Interval timing by neural integration: Supplementary Materials

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

Princeton Neuroscience Institute Princeton University Princeton, NJ

In these Supplementary Materials, we describe additional neural mechanisms needed to justify certain details of the timing model described in the main text, and we supplement the human behavioral evidence for rapid duration learning with evidence from mice performing a timing task. As discussed in section 1, the additional neural mechanisms include a network that performs temporal differentiation to decode the time left before an upcoming event, and a feedback controller that uses differentiation to stabilize the integrator dynamics of the timing model. In section 2, we derive the linear noise assumption used in the main text from the sigmoidal activation function of our leaky integrator model of neural population activity. Section 3 demonstrates that the presence of a lower reflecting boundary in a drift-diffusion model of neural activity still leads to an approximately inverse Gaussian distribution of response times (RTs), despite the strict dependence of this inverse Gaussian shape on the absence of such a bound. In section 4, we consider a spectrum of alternative models of noise in neural activity, and we consider the ways in which these alternatives could affect our explanation of scalar invariance in terms of competitive Poisson shot noise (provided in the main text). In section 5, we present additional evidence for rapid learning of new interval durations in mice. Finally, section 6 shows the performance of the model on the Beat-the-Clock task.

1. Neural mechanisms for differentiation and integration

1.1. A neural mechanism for temporal differentiation

It may be useful to take derivatives of ramping activation in a neural integrator to recover the inputs to that integrator, and to tune the properties of that integrator (one means of which we demonstrate in Sec. 1.2).

We compute ramp slopes by differentiating ramp-unit activations with respect to time. The differentiator circuit we propose is quite simple, involving one extra stage of leaky integration followed by a weighted subtraction operation. This mechanism implements Euler's method for approximating a derivative by sampling a function g(u) at two different points u_1 and u_2 and dividing $g(u_2) - g(u_1)$ by $u_2 - u_1$. The circuit accomplishes this by imposing a delay on a signal, then subtracting the delayed

version from the original and dividing the result by the delay magnitude:

$$\frac{dx}{dt} \approx \frac{x(t + \Delta t) - x(t)}{\Delta t}.$$
(1)

This circuit takes advantage of the 'group delay' that a leaky integrator imposes on the signals that propagate through it. Group delay (Oppenheim and Willsky, 1996) is the duration of the time delay that is imposed on each narrow frequency band of the input signal to a filter. When the input signal is a ramp, the output of a leaky integrator asymptotically approaches the same ramp delayed in time.

The system of 3 units shown in Supplementary Fig. 1a, with activations $x_1(t)$, $x_2(t)$ and $x_3(t)$, can compute an approximation of the derivative of $x_1(t)$ as long as $x_1(t)$ does not vary too quickly. Specifically, given a constant input signal \tilde{A} , (corresponding to drift A), which is the normal operating regime of the SRT timing model, the output of this derivative circuit asymptotically approaches A over time. The circuit is formally specified as follows:

$$\dot{x}_{1} = A$$

$$\dot{x}_{2} = -x_{2} + \frac{k-1}{k} \cdot x_{2} + \frac{1}{k} \cdot x_{1}$$

$$\dot{x}_{3} = -x_{3} + \frac{x_{1} - x_{2}}{k}$$
(2)

It is clear that the solution for x_1 when A is constant is $x_1(t) = A \cdot t + c_1$ (with $x_1(0) = c_1$). The system is simple enough that it can be solved one variable at a time.

For $x_2(t)$ we obtain:

$$_{2}(t) = A - \frac{c_{2}}{k}e^{-t/k} - x_{3},$$
(3)

and for $x_3(t)$, we obtain:

x

$$x_3(t) = A - \frac{c_2}{k-1}e^{-t/k} + c_3 e^{-t}.$$
 (4)

Therefore $\lim_{t\to\infty} x_3(t) = A$, as illustrated in Supplementary Fig. 1b. In this way, the average duration-encoding weight of the SRT timing model can be translated into the activation $x_3(t)$



Figure 1: **a**: A circuit for differentiating low-frequency signals with respect to time. Labels indicating leaky integration indicate the shape of leaky units' exponential impulse response, which is more rapidly decaying for the red output unit. **b**: Timecourses of unit activations without noise. **c**, **d**, **e**: The effect of interconnection noise in the differentiator mechanism. **c**: A short delay between the original and delayed signals (k small) amplifies noise in the input signal, leading to wild fluctuations in the derivative estimate (red), with average behavior near the value of the time derivative of the noiseless input signal, plotted as the black, dashed horizontal line. **d**: Longer delays (larger k) lead to a better estimate. **e**: A functionally useful estimate. Larger k leads to greater smoothing but slower asymptotic approach to the desired derivative value, illustrating a tradeoff between speed and precision.

by the derivative circuit. Faster approach to the derivative value can be achieved by reducing the delay of the delayed signal. However, in this case, noise in the input signal begins to have a detrimental effect. Supplementary Fig. 1c, d and e show the effects of a constant level of noise for three different signal delays. A speed-precision tradeoff is therefore evident, but with enough averaging of the ramp prior to differentiating, any level of interconnection noise can be handled by some parameterization of this circuit with a given delay.

If we additionally assume that the operation of division can be implemented neurally, then the time remaining until an upcoming event, or 'time left' (cf. Gibbon and Church, 1981), can be computed by dividing the remaining distance between a response threshold z and the current position of a ramp variable $x_1(t)$ by the slope of the timer ramp, $A \approx x_3(t)$:

Time Left(t) =
$$\frac{z - x_1(t)}{x_3(t)}$$
. (5)

Gibbon and Church (1981) showed that rats and pigeons were able to compute the time remaining before each of two possible rewards in their time-left task, and to select more frequently the actions that led to the earlier reward. Since one of the two reward-delays continuously ran down during the choice procedure, animals needed to be able to compute the quantity in Supplementary Eq. 5. We have shown that this computation can easily be achieved with a ramping firing-rate model using the proposed temporal differentiation circuit (although it can also be achieved by using the inputs to the integrator — the outputs of the active Clock-Pulse Generator switches in the SRT model — as an estimate of the ramp slope A).

The derivative operation we have described can theoretically be applied to any input signal $x_1(t)$. However, in general, the output of the circuit will become an increasingly poor approximation of its derivative as higher frequencies appear in $x_1(t)$, since leaky integration in the delay unit with activation x_2 attenuates high frequencies (but see Tripp and Eliasmith, 2010, for alternative models of neural differentiation). For our purposes, applying a differentiator to a ramp to derive its slope is both relatively simple, since high frequencies are not present in gradual ramps, and important for tuning the properties of a neural integrator, as we now demonstrate.

1.2. Feedback control for achieving nearly perfect integration

A critical aspect of neural models that implement computations such as integration by precisely tuning parameters is that mistuned parameter values will destroy their operation (Seung et al., 2000). Models that involve a 'line attractor' for memorizing an analog quantity over a delay period, for example (e.g., Machens et al., 2005), require exquisite precision in the shape of the output functions of the units that produce the attractor.

Here, line attractors are not used, but as in the autapse model of Seung et al. (2000), precise balance must be achieved in the integrator unit between leak and positive feedback (cf. the similarly precise balance between leak and lateral inhibition that must be achieved in the neural drift-diffusion implementation of Bogacz et al., 2006). A simple and effective method for precise tuning in the current case — i.e., the case in which a constant signal is being integrated — is to use feedback control of the relevant parameters based on deviations from linearity in the resulting ramp.

$$\frac{dw}{dt} = -q \cdot \frac{d^2 x_1(t)}{dt^2}.$$
(6)

Supplementary Eq. 6 states that if the second derivative with respect to time of the ramping activation in the SRT timer is greater than zero, then the activation is accelerating and selfexcitation is too strong. In this case, the recurrent, excitatory weight of the ramp unit is weakened in proportion to the ramp's second derivative (or in proportion to a leaky integral of this value). If the second derivative is negative, the excitatory weight is strengthened in proportion to the absolute value of the second derivative. With a high enough proportionality constant, or gain, this linear feedback control law is guaranteed to keep the self-excitatory weight close enough to zero to ensure approximately linear ramps (the weight is subject to a steady-state error that is small when the gain is large, although a very large gain can lead to instability; Franklin et al., 1994). The key to this approach is the ability to differentiate with respect to time, which we can accomplish using the differentiator mechanism previously discussed.

It should be noted that the derivative circuit in Supplementary Fig. 1 itself requires precise tuning of three weights that all depend on the value k. However, whether the estimated derivative is positive or negative is all that matters for the feedback control law in Supplementary Eq. 6; the self-exciting, recurrent weight equal to (k - 1)/k on unit 2 merely determines the amount of delay that the circuit imposes on signals and does not affect the sign of the circuit's output. Thus the only truly critical requirement for ramp-tuning is that the two weights from units 1 and 2 to unit 3 in Supplementary Fig. 1 are of approximately equal magnitudes and opposite sign, a condition which we assume can be achieved by any of a number of simple mechanisms (e.g., synaptic weight normalization, or other forms of feedback control).

2. Simplified equation for population activity

We now derive the simplified stochastic differential equation for neural population activity on which the SRT model is based from the type of two-equation systems found in, e.g., Hopfield and Tank (1985). Such two-equation systems make it possible to model sigmoidal squashing in the most plausible manner. Specifically, activations are based on a sigmoidally squashed sum of weighted inputs that include additive, Gaussian, interconnection noise. In contrast, Eq. 6 in Methods applies sigmoidal squashing only to the noiseless component of the inputs. The model of Hopfield and Tank (1985), represented by the deterministic component of Eq. 8 and Eq. 9, contains an equation for leaky integration of weighted sums of input from other units (Eq. 8) followed by amplification with saturating upper and lower bounds (Eq. 9). We add white noise to the input computation to represent the effect of Poisson shot noise, transmitted through stochastic synapses:

$$I_{i} = \sum_{j=1}^{n} w_{ij} f_{j}(x_{j})$$
(7)

$$\tau \cdot dx_i = \left[-x_i + I_i + \sqrt{I_i} \cdot \frac{dB_i}{dt} \right] dt$$
$$= \left(-x_i + I_i \right) dt + \sqrt{I_i} \cdot dB_i \tag{8}$$

$$V_i(t) = f_i(x_i(t)) = \frac{1}{1 + \exp(-\lambda \cdot (x_i - \beta))}.$$
 (9)

For simplicity here, we assume that the balance between excitation and inhibition, represented in the main text by γ , is folded into the weights w_{ij} .

In general, we will not require the connection strengths, w_{ij} between any two units, *i* and *j*, to be equal. However, it will simplify matters to assume that all units have the same activation functions, $f_i \equiv f$. Assuming that amplification of inputs by all units is instantaneous, we can replace $f(x_j)$ with V_j . Using Ito's Lemma (i.e., the chain rule, for stochastic systems), we then arrive at a single equation for the system, Eq. 10:

$$\tau \cdot dV_i = \frac{df}{dx_i} \cdot \left[-f^{-1}(V_i) + \sum_{j=1}^n w_{ij}V_j \right] dt + \dots$$
$$\frac{1}{2} \frac{d^2f}{dx_i^2} \cdot I_i dt + \frac{df}{dx_i} \cdot \sqrt{I_i} dB_i. \tag{10}$$

The second term of Eq. 10 is 0 at the inflection point ($x_i = \beta$) of the sigmoid function f, since the second derivative is 0 there, and is nearly 0 in a range around the inflection point. Examining the first term of Eq. 10, we see that as $f^{-1}(V) = x$ approaches I for some constant input I, V must approach f(I).

Furthermore, when x is in the linear range of f, $df/dx \approx \lambda/4$, and $f(x) \approx (\lambda/4) \cdot x + (2 - \lambda\beta)/4$. Thus $f^{-1}(V) \approx (4V - 2 + \lambda\beta)/\lambda$. We therefore get the following deterministic equation:

$$\tau \cdot dV_i \approx \frac{\lambda}{4} \cdot \left(-\frac{4}{\lambda}V_i + \frac{2}{\lambda} - \beta + I\right) dt$$
$$\approx \left(-V_i + \frac{\lambda}{4}I + \frac{2 - \lambda\beta}{4}\right) dt$$
$$\approx \left(-V_i + f(I)\right) dt \tag{11}$$

In Eq. 11, we have returned the nonlinear term f after inverting its linearization, so it is not obvious that this system closely approximates Eq. 4 in the main article, in Materials and Methods. However, it is only when V approaches 1 or 0 that nonlinearities affect the approximation, and they will have a small effect, as numerical simulations confirm. Thus we can use the following more manageable equation in our simulations and analyses:

$$\tau \cdot dV_i = \left(-V_i + f\left(\sum_{j=1}^n w_{ij}V_j\right)\right) dt + \frac{df}{dx} \cdot \sqrt{\sum_{j=1}^n w_{ij}V_j} \, dB_i.(12)$$

If we assume that $\lambda = 4$, we get the following simpler equation (equivalent to Eq. 4 in the main article), on which the SRT model is based:

$$\tau \cdot dV_i = \left(-V_i + f\left(\sum_{j=1}^n w_{ij}V_j\right)\right) dt + \sqrt{\sum_{j=1}^n w_{ij}V_j} \, dB_i.(13)$$

(If $\lambda \neq 4$, other parameters can be adjusted to obtain the same result.)

3. Negligible effects of a lower bound on activation on the first-passage time distribution

One gap in our neural network approximation of a driftdiffusion process derives from the assumption that the integrator unit applies a hard lower bound of 0 to its activation. The inverse Gaussian distribution generated by an equivalent driftdiffusion process with only one absorbing boundary assumes no such lower bound, so that the process can dip arbitrarily far below 0. It is therefore important to verify that applying a lower bound does not make the inverse Gaussian a poor approximation to the first-passage time distribution of the bounded integrator. Supplementary Fig. 2 shows a simulation of a bounded, noisy integrator that justifies use of the inverse Gaussian as long as the constant noise coefficient c is not too large relative to the threshold/drift ratio z/A. Formally, this model is as follows:

$$dx = A \cdot dt + c \cdot dB. \tag{14}$$

4. Scalar invariance is robust to variations in the function relating firing-rate mean to firing-rate variance

In this section, we examine alternative models of the functional relationship between firing rates and noise, since the Poisson shot noise assumption made in the main text, although plausible, is possibly too restrictive. It implies, for example, a Fano



Figure 2: Neural network units with sigmoid nonlinearities cannot take on negative activation values, but the inverse Gaussian distribution arises for drift-diffusion processes whose sample paths can be negative, with arbitrarily large absolute values. Drift-diffusion simulations shown in the bottom panel (c = 0.25, z = 1, A = 0.5), in contrast, have a reflecting boundary at 0, and are a better representation of a noisy sigmoid unit's activation. The inverse Gaussian density (red curve in top panel) is approximated quite well by first-passage time histogram as long as noise is sufficiently small relative to the threshold/drift ratio. When noise is larger, the approximation gets worse, but in such cases, the distribution is in any case too strongly skewed to account for typical behavioral data in timing.

factor that is exactly equal to 1 (i.e., the firing-rate variance should equal the firing-rate mean), but this condition is often violated in electrophysiological data (cf. Churchland et al., 2010). Nevertheless, if neural resources are allocated within the SRT model so as to minimize response time variance, then approximate scalar invariance still holds over a large range of interval durations. To prove this, we first review how scalar invariance arises in a simplified version of our model that assumes deterministic ramping. Approximating this deterministic ramping condition by averaging over many units allows the SRT model to preserve approximate scalar invariance over a range of noise assumptions. Finally, we formally define the neuronal computing elements of the SRT model and derive an analytical approximation of its response time (RT) variance. This expression will be used in the Results to determine the variance-minimizing allocation of elements to each layer of the model.

4.1. Noise model spectrum

The basic model is defined by Supplementary Eq. 14. On each trial, the process x(t) starts with x(0) = 0, and continues until first-passage through an absorbing boundary at x = z. The drift is *A*.

The following assumptions define a spectrum of possibilities regarding the functional relationship between noise and mean firing rate that encompass the range of what seems plausible. At the lower end of the spectrum (assumption 1 below), noise is constant, and has no relationship with firing rates. At the upper end (assumption 3 below), the variance of firing rates grows as the square of the mean. Assumption 3 seems quite implausible since observed Fano factors in cortical neurons are rarely much greater than 2 (cf. Churchland et al., 2010), but it provides a convenient upper bound for our analysis. The competitive Poisson and SRT models of the main text fall exactly in the middle of this spectrum.

Assumption 1: The noise coefficient is constant: c = k. This leads to CVs that grow as the timed duration T grows.

Assumption 2: The noise coefficient scales linearly with the square root of the drift: $c = k \cdot \sqrt{A}$. This leads to scalar invariance — i.e., a constant coefficient of variation (CV) of response times (RTs) — across all timed durations.

Assumption 3: The noise coefficient scales linearly with the drift: $c = k \cdot A$. This leads to CVs that decrease with increasing duration.

While the skewness-to-CV ratio is predicted to be constant and equal to 3 (see main text, section 3.1), regardless of the timed interval duration, the CV and skewness are individually constant only under assumption 2. As shown in the main text, section 3.1, the CV for the drift-diffusion model is as follows:

$$CV = \frac{c}{\sqrt{Az}}.$$
 (15)

In general, if we substitute $c = kA^r$, we get the following:

$$CV = \frac{kA^r}{\sqrt{Az}} = kA^{r-1/2}z^{-1/2}.$$
 (16)

This CV function is plotted for different values of r, with k = 0.3 and z = 5, in Supplementary Fig. 3.



Figure 3: Coefficient of variation (CV) in response times for alternative noise models ($0 \le r \le 1$), as a function of the timed interval duration.

To implement assumption 1 (corresponding to r = 0), we substitute the constant *k* for *c*:

$$CV_1 = \frac{k}{\sqrt{Az}} = \frac{k}{z} \cdot \sqrt{\frac{z}{A}} = \frac{k}{z} \sqrt{E(T)}.$$
 (17)

For this model, the CV grows with the square root of the duration, in violation of scalar invariance. As shown in the main text, under assumption 2 (r = 0.5), the CV is independent of A:

$$CV_2 = \frac{k}{\sqrt{z}}.$$
(18)

For this model, as we have shown, setting z to be fixed at its maximum allowable value for all interval durations minimizes the CV and results in scalar invariance.

To implement assumption 3 (r = 1), we set *c* equal to *kA*:

$$CV_3 = \frac{kA}{\sqrt{Az}} = k\sqrt{\frac{A}{z}} = \frac{k}{\sqrt{E(T)}}.$$
(19)

Thus the CV of this model decreases with increasing duration T (it is still the case that E(t) = T). This decreasing CV result holds whether or not we assume that z is constant. Since it is so far from what is observed experimentally, we consider only the range of r < 0.5.

Values of *r* that are far from 0.5 seem quite implausible with respect to physiological and behavioral data. For *r* values less than 0.5, however, an optimization argument shows that minimizing the CV by allocating more neural resources to the ramping mechanism of the model leads to approximate scalar invariance over a large range of interval durations. This allocation process effectively makes the model approximate a model with a deterministic, linear ramp up to a threshold that is normally distributed across trials. We show this specifically for the most extreme case of r = 0. For *r* values slightly greater than 0.5, the CV can be minimized for very short intervals by again averaging out the noise in the ramping process, but we do not further consider the case of $r \gg 0.5$.

4.2. The deterministic ramp model

An abstract model that can serve as a bridge between theories of decision making and timing is the model of Grice (1968). In this model of stimulus categorization, a fixed-intensity perceptual input drives an integrator to ramp linearly and deterministically toward a normally distributed response threshold. The stimulus intensity determines the slope of the ramp and thus the duration of the process — the RT — while variability in the threshold produces variance in the RTs.

Supplementary Fig. 4 shows how this model can be used to time two different durations, by ramping at rates inversely proportional to the durations to be timed. We refer to Grice's model as the 'deterministic ramp model' when we apply it in this way to interval timing. The deterministic ramp model is essentially a continuous-time version of classic, digital counting models of timing, in which the integration of a constant stimulus replaces the counting of periodic or Poisson clock pulses. That the deterministic ramp model predicts scalar invariance can be seen by reflecting the threshold distribution across the ramp used for a particular duration onto the time axis. The standard deviation of each RT distribution is equal to the standard deviation of the threshold value divided by the ramp slope, and thereby increases linearly with duration. Furthermore, when the time axis is divided by the duration of the timed interval, the RT distributions for long and short intervals superimpose on each other, since they are both Gaussian.

We show that CV minimization produces a mix of integrator and trigger units in the SRT model that is sufficient to reduce noise in the integration process (approximating a deterministic ramp), while retaining substantial variability in the trigger process. Furthermore, threshold variability (i.e., variability in the level to which ramps rise by the time of a response) exhibits a normal distribution as an emergent property of this system. Thus, in the limit of large numbers of units, the SRT model approximates the simpler deterministic ramp model, reproducing scalar invariance, whether or not r is exactly 0.5.



Figure 4: A noiseless ramp with a normally distributed threshold — namely, the decision making model of Grice (1968), which we refer to as the 'deterministic ramp' model when applied to interval timing. Scalar invariance results automatically from this model, since RT standard deviation equals the threshold standard deviation divided by the ramp slope.

4.3. Threshold statistics

To determine the value of an effective threshold, we simulated the SRT model with r = 0 many times, taking the time \hat{T} of behavioral response to be equal to the time at which the trigger unit exceeds an activation value of 0.6 (near this level, bistable units experience their maximal rates of activation change and thus function best as response triggers). We then took the value of the threshold on each trial to be the value of the ramp unit's activation at \hat{T} . Gaussian, analytical approximations of threshold distributions exist for the deterministic ramp case (e.g., Lim and Rinzel, 2009), and simulated spiking models of decision making circuits also produce approximately normal thresholds (Lo and Wang, 2006). Supplementary Fig. 5a shows that this pattern of approximately normally distributed thresholds is preserved in our model (it is seen in all our simulations). Thus the normally distributed thresholds assumed in many timing models (e.g., Gibbon, 1992) appear to arise naturally from the dynamics of a nonlinear readout/comparison process in both simplified and more biophysically detailed neural models.

We now note an important limitation of the trigger mechanism. Supplementary Fig. 5b shows that thresholds are only



Figure 5: A: A typical threshold distribution across 1500 timing trials for the SRT model with r = 0, showing the approximate normality seen in simulations across a wide range of parameters (here the drift-diffusion process used c = 0.05 and A = 0.6, and trigger unit parameters in Supplementary Eq. 14, were c = 0.05, $\alpha = 4$, $\beta = 1.3$, $w_{ii} = 2$, $\tau = 1$). The best-fitting normal distribution has mean 0.75 and standard deviation 0.075, implying a CV of 0.075/0.75 = 0.1. B: Average threshold value across 100 trials for different ramp rates, plotted vs. the corresponding RT. Note that this threshold value is mostly constant but increases abruptly at very small RTs, where the separation of ramp and trigger time scales breaks down (see text).

approximately constant across conditions: as the timed duration grows short, the threshold increases. With very high ramp rates, the speed of the trigger mechanism can no longer be considered effectively infinite relative to ramp speed, and the ramp reaches a much higher level by the time a response is generated. Even with sigmoidal squashing of ramp activation, this implies a lower limit on the durations that can be encoded and estimated by this mechanism — a duration range (e.g., 0—500 msec) in which a timing mechanism based more heavily on inter-spike intervals by individual neurons seems likely to take over (e.g., Haß et al., 2008).

4.4. Response time density for the drift-diffusion model with a single, stochastic threshold

For the drift-diffusion process defined by Supplementary Eq. 14, with x(t) ($0 \le t < \infty$) with a single, deterministic threshold or absorbing boundary z > 0, drift A > 0, noise c > 0 and starting point x(0) = 0, the probability distribution of the 'first-passage time' at which x(t) first becomes larger than z is the Wald or inverse Gaussian distribution (Luce, 1986). The Wald density (Wald, 1947), plotted in Fig. 3 of the main article, is typically specified as follows, with $\mu = z/A$, and $\gamma = (z/c)^2$:

$$p(t,\mu,\gamma) = \sqrt{\frac{\gamma}{2\pi t^3}} \cdot \exp\left(\frac{-\gamma \cdot (t-\mu)^2}{2\mu^2 t}\right).$$
(20)

The expected first-passage time is μ , and the variance is μ^3/γ (Luce, 1986).

When the assumption of a normally distributed threshold is added to the model just described, the resulting first-passage density arises from convolving the threshold distribution with Supplementary Eq. 20, replacing z in Supplementary Eq. 20 with the normally distributed random threshold value ζ , whose mean and standard deviation we label z and σ_z :

$$p(t, A, c, z, \sigma_z) = \dots$$

$$\int_{-\infty}^{\infty} \left\{ \sqrt{\frac{(\zeta/c)^2}{2\pi t^3}} \cdot \exp\left(\frac{-(\zeta/c)^2 \cdot (t - \zeta/A)^2}{2(\zeta/A)^2 t}\right) \times \dots \right.$$

$$\sqrt{\frac{1}{2\pi\sigma_z^2}} \cdot \exp\left(-\frac{(\zeta - z)^2}{2\sigma_z^2}\right) \right\} d\zeta. \quad (21)$$

The effect of this convolution is to smear out the Wald distribution further as σ_z grows larger (see Fig. 3 of the main article). However, the expected first-passage time for both models (Supplementary Eqs. 20 and 21) is the same (E[t] = z/A) for equal values of A, c, x(0) and z, since convolution with a Gaussian redistributes probability mass symmetrically on either side of the Wald's mean (which we confirmed with numerical simulations).

The expected first-passage time of the stochastic threshold model, conditioned on a particular threshold ζ , is just μ (Luce, 1986):

$$\mathbf{E}[t|\zeta] = \zeta/A. \tag{22}$$

By the law of total expectation, $E[t] = E[E[t|\zeta]] = z/A$. Thus the average RT of a drift-diffusion process with a normally distributed threshold will be the same as we would get with a single, deterministic threshold: the average threshold divided by the drift.

The variance of the inverse Gaussian is μ^3/γ (Luce, 1986), or in terms of ζ , *c* and *A*:

$$\operatorname{Var}(t|\zeta) = \frac{c^2 \zeta}{A^3}.$$
(23)

The variance of the model's RT distribution can now be determined from the law of total variance by substituting $E[t|\zeta]$ and $Var(t|\zeta)$ into the total variance formula (Rice, 1995):

$$Var(t) = E[Var(t|\zeta)] + Var[E(t|\zeta)].$$
(24)

First we derive $E[Var(t|\zeta)]$. Since $Var(t|\zeta)$ is just proportional to ζ by Supplementary Eq. 23, we get:

$$E[Var(t|\zeta)] = \frac{c^2}{A^3} E(\zeta) = \frac{c^2 z}{A^3}.$$
 (25)

Next we derive $Var[E(t|\zeta)]$, assuming that the standard deviation of the threshold distribution is some real multiple *K* of the ramp noise coefficient:

$$\operatorname{Var}[\mathrm{E}(t|\zeta)] = \operatorname{Var}\left(\frac{\zeta}{A}\right) = \frac{1}{A^2}\operatorname{Var}(\zeta) = \frac{1}{A^2}K^2c^2. \tag{26}$$

SRT simulations with r = 0 suggest that, consistent with Supplementary Eq. 26, threshold variance $Var(\zeta)$ does indeed grow approximately linearly with the square of the noise coefficient. In our simulations, we always obtained a value of *K* on the order of $0.1 \cdot z$, although the specific value appears to depend on the number of units in the model, the amount of noise at each

connection, and the self-excitation strength of the trigger layer. We thus define $\alpha = K/z$.

Putting these results together yields:

$$\operatorname{Var}(t) = \frac{c^2 z}{A^3} + \left(\frac{\alpha c z}{A}\right)^2,\tag{27}$$

where the first term represents the contribution of the ramp to the variance, and the second, the contribution of the trigger mechanism.

For all values of r between 0 and 1, averaging across uncorrelated noise sources reduces the value of k and thereby reduces Var(t). Since A and z are free parameters aside from the requirement that z/A = T, we consider what values these should take on to minimize Var(t). For Supplementary Eq. 27, maximizing both A and z is the best way to reduce variability for long intervals, given the A^3 in the denominator of the first term contributed by the ramp.

4.5. Approximating scalar invariance by constrained variance minimization

We now investigate the effect of plausible noise-reduction methods. To reduce noise, we exploit the weighted sum operation that each unit applies to its inputs in order to compute averages across N redundant ramp and M redundant trigger units (cf. Shadlen and Newsome, 1998). Minimizing RT variance with this technique reduces to solving a resource allocation problem, for which the optimal solution reduces noise in the ramp but leaves trigger noise mostly unchanged.

For ease of analysis, we assume that the white noise added at the connection between any pair of units in Supplementary Fig. 7a has zero cross-correlation (each synapse acts as an independent source of white noise). However, despite this assumption, recurrent connections among the trigger units act to correlate their responses to this independent interconnection noise. This leads to inputs from other trigger units that include relatively strongly correlated stochastic components - components for which averaging is not as effective a remedy. (Assuming independent interconnection noise simplifies our presentation, but a small amount of interconnection correlation does not drastically change the model's properties - indeed, simulations suggest that small, positive noise correlations act primarily to increase the effective noise coefficient at each connection, thereby increasing the incentive to allocate units for averaging out the remaining, uncorrelated components of the noise.)

The expression for the SRT model's RT variance (Supplementary Eq. 27) shows that increasing the drift A will diminish the variance, all other parameters being held constant. Furthermore, when A becomes very large, A^3 in the denominator of



Figure 7: **a:** Basic interval timing network, showing typical interconnection and recurrent connection strengths (for $\alpha = 4$ in Eq. 5, main text). **b:** Example of a network that implements variance reduction. The simple timing circuit in A is re-implemented by *N* redundant ramp units and *M* redundant trigger units. RT is defined as the first time at which average trigger unit activation exceeds an arbitrary value between roughly 0.5 and 0.8, and the threshold value is defined as the average ramp unit activation level at the time of this event (individual trigger units are leaky integrators, while the trigger population as a whole is bistable; when trigger units' recurrent connection strengths are not reduced in proportion to *M*, then bistability in individual units makes noise reduction even more difficult). **c-f:** Statistics of the threshold implementation as a function of trigger population size *M*, with deterministic ramping ($N = \infty$), and noisy interconnections within the trigger population ($c = 1, \alpha = 4, \beta = 1.3, w_{ii} = 2/M, \tau = 1, I = 0.05$ in Eq. 5, main text). **c:** The value of the threshold for this particular simulation increased asymptotically toward a constant and became less variable as the number of units *M* increased (error bars indicate s.e.m.). **d:** The variance of the threshold distribution of final ramp values) dropped in proportion to 1/M. **e:** When this threshold distribution is applied to a noiseless ramp, it results in a coefficient of variation (CV) equal to standard deviation/threshold, giving rise to a psychologically plausible range of CVs (0.08–0.3). **f:** The RT distribution is relatively constant in location and dispersion, except for small *M*; error bars indicate RT s.e.m.

the first term of Supplementary Eq. 27 makes this term negligible for large A. This leaves a variance approximately proportional to $1/A^2$, hence standard deviation proportional to 1/A. In fact, we get scalar invariance as long as the threshold z is held constant for all different interval lengths, and A is reduced for longer intervals, since interval duration in that case is also inversely proportional to A. However, for long intervals (and small values of A), the first term will dominate the variance formula, and the first-passage time distribution will take on a greater skew. Scalar invariance will then be violated, since the CV will begin to grow as the square root of the mean interval duration rather than remaining constant. This implies an upper limit on the durations that can be timed with proportional precision.

Supplementary Eq. 27 does not rule out encoding durations by threshold adaptation, or by simultaneous drift and threshold adaptation, since both of these approaches can lead to accurate mean duration estimates. However, since A represents the rate of increase of a physical quantity, it is intrinsically bounded and has a maximum value at which approximate linearity holds. More importantly, the A^3 term in Supplementary Eq. 27 makes it far easier to reduce variance by increasing A than by decreasing z. Since we want to maximize A, and since z must always equal A times duration, the best strategy is therefore to hold z fixed at the maximum value that the ramp can attain in its linear range, and to adapt drift alone in order to time different intervals.

Assuming that *A* is maximized and *z* remains fixed, the best way to reduce variance further in Supplementary Eq. 27 is to reduce the noise coefficient *c* in some way (we assume for simplicity that *c* is the same for every interconnection in the model, although it is potentiated by the connection strength w_{ij} ; see Eq. 4). The simplest way to do this is by averaging. By increasing the number of units that perform each task (ramping or triggering) and computing a weighted sum of their activations, the same ramp and trigger functions can be carried out with the effective noise coefficient thereby reduced. Specifically, for *N* independent drift-diffusion processes, the variance of the average process \overline{X} is the following:

$$\operatorname{Var}(\overline{\mathbf{X}}) = \frac{1}{N^2} \sum_{i=1}^{N} \operatorname{Var}(\mathbf{X}_i) = \frac{1}{N} \operatorname{Var}(\mathbf{X}_i).$$
(28)

Simulations over a range of parameters show that the variance of the threshold distribution implemented by a fully connected, *M*-unit trigger population, with interconnection strengths equal to 2/M (as in Supplementary Fig. 7 B), is proportional to 1/M (see Supplementary Fig. 7d). The average behavior of such a population of units with equal connection strengths is the same as that of a single unit with self-excitatory strength k = 2,



Figure 6: a, b: State space plots of a neural network unit's rate of activation change (dV/dt in Eqs. 4-5, Materials and Methods) for units with fixed external input and balanced (a) or strong (b) excitatory, recurrent connections. Equilibrium curves are solid; velocities dV/dt are indicated by arrows and shading (light = positive, dark = negative). c: The catastrophe manifold formed by the equilibrium curves of Eq. 6, main text, as the self-excitatory, recurrent weight strength k ranges from 0 to 2; and the corresponding surface formed by approximately normal threshold densities at each value of k. Three network symbols are also illustrated: an integration symbol corresponding to a recurrent connection strength tuned to achieve perfect integration (k = 1 for a unit with $\alpha = 4$ in Eq. 6); a symbol representing the unit's exponentially decaying impulse response in the leaky integration regime (k < 1); and a sigmoidal symbol representing recurrent excitation that is strong enough to produce hysteresis and bistable switching (or 'latch') behavior (k > 1). **d:** A latch based on hysteresis. Solid curves denote stable attractors. Where a dashed curve is plotted, two possible attractors exist. States above the dashed curve converge to the upper solid curve; states below converge to the lower solid curve. This latch can store a 1 (upper gray region) or a 0 (lower gray region) as long as the input is held between A and B, and will be least susceptible to bit-flipping during a hold period at the midpoint, (A + B)/2.

which is strongly bistable and hysteretic when $\lambda = 4$ in Eq. 4, main text.

The system now faces a resource allocation problem: Given R = N + M units to work with, should more units be used to implement ramps or triggers? In this case, the formula for RT variance (Supplementary Eq. 27) becomes:

$$\operatorname{Var}(t) = \frac{c^2 z}{N \cdot A^3} + \frac{\alpha^2 c^2 z^2}{M \cdot A^2}.$$
(29)

Supplementary Eq. 29 allows us to find the value of M that minimizes variance for any given interval duration. The solution M_{opt} is the following:

$$\frac{M_{\text{opt}} = \dots}{\frac{-2R\alpha^2 z^2 A \pm \sqrt{(2R\alpha^2 z^2 A)^2 + 4(z - \alpha^2 z^2 A)(R^2 \alpha^2 z^2 A)}}{2 \cdot (z - \alpha^2 z^2 A)}}.$$
 (30)

Supplementary Fig. 8 shows a representative case of RT standard deviation plotted as a function of trigger allocation M for a set of durations. It is clear that for small A, corresponding



Figure 8: The envelope of optimal trigger allocations (red curve) superimposed on a sequence of black curves representing the standard deviation of duration estimates as a function of the number of units out of 1000 allocated to the trigger mechanism (with the remainder allocated to the ramp). Black curves correspond, from bottom to top, to a logarithmically increasing sequence of interval durations: 7.5, 15, 30, 60, 120 and 240 seconds. Intersections of the red curve with the black curves indicate the variance minimizing trigger allocations. Allocations are based on an available pool of 1,000 units with noise coefficient c = 1, and threshold-noise potentiation factor K = 0.1 (Supplementary Eq. 26). The envelope continues leftward as interval duration goes to infinity, with the optimal allocation to the trigger mechanism equalling 1 at some finite interval. Furthermore, the range of feasible allocations which produce a standard deviation less than the expected value collapses to 0 at durations greater than or equal to $(z^2 R)/c^2$ (where R is the size of the available resource pool). In general, allocating almost all units to the ramp mechanism for all intervals results in nearly minimal overall RT variance.

to long intervals, the best value for M is small. The optimal trigger allocation M_{opt} in fact approaches a single unit as interval duration grows large (longer than 120 seconds). Furthermore, when interconnection noise is positively correlated (e.g., with pairwise correlations of 0.1), simulations suggest that M_{opt} is even smaller than in the case of independent noise, and that RT variance rises much more steeply as a function of M.

When sufficiently many units are allocated to the ramp mechanism, but the trigger population size remains roughly constant (as variance minimization requires), the SRT model approximates the deterministic ramp model with a deterministic ramp. The model therefore accounts for scalar invariance over a range of durations whose extent depends on the number of units available to act as integrators. This range is bounded below by the tendency of thresholds to blow up as interval durations become very short. It is bounded above by the tendency of diffusive noise to grow relative to drift as interval durations grow very large, resulting in growing CVs. This property is in fact consistent with evidence for non-constant CVs in the literature (cf. Bizo et al., 2006).

4.6. The effect of correlated noise on RT variance

Correlations among the synaptic noise sources of the SRT model invalidate the independence assumption on which our analytical approximation to its RT variance was based. To determine whether the model's behavior was robust to violations



Figure 9: The effect of correlated noise on the coefficient of response time variation as a function of trigger-unit proportion.

of this independence assumption, we explored the effect of correlated noise through computer simulations. One independent source of noise was associated with each interconnection in the model; then a weighted sum of noise was added to the signal transmitted through each connection, with a strong weight on the corresponding noise source, and weaker weights on the noise sources corresponding to other connections (the correlation weighting matrix consisted of 1s along the diagonal, and smaller values off the diagonal).

Because the correlation matrix consumed a great deal of memory, we were limited to simulating populations that were small relative to the population sizes simulated with independent noise. Supplementary Fig. 9 nevertheless shows a clear and consistent effect of correlations. There, the coefficient of RT variation is plotted as a function of the proportion of 60 total units allocated to the trigger layer, for three different interval durations (drift rates) and three different correlation strengths.

As the noise correlation increased, the CV typically increased at all trigger allocations while the mean RT remained roughly constant. Importantly, the CV increased more rapidly for large trigger allocations. Thus the incentive to minimize variance by allocating more units to the SRT model's ramp layer appears even stronger than when noise is independent.

5. Evidence for abrupt duration learning

In order to assess whether rapidity and abruptness in adjusting to new temporal parameters could be observed in individual, timed responses under positive reinforcement, we reanalyzed the data reported by Balci et al. (2008). In one experiment (Experiment 2), four mice (C57BL/6N) were trained to respond at



Figure 10: Switch latencies observed in three different phases. Each plot shows the data from an individual subject. The coordinates of the dots indicate the switch latencies (y-axis) and the corresponding trial number (x-axis). When a switch was not observed the corresponding trial was not assigned any symbol. The vertical tick marks located at the bottom of each plot indicate the short trials. The vertical dashed lines indicate the boundaries between phases, i.e., trials at which the new duration pairs were introduced.

one location in anticipation of the reinforcement after a short interval (in 'short' trials) and at another location in anticipation of the reinforcement after a long interval (in 'long' trials). Mice self-initiated their trials. Short and long trials were presented in random order. Trial type was not signaled by any discriminative stimulus and thus the subject could only rely on its estimate of elapsed time in distributing anticipatory responses across two response locations.

In this paradigm the subject typically started anticipating the reinforcement at the short latency location, and when sufficient time elapsed without delivery of reinforcement, switched to the long-latency hopper (Balci et al., 2008). Mice were provided with correction trials: that is, if they responded erroneously in a trial, that particular trial type was presented repeatedly until the mouse emitted the correct response, after which, trial types were presented in random order. Mice were trained in three, multiple-session phases. Phase 1 consisted of a 6 s and 18 s pair. Phase 2 involved training with either a 3 s/9 s or a 12 s/36 s pair, and Phase 3 involved the pair not used in Phase 2. The order of training duration pairs and response locations were counter-balanced across subjects. Phase 1 training lasted for 19 sessions, Phase 2 training lasted for 10 sessions and Phase 3 training lasted for 21 sessions. This paradigm has been further used in rats (Ludvig et al., 2008) and humans (Balci et al., 2009) without correction trials.

In Supplementary Fig. 10, we present trial by trial switch latencies (observed in long trials) separately for four mice and for three different phases (delimited by dashed lines; compare to model performance in Fig. 2 of the main text). This figure suggests that each mouse adjusted relatively abruptly to the new duration pairs. Specifically, once they exhibited a switch at a latency in between the two durations of a new pair, their switch latency was more or less stable (also see Fig. 2 in Balci et al., 2008).

Abrupt, immediate changes in behavior can be explained by the learning rules we have presented when they are parameterized to learn durations after a single training exposure (i.e., with the learning rate set to 1). Gradual changes can be explained by a smaller learning rate. In contrast, abrupt, but delayed, changes in behavior after several exposures to a new duration require additional machinery to explain. However, a simple elaboration of these learning rules that could model the data might involve allocating a completely new timer after several exposures to a new duration, with the new timer's learning rate set to 1.

6. Model performance in the Beat-the-Clock task

Supplementary Figs.11 and 12 demonstrate the performance of the model in the Beat-the-Clock task. The model is parameterized to learn from single exposures to new durations, so that on average, the model's ramping integral hits the fixed threshold z at the correct time. It responds at a constant proportion of this time, however, using a response threshold z' that is lower than z. We chose this threshold to give response times that were below the target interval duration by one standard deviation of the model's response time distribution, which is close, but not identical, to the value that maximizes rewards.

Supplementary Fig. 11 shows autocorrelation in the response times that arises as a result of the learning rules. Supplementary Fig. 12 shows the trial-by-trial performance of the model.



Figure 11: Comparison of the autocorrelation of response times from a human participant and the model in the Beat-the-Clock task. The human participant is the same one whose performance is shown in Fig. 6 of the main text.

References

- Balci, F., Freestone, D., Gallistel, C. R., 2009. Risk assessment in man and mouse. Proceedings of the National Academy of Science, USA. 106, 2459– 2463.
- Balci, F., Papachristos, E. B., Gallistel, C. R., Brunner, D., Gibson, J., Shumyatsky, G. P., August 2008. Interval timing in genetically modified mice: a simple paradigm. Genes, Brain and Behavior 7, 373–84.
- Bizo, L. A., Chu, J. Y. M., Sanabria, F., Killeen, P. R., 2006. The failure of Weber's law in time perception and production. Behavioral Processes 71, 201–210.
- Bogacz, R., Brown, E., Moehlis, J., Holmes, P., Cohen, J. D., 2006. The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced choice tasks. Psychological Review 113 (4), 700–765.
- Churchland, M. M., Yu, B. M., Cunningham, J. P., Sugrue, L. P., Cohen, M. R., Corrado, G. S., Newsome, W. T., Clark, A. M., Hosseini, P., Scott, B. B., Bradley, D. C., Smith, M. A., Kohn, A., Movshon, J. A., Armstrong, K. M., Moore, T., Chang, S. W., Snyder, L. H., Lisberger, S. G., Priebe, N. J., Finn, I. M., Ferster, D., Ryu, S. I., Santhanam, G., Sahani, M., Shenoy, K. V., 2010. Stimulus onset quenches neural variability: a widespread cortical phenomenon. Nature Neuroscience 13, 369–378.
- Franklin, G. F., Powell, J. D., Emami-Naeini, A., 1994. Feedback Control of Dynamic Systems. Addison-Wesley, New York.
- Gibbon, J., 1992. Ubiquity of scalar timing with a Poisson clock. Journal of Mathematical Psychology 35, 283–293.
- Gibbon, J., Church, R. M., 1981. Time left: linear versus logarithmic subjective time. Journal of Experimental Psychology: Animal Behavior Processes 7 (2), 87–108.
- Grice, G. R., 1968. Stimulus intensity and response evocation. Psychological Review 75, 359–373.
- Haß, J., Blaschke, S., Rammsayer, T., Herrmann, J. M., 2008. A neurocomputational model for optimal temporal processing. Journal of Computational Neuroscience 25, 449–464.
- Hopfield, J. J., Tank, D. W., 1985. "neural" computation of decisions in optimization problems. Biological Cybernetics 52, 141–152.
- Lim, S., Rinzel, J., 2009. Noise-induced transitions in slow wave neuronal dynamics. Journal of Computational Neuroscience.
- Lo, C.-C., Wang, X.-J., 2006. Cortico-basal ganglia circuit mechanism for a decision threshold in reaction time tasks. Nature Neuroscience 9 (7), 956– 963.
- Luce, R. D., 1986. Response Times: Their Role in Inferring Elementary Mental Organization. Oxford University Press, New York.
- Ludvig, E. A., Balci, F., Longpre, K., October 2008. Timescale dependence in a conditional temporal discrimination procedure. Behavioural Processes 77 (3), 357–63.
- Machens, C. K., Romo, R., Brody, C. D., 2005. Flexible control of mutual inhibition: a neural model of two-interval discrimination. Science 307, 1121– 1124.
- Oppenheim, A. V., Willsky, A. S., 1996. Signals and systems, 2nd Edition. Prentice Hall, New York, NY.
- Rice, J. A., 1995. Mathematical Statistics and Data Analysis. Duxbury Press, Belmont, CA.
- Seung, H. S., Lee, D. D., Reis, B. Y., Tank, D. W., 2000. The autapse: a simple illustration of short-term analog memory storage by tuned synaptic feedback. Journal of Computational Neuroscience 9, 171–185.
- Shadlen, M. N., Newsome, W. T., 1998. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. Journal of Neuroscience 18 (10), 3870–3896.
- Tripp, B., Eliasmith, C., 2010. Population models of temporal differentiation. Neural Computation 22, 621–659.
- Wald, A., 1947. Sequential Analysis. Dover, New York.



Figure 12: Model performance in the Beat-the-Clock task. All plotting conventions are the same as in Fig. 4 of the main text.