. Wynne and Staddon (1988), for example, showed that pigeons set wait time durations to a fixed proportion of the scheduled inter-food interval, and that they can update wait times within a single interval after a change in interval duration. In rats, Higa (1997)

showed asymptotic convergence on a new waiting time over approximately five trials.

While these data all suggest an ability to update temporal representations rapidly, a more powerful test of learning speed would use a much larger number of target interval durations, and transition unpredictably among these target durations after each is held constant for a relatively small number of trials. The beat-the-clock task introduced in Materials and Methods has these properties, and humans do show evidence of extremely rapid learning in this task, as we now describe.

Human beat-the-clock performance

All 17 participants in the beat-the-clock task showed evidence of extremely rapid learning. Figure 6a shows plots of response times as a function of trial number for a typical participant (blue points, with connecting lines indicating consecutive sequences of on-time responses for one participant). These are superimposed on indicators of the stimulus duration assigned to the miniblock (green circles). Late trials are indicated with red stars inside the circles. The on-time responses closely track the actual stimulus durations, as shown in Figure 6b. There the average, normalized timing error for on-time responses (i.e., timing error divided by stimulus duration) is plotted as a function of the trial number into a new miniblock. Figure 6c shows the proportion of failures to respond on time (including omitted responses and late responses within 200 ms of the deadline) as a function of the same trial numbers. Significant improvements in timing error and on-time proportion occurred in the first trial for all participants, and on the second for a subset of participants (Wilcoxon ranksum test, $p < 0.05$), but not thereafter. The great majority of this improvement occurred in the first trial following a transition.

Figure 6d shows scatterplots of the response time on a given trial versus the stimulus duration in the current miniblock. On the first trial of the miniblock, there is no obvious relationship between response times and stimulus durations. By the second trial of the miniblock (Fig. 6e), however, the response times fall close to the identity line. Thus, one trial sufficed for timing performance to reach asymptote for some participants, and even the lowest-earning participants reached asymptote by the third trial of a new miniblock. Across participants, the mean R^2 values for trials 1 to 4 after the start of a miniblock were, respectively, 0.41, 0.72, 0.79, and 0.80. After Bonferroni correction for multiple comparisons, the R^2 value on trial 1 was significantly different than in all other trials, but no other pairwise differences were significant at $p < 0.05$. Similar results held for normalized earliness and proportion of late responses on each trial.

All of these improvements occurred despite the secondary task designed to prevent counting (described in Materials and Methods). Preventing explicit, subvocal counting is critical, since
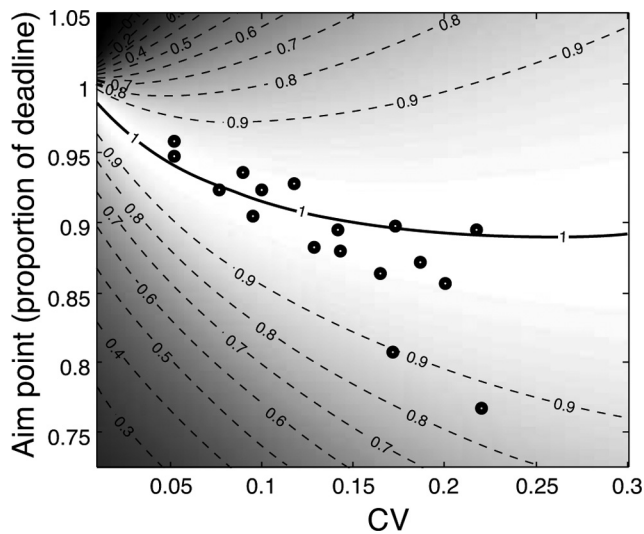
Figure 7. Optimality analysis in the beat-the-clock task. The estimated aim point, as a proportion of the deadline duration, is plotted with an open circle for each participant as a function of individual CV. The reward-maximizing aim proportion is denoted by the solid curve (the ridge of the expected gain surface). For each value of CV, the grayscale background and contour lines represent the proportion of maximum reward expected, evaluated at every aim point.

rapid changes in time estimates could be achieved by counting alone, rather than by adapting an implicit timing mechanism. Although it is possible that some participants refused to follow our requests not to count and managed to do so despite the secondary distractor task, pilot data from two of the authors (F. Balci and P. Simen) who did not count showed the same pattern of rapid learning. Indeed, all participants showed similar learning patterns, suggesting that explicit counting is unlikely to provide a general explanation for our results.

The data also show that participants timed their responses so as to maximize rewards in the task. Assuming that response times are scale invariant, the optimal (i.e., reward-maximizing) aim point for a participant is a fixed proportion of the interval duration. This proportion is closer to 1 for participants with smaller CVs, who therefore have a smaller chance of missing the deadline than those with larger CVs, given the same aim point. For each participant, the mean on-time response time was indeed close to a fixed proportion of the interval duration, as shown by the high R^2 values for response times versus interval durations. Empirical aim points were estimated from the on-time response times, normalized by the target duration, using the mean of the best-fitting inverse Gaussian distribution, truncated to exclude probability mass beyond the deadline. These aim points were then compared to the optimal aim points for the corresponding CVs. As Figure 7 shows, all but two participants come remarkably close to aiming at the reward-maximizing proportion of each interval (the small size of beat-the-clock data sets and the artificial negative skewness imposed on the response time data by the deadline cutoff combined with frequent changes in target duration make it difficult for any distribution with zero or positive skewness to fit these data well, but aim points are similarly close to their optimal values under fits of truncated normal and gamma distributions). These results are consistent with other studies of optimal temporal risk assessment (e.g., Hudson et al., 2008; Balci et al., 2009a; Jazayeri and Shadlen, 2010), but they extend the durations timed by an order of magnitude relative to these studies.

Supplemental Figure 12 (available at www.jneurosci.org as supplemental material) shows that the model behaves in the same

way. Supplemental Figure 11 (available at www.jneurosci.org as supplemental material) demonstrates that both human participants and the model produce similar, positive autocorrelations of response time out to lags of five trials. Since late responses are extremely costly, we parameterized the model to aim for a time that was short of the deadline by 1 SD of the set of first-passage times (the reward-maximizing aim point differs from this value slightly, but only a relatively small amount of reward is lost by this simple choice). To time the interval, the learning rules were applied on every trial. Biasing toward early responding was then accomplished by reducing the threshold for responding (but without changing the threshold z used in the learning rules to estimate the deadline durations).

The learning speed exhibited by human participants in this task poses a challenge for any model of implicit interval timing that does not possess an easily implemented, one-shot learning rule capable of perfect accuracy (on average). The opponent Poisson diffusion model, however, can perform with the speed and accuracy that the data demand, using extremely simple learning rules.

Discussion

We have shown that a noisy neural network can time intervals in a way that is both psychologically and physiologically plausible. This network—the SRT model at level 3 of our modeling hierarchy—times responses by controlling the rate at which average activation in its integrator layer rises linearly to a threshold. This pattern of ramping activity is consistent with firing rate data from monkey parietal cortex (Leon and Shadlen, 2003) and presupplementary motor cortex (Mita et al., 2009) in timing tasks. Rapid updating of the model's time estimates, and correspondingly rapid changes in its ramp slopes, are consistent with behavioral data in rats (Davis et al., 1989; Bevins and Ayres, 1995; Higa, 1997), mice (Balci et al., 2008b), and pigeons (Wynne and Staddon, 1988); with human behavioral data from the beat-the-clock task; and with firing rate data from rat thalamic neurons (Komura et al., 2001).

Because it effectively computes a running difference between an excitatory spike count and a proportionally smaller inhibitory spike count, the SRT model approximates a drift-diffusion process with a noise coefficient proportional to the square root of the drift, i.e., the opponent Poisson diffusion model at the top of our modeling hierarchy. For this model, setting the threshold for detecting the end of an interval to the maximum possible value and adapting the drift parameter to time different durations yields unbiased, maximally precise temporal estimation. This parameterization of the model also predicts time-scale-invariant first-passage time distributions. It thus provides a simple, normative explanation for the time scale invariance observed in response time data from timing experiments—a phenomenon spanning multiple species, timing tasks, and orders of magnitude of timed durations (Buhusi and Meck, 2005).

A key strength of timing models based on temporal integration is that they exploit the same computational principle of accumulation to a bound that defines the sequential sampling family of decision-making models (Ratcliff and McKoon, 2008). Response time distributions from perceptual decision-making tasks are typically much more skewed than response time distributions from timing tasks, but this discrepancy presumably occurs because the noise in perceptual decision making includes a large external component provided by the stimulus. In contrast, the noise in timing may be largely internal and a consequence of noisy neural processing.

A second distinction between timing and perceptual decision making involves response thresholds: previously, we have observed reward-maximizing adjustments to response bias and speed–accuracy trade-offs in perceptual decision making that are best modeled as a process of strategic threshold adaptation in a drift-diffusion model (Simen et al., 2006, 2009; Balci et al., 2011). One of the theoretical strengths of sequential sampling models is precisely their ability to accommodate speed–accuracy trade-offs in terms of adjustable thresholds (Luce, 1986). Optimal timing, in contrast, requires that thresholds must be held fixed and drift strategically adapted to maximize precision. The presence or absence of an exogenous signal and a particular response–reward contingency must therefore be capable of switching which parameter—threshold or drift—is adapted if reward-maximizing performance is to be achieved more generally.

Nevertheless, a single set of integrator mechanisms already used to model the neural basis of decision making through sequential sampling (cf. Gold and Shadlen, 2001; Usher and McClelland, 2001; Bogacz et al., 2006) can clearly be parameterized both to time intervals and to make perceptual decisions, suggesting that common brain areas may subserve these arguably distinct behavioral functions. Because of the ease with which they can be applied to both types of tasks, the SRT and opponent Poisson diffusion models can also be used to perform a variety of tasks that combine decision making and time estimation, such as decision making under a response deadline, shorter–longer interval comparison (Leon and Shadlen, 2003), and the "time-left" task (Gibbon and Church, 1981) (supplemental Eq. 5, available at www.jneurosci.org as supplemental material). Importantly, it can also learn the durations on which these computations are based in a single exposure.

Among existing timing models, the opponent Poisson diffusion model is most closely related to scalar expectancy theory (Gibbon, 1977) and the behavioral theory of timing (Killeen and Fetterman, 1988). However, it modifies these models by adding an inhibitory spike train, and it further departs from SET in assuming a fixed, rather than varying, threshold. As a result, the model predicts inverse Gaussian distributions of timed responses that fit behavioral data closely while accommodating much higher spike rates than are allowed by a purely excitatory Poisson counter model. We have shown, in fact, that inverse Gaussians fit timed response distributions better than competing distributions with the same number of parameters (the gamma distribution of BeT and the approximately normal distribution of SET) when behavioral CVs are large enough to distinguish these distributions from each other.

Among decision-making models, the opponent Poisson diffusion model is most closely related to the "wave theory" model of Link (1992). In Link's (1992) model, a Poisson spike train represents stimulus intensity as a spike rate proportional to the intensity. In a greater–lesser comparison of two stimuli, the difference between spike counts for each stimulus is the decision variable, which performs a random walk toward one of two boundaries. When intensity A is multiplied by a constant K, the just-noticeable difference between $K \times A$ and the comparison is also multiplied by K. This is the classic form of Weber's law. Thus, opponent Poisson processing appears to account both for the nontemporal form of Weber's law and for time scale invariance in interval timing.

In this article, we have restricted our attention to the type of timed behavior that the model most naturally explains, which is the planned production of a single timed response. This restriction excludes a type of behavior for which the model can also

easily give an account: this is the emission of a sequence of responses in each trial of a task, as animals produce in the peak-interval task, for example (Catania, 1970; Roberts, 1981). In unpublished work, we have analyzed the performance of hierarchical compositions of timing circuits that fit the data in such tasks. A detailed account of the operational details of such a model is beyond the scope of the current paper, but its basic principle is simple: for each response of the current SRT model, the trigger layer can control the emission of packets of high-rate responses generated by a hierarchically subordinate SRT circuit with a much faster ramp rate (cf. Kirkpatrick, 2002).

The primary virtue of the diffusion and SRT timing models is their simplicity: in addition to accounting for empirical data, the SRT model is constructed entirely from the same components used in computational models of decision-making circuits based on neural integration. These in turn are hypothesized to implement classic psychological models of decision making based on diffusion processes. Therefore, relatively little needs to be added to these models to provide neurally plausible accounts of behavior across a wide range of timing and perceptual decision-making tasks. Furthermore, the rules governing adaptation in the model give it flexibility and robustness, and a simple expression of the statistics of its behavior is available based on the closed-form, inverse Gaussian distribution, which closely fits response time data from timing tasks. These results suggest that networks of leaky integrators whose connection strengths are precisely tuned by appropriate feedback mechanisms may subserve both temporal and perceptual processing in the brain.

References

Ahrens MB, Sahani M (2008) Inferring elapsed time from stochastic neural processes. Adv Neural Inform Processing 20:1–8.

Balci F, Ludvig EA, Gibson JM, Allen BD, Frank KM, Kapustinski BJ, Fedolak ET, Brunner D (2008a) Pharmacological manipulations of interval timing on the peak procedure in mice. Psychopharmacology 201:67–80.

Balci F, Papachristos EB, Gallistel CR, Brunner D, Gibson J, Shumyatsky GP (2008b) Interval timing in genetically modified mice: a simple paradigm. Genes Brain Behav 7:373–384.

Balci F, Freestone D, Gallistel CR (2009a) Risk assessment in man and mouse. Proc Natl Acad Sci U S A, USA 106:2459–2463.

Balci F, Gallistel CR, Allen BD, Frank KM, Gibson JM, Brunner D (2009b) Acquisition of peak responding: what is learned? Behav Processes 80:67–75.

Balci F, Simen PA, Niyogi R, Saxe A, Hughes J, Holmes P, Cohen JD (2011) Acquisition of decision making criteria: reward rate ultimately beats accuracy. Atten Percept Psychophys 73:640–657.

Bevins R, Ayres JJB (1995) One-trial context fear conditioning as a function of the interstimulus interval. Anim Learn Behav 23:400–410.

Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD (2006) The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced choice tasks. Psychol Rev 113:700–765.

Brody CD, Hernandez A, Zainos A, Romo R (2003) Timing and neural encoding of somatosensory parametric working memory in macaque prefrontal cortex. Cereb Cortex 13:1196–1207.

Brunner D, Kacelnik A, Gibbon J (1992) Optimal foraging and timing processes in the starling (*Starling vulgaris*): effect of inter-capture interval. Anim Behav 44:597–613.

Buhusi CV, Meck WH (2005) What makes us tick? Functional and neural mechanisms of interval timing. Nat Rev Neurosci 6:755–765.

Buhusi CV, Aziz D, Winslow D, Carter RE, Swearingen JE, Buhusi MC (2009) Interval timing accuracy and scalar timing in C57BL/6 mice. Behav Neurosci 123:1102–1113.

Catania AC (1970) Reinforcement schedules and psychophysical judgments: a study of some temporal properties of behavior. In: The theory of reinforcement schedules (Schoenfeld WN, ed), pp. 1–42. New York: Appleton-Century-Crofts.

Cheng K, Westwood R (1993) Analysis of single trials in pigeons timing performance. J Exp Psychol Anim Behav Process 19:59–67.

Cheng K, Miceli P (1996) Modeling timing performance on the peak procedure. Behav Process 37:137–156.

Church RM, Meck WH, Gibbon J (1994) Application of scalar timing theory to individual trials. J Exp Psychol Anim Behav Process 20:135–155.

Church RM, Lacourse DM, Crystal JD (1998) Temporal search as a function of the variability of interfood intervals. J Exp Psychol Anim Behav Process 24:291–315.

Church RM, Broadbent HA (1992) A connectionist model of timing. In: Neural network models of conditioning and action (Commons ML, Grossberg S, Staddon JER, eds), pp. 225–240. Hillsdale, NJ: Erlbaum.

Cohen M, Grossberg S (1983) Absolute stability of global pattern formation and parallel memory storage by competitive neural networks. IEEE Trans Systems Man Cybern 13:815–826.

Cragg BG, Temperley HNV (1955) Memory—the analogy with ferromagnetic hysteresis. Brain 78:304–315.

Davis M, Schlesinger LS, Sorenson, CA (1989) Temporal specificity of fear conditioning: Effects of different conditioned stimulus-unconditioned stimulus intervals on the fear-potentiated startle effect. J Exp Psychol Anim Behav Process 15:295–310.

Dragoi V, Staddon JE, Palmer RG, Buhusi VC (2003) Interval timing as an emergent learning property. Psychological Review 110:126–144.

Drew MR, Zupan B, Cooke A, Couvillon PA, Balsam, PD (2005) Temporal control of conditioned responding in goldfish. J Exp Psychol Anim Behav Process 31:31–39.

Durstewitz D (2003) Self-organizing neural integrator predicts interval times through climbing activity. J Neurosci 23 (12):5342–5353.

Eckhoff P, Holmes P, Law C, Connolly P, Gold J (2008) On diffusion processes with variable drift rates as models for decision making during learning. New J Physics 10:015006.

Gallistel CR, Gibbon J (2000) Time, rate and conditioning. Psychol Rev 107:289–344.

Gardiner CW (2004) Handbook of stochastic methods, Ed 3. New York: Springer.

Gerstein GL, Mandelbrot B (1964) Random walk models for the spike activity of a single neuron. Biophysical J 4:41–68.

Gerstner W (1999) Population dynamics of spiking neurons: fast transients, asynchronous states, and locking. Neural Comput 12:43–89.

Gibbon J (1977) Scalar expectancy theory and Weber's law in animal timing. Psychol Rev 84:279–325.

Gibbon J (1992) Ubiquity of scalar timing with a Poisson clock. J Math Psychol 35:283–293.

Gibbon J, Church RM (1981) Time left: linear versus logarithmic subjective time. J Exp Psychol Anim Behav Process 7:87–108.

Gibbon J, Church RM (1990) Representation of time. Cognition 37:23–54.

Gibbon J, Church RM, Meck WH (1984) Scalar timing in memory. Ann N Y 423:52–77.

Gold JI, Shadlen MN (2000) Representation of a perceptual decision in developing oculomotor commands. Nature 404:390–394.

Gold JI, Shadlen MN (2001) Neural computations that underlie decisions about sensory stimuli. Trends Cogn Sci 5:10–16.

Goldman MS (2009) Memory without feedback in a neural network. Neuron 61:621–634.

Goldman MS, Levine JH, Major G, Tank DW, Seung HS (2003) Robust persistent neural activity in a model integrator with multiple hysteretic dendrites per neuron. Cereb Cortex 13:1185–1195.

Grossberg S, Schmajuk NA (1989) Neural dynamics of adaptive timing and temporal discrimination during associative learning. Neural Netw 2:79–102.

Harth EM, Csermely TJ, Beek B, Lindsay RD (1970) Brain functions and neural dynamics. J Theor Biol 26:93–120.

Haß J, Blaschke S, Rammsayer T, Herrmann JM (2008) A neurocomputational model for optimal temporal processing. J Comp Neurosci 25:449–464.

Hertz J, Krogh A, Palmer RG (1991) Introduction to the theory of neural computation. In: Santa Fe Institute studies in the sciences of complexity lecture notes, Vol. 1. Redwood City, CA: Addison Wesley.

Higa JJ (1997) Rapid timing of a single transition in interfood interval duration by rats. Anim Learn Behav 25:177–184.

Higham DJ (2001) An algorithmic introduction to numerical simulation of stochastic differential equations. Soc Indust Appl Math Rev 43:525–546.

Hudson TE, Maloney LT, Landy MS (2008) Optimal compensation for temporal uncertainty in movement planning. PLoS Comput Biol 4:e1000130.

Ivry RB, Richardson TC (2002) Temporal control and coordination: the multiple timer model. Brain Cogn 48:117–132.

Jackson AS (1960) Analog computation. New York: McGraw-Hill.

Jazayeri M, Shadlen, MN (2010) Temporal context calibrates interval timing. Nat Neurosci 13:1020–1026.

Jordan DW, Smith P (1999) Nonlinear ordinary differential equations, Ed 3. New York: Oxford UP.

Jorgensen B (1987) Exponential dispersion models. J R Stat Soc B Stat Methodol 49:127–162.

Karmarkar UR, Buonomano, DV (2007) Timing in the absence of clocks: encoding time in neural network states. Neuron 53:427–438.

Killeen PR, Fetterman JG (1988) A behavioral theory of timing. Psychol Rev 95:274–295.

Kirkpatrick K (2002) Packet theory of conditioning and timing. Behav Process 57:89–106.

Komura Y, Tamura R, Uwano T, Nishijo H, Kaga K, Ono T (2001) Retrospective and prospective coding for predicted reward in the sensory thalamus. Nature 412:546–549.

Koulakov AA, Raghavachari S, Kepecs A, Lisman JE (2002) Model for a robust neural integrator. Nat Neurosci 5:775–782.

Lapique L (1907) Recherches quantitatives sur l'excitation électrique des nerfs traitée comme une polarization. J Physiol Pathol Genet 9:620–635.

Leon MI, Shadlen MN (2003) Representation of time by neurons in the posterior parietal cortex of the macaque. Neuron 38:317–327.

Link SW (1992) The wave theory of difference and similarity. Hillsdale, NJ: Lawrence Erlbaum.

Lo C-C, Wang X-J (2006) Cortico-basal ganglia circuit mechanism for a decision threshold in reaction time tasks. Nat Neurosci 9:956–963.

Luce RD (1986) Response times: their role in inferring elementary mental organization. New York: Oxford UP.

Ludvig EA, Sutton RS, Kehoe EJ (2008) Stimulus representation and the timing of reward-prediction errors. Neural Comput 20:3034–3054.

Macar F, Vidal F (2003) The CNV peak: an index of decision making and temporal memory. Psychophysiology 40:950–954.

Machado A (1997) Learning the temporal dynamics of behavior. Psychol Rev 104:241–265.

Major G, Tank DW (2004) Persistent neural activity: prevalence and mechanisms. Curr Opin Neurobiol 14:1–10.

Matell M, Meck WH (2004) Cortico-striatal circuits and interval timing: coincidence detection of oscillatory processes. Cognitive Brain Res 21 (139–170).

Matell MS, Portugal GS (2007) Impulsive responding on the peak-interval procedure. Behav Process 74:198–208.

Meck WH, Komeily-Zadeh FN, Church RM (1984) Two-step acquisition: modification of an internal clock's criterion. J Exp Psychol Anim Behav Process 10:297–306.

Miall RC (1989) The storage of time intervals using oscillating neurons. Neural Comput 1:359–371.

Mita A, Mushiake H, Shima K, Matsuzaka Y, Tanji J (2009) Interval time coding by neurons in the presupplementary and supplementary motor areas. Nat Neurosci 12:502–507.

Onn SP, Grace AA (2000) Amphetamine withdrawal alters bistable states and cellular coupling in rat prefrontal cortex and nucleus accumbens neurons recorded *in vivo*. J Neurosci 20:2332–2345.

Rakitin BC, Gibbon J, Penney TB, Malapani C, Hinton SC, Meck WH (1998) Scalar expectancy theory theory and peak-interval timing in humans. J Exp Psychol Anim Behav Process 24:15–33.

Ratcliff R, Rouder J (1998) Modeling response times for two-choice decisions. Psychol Sci 9:347–356.

Ratcliff R, McKoon G (2008) The diffusion decision model: theory and data for two-choice decision tasks. Neural Comput 20:873–922.

Renart A, de la Rocha J, Bartho P, Hollender L, Parga N, Reyes A, Harris K (2010) The asynchronous state in cortical circuits. Science 327:587–590.

Reutimann J, Yakovlev V, Fusi S, Senn W (2004) Climbing neuronal activity as an event-based cortical representation of time. J Neurosci 24:3295–3303.

Roberts S (1981) Isolation of an internal clock. J Exp Psychol Anim Behav Process 7:242–268.

Roitman J, Shadlen M (2002) Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. J Neurosci 22:9475–9489.

Ross S (1998) A first course in probability, Ed 5. Upper Saddle River, NJ: Prentice Hall.

Seung HS (1996) How the brain keeps the eyes still. Proc Natl Acad Sci U S A 93:13339–13344.

Seung HS, Lee DD, Reis BY, Tank DW (2000) The autapse: a simple illustration of short-term analog memory storage by tuned synaptic feedback. J Comput Neurosci 9:171–185.

Shadlen MN, Newsome WT (1998) The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. J Neurosci 18:3870–3896.

Shadlen MN, Newsome WT (2001) Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. J Neurophysiol 86:1916–1936.

Shea-Brown E, Rinzel J, Rakitin B, Malapani C (2006) A firing-rate model of parkinsonian deficits in interval timing. Brain Res 1070:189–201.

Simen P, Cohen JD, Holmes P (2006) Rapid decision threshold modulation by reward rate in a neural network. Neural Netw 19:1013–1026.

Simen P, Contreras D, Buck C, Hu P, Holmes P, Cohen JD (2009) Reward rate optimization in two-alternative decision making: empirical tests of theoretical predictions. J Exp Psychol Human Percept Perform 35:1865–1897.

Skellam JG (1946) The frequency distribution of the difference between two Poisson variates belonging to different populations. J R Stat Soc Ser A 109:296.

Smith PL (2010) From Poisson shot noise to the integrated Ornstein-Uhlenbeck process: Neurally principled models of information accumulation in decision-making and response time. J Math Psychol 54:266–283.

Treisman M (1963) Temporal discrimination and the indifference interval: implications for a model of the "internal clock." Psychol Monographs 77:1–31.

Usher M, McClelland JL (2001) The time course of perceptual choice: the leaky, competing accumulator model. Psychol Rev 108:550–592.

van Vreeswijk C, Sompolinsky H (1996) Chaos in neuronal networks with balanced excitatory and inhibitory activity. Science 274:1724–1726.

Wackermann J, Ehm W (2006) The dual klepsydra model of internal time representation and time reproduction. J Theoretical Biol 239:482–493.

Wald A (1947) Sequential analysis. New York: Wiley.

Wang XJ (2001) Synaptic reverberation underlying mnemonic persistent activity. Trends Neurosci 24:455–463.

Wilson HR, Cowan JD (1973) A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. Kybernetik 13:55–80.

Wynne CDL, Staddon J (1988) Typical delay determines waiting time on periodic-food schedules: static and dynamic tests. J Exp Anal Behav 50:197–210.

Yi L (2007) Applications of timing theories to a peak procedure. Behav Process 75:188–198.