

A biologically plausible model of time-scale invariant interval timing

Rita Almeida · Anders Ledberg

Received: 27 July 2009 / Revised: 2 October 2009 / Accepted: 12 October 2009 / Published online: 28 October 2009
© The Author(s) 2009. This article is published with open access at Springerlink.com

Abstract The temporal durations between events often exert a strong influence over behavior. The details of this influence have been extensively characterized in behavioral experiments in different animal species. A remarkable feature of the data collected in these experiments is that they are often time-scale invariant. This means that response measurements obtained under intervals of different durations coincide when plotted as functions of relative time. Here we describe a biologically plausible model of an interval timing device and show that it is consistent with time-scale invariant behavior over a substantial range of interval durations. The model consists of a set of bistable units that switch from one state to the other at random times. We first use an abstract formulation of the model to derive exact expressions for some key quantities and to demonstrate time-scale invariance for any range of interval durations. We then show how the model could be implemented in the nervous system through a generic and biologically plausible mechanism. In particular, we

show that any system that can display noise-driven transitions from one stable state to another can be used to implement the timing device. Our work demonstrates that a biologically plausible model can qualitatively account for a large body of data and thus provides a link between the biology and behavior of interval timing.

Keywords Scale-invariant timing · Mathematical model · Bistability · Noise-driven transitions · Linear death process

1 Introduction

The duration of time intervals between behaviorally meaningful events, such as a stimulus predicting food and the actual access to the food, is known to influence behavior both at short and long time-scales. For example, when food is made conditionally available at a fixed time interval after the previous collection of food (on a so called fixed interval schedule of reinforcement), animals will adapt the responses on single trials to the current interval duration. When the duration of the fixed interval is changed, animals learn to adapt their responses in accordance with the new duration (Ferster and Skinner 1957, Ch. 5). This ability, to change a response behavior as a function of the arbitrary duration of a time interval in the range of seconds to minutes, is often referred to as ‘interval timing’ (e.g. Staddon and Cerutti 2003) to distinguish it from other types of ‘timing’ behaviors (e.g. Mauk and Buonomano 2004). In experimental studies of interval timing the subjects are exposed to time intervals of different durations and the variable of interest is typically how the temporal distribution of responses varies as a function

Electronic supplementary material The online version of this article (doi:10.1007/s10827-009-0197-8) contains supplementary material, which is available to authorized users.

Action Editor: David Golomb

R. Almeida
Intitut d’Investigacions Biomèdiques August Pi i Sunyer
(IDIBAPS), C. Mallorca 183, 08036 Barcelona, Spain
e-mail: ralmeida@clinic.ub.es

A. Ledberg (✉)
Center for Brain and Cognition, Department of Information
and Communication Technologies, Universitat Pompeu
Fabra, Roc Boronat 138, 08018 Barcelona, Spain
e-mail: anders.ledberg@upf.edu

of interval duration. A remarkable fact of the results obtained in these experiments is that they are often time-scale invariant (Gibbon 1977). This means that the temporal distributions of responses for two different interval durations are the same if the time-axis is scaled (divided) by the duration of the interval. Time-scale invariant response distributions have been reported in many different types of interval timing tasks and in several different species including some mammals, fishes, birds (Gibbon 1977; Lejeune and Wearden 2006), and humans (Wearden and Lejeune 2008), indicating that time-scale invariant behavior reflects a fundamental property of the functional organization of perhaps all vertebrates. The aim of the work presented here is to show that time-scale invariance is a generic property of a certain type of computational architecture and moreover to demonstrate how such an architecture can be implemented in the nervous system, possibly accounting for the behavior in interval timing tasks.

There is substantial evidence indicating that cerebellum, basal ganglia and the cerebral cortex are participating in different aspects of interval timing (for reviews see Ivry 1996; Meck 1996; Gibbon et al. 1997; Mauk and Buonomano 2004; Buhusi and Meck 2005), but less is known about the actual neural mechanisms supporting this behavior. One of the more difficult aspects of the data to account for is the substantial range of interval durations over which time-scale invariance has been demonstrated. Indeed in some tasks this range covers two orders of magnitude (Gibbon 1977). This flexibility in timing temporal durations makes it seem implausible that the neuronal mechanisms involved relies on dedicated, and fixed, time constants, something that for example has been postulated in models of timing behavior in the context of classical conditioning (Grossberg and Schmajuk 1989; Fiala et al. 1996). It rather seems more likely that interval timing behavior results from a system having time constants that can be changed depending on the reinforcement history of the organism (cf Killeen and Fetterman 1988; Machado 1997).

Several models of interval timing have been previously proposed. Some models aim primarily at accounting for the behavioral data and are therefore often not concerned with implementational issues (e.g. Gibbon 1977; Gibbon et al. 1984; Killeen and Fetterman 1988; Machado 1997; Staddon and Higa 1999). These are typically abstract models of general mechanisms that can explain the main features of the behavioral data such as time-scale invariance. Other models primarily aim at accounting for some particular aspects of neural activity believed to be involved in interval timing (e.g. Kitano et al. 2003; Durstewitz 2003; Reutimann et al.

2004). These are concrete models of neural mechanisms and therefore not concerned with behavior. However, a more complete understanding of interval timing must come through models that bridge the behavioral and neural levels. One model that both postulates a general (abstract) mechanism of interval timing and provides anatomical and physiological details about the implementation has been proposed by Meck and co-workers (e.g. Matell and Meck 2004; Meck et al. 2008). However, it remains to be seen if this model can be implemented in a more realistic setting. We will return to this and other models in Discussion.

In the work presented here we describe and analyze a model of an interval timing device, the “stop-watch”, and demonstrate that it is consistent with time-scale invariant behavior over a substantial time-range. This device consists of a set of bistable units that switch from one state to the other at random times and the rate at which they switch will be used to “encode” time intervals of different durations. We will describe two versions of the stop-watch below. First an abstract version is described where the units are modeled probabilistically. This abstract version can be fully analyzed by elementary methods and will be used to isolate a minimal mechanism of time-scale invariance. The second version of the stop-watch will implement the same mechanism of interval timing in terms of variables that could be instantiated in the nervous system. This implementation is generic in the sense that it could be realized at different levels of nervous activity. We will give two examples illustrating this fact: first each unit in the stop-watch will be modeled as a single neuron; in the second example each unit will be modeled as a micro-circuit.

To use a set of bistable units with random state changes in the context of interval timing was first suggested some while ago (Miall 1993). More recent models have suggested specific implementations of this basic setup and shown that it can support time-scale invariance (Okamoto and Fukai 2001; Miller and Wang 2006). In particular, in Miller and Wang (2006) the authors suggest a particular architecture that supports time-scale invariance through the same mechanism we propose here. Our work extends and generalizes the work of Miall (1993) and Miller and Wang (2006) in several ways. We introduce an abstract and general formulation of the original idea, through which it is possible to derive simple analytical expressions for the variables of interest and demonstrate a general form of time-scale invariance. We show that this abstract model can be implemented in the nervous system through a generic mechanism and thus establish a strong link between the behavioral and neural levels.

2 Overview of the computational architecture

The stop-watch has the following general characteristics: (i) it consists of M identical bistable units that can switch from one state to the other at random times. We will refer to these two states as *spontaneous* and *activated* respectively; (ii) The rate at which the units switch from the spontaneous state to the activated state (the activation rate) is the same for all units, and is much higher than the rate of switching in the other direction; (iii) The activation rate can be adjusted through learning (between trials); (iv) The state of the stop-watch at time t is fully characterized by the number of units in a particular state at this point in time; (v) The states of the units can be simultaneously reset to the spontaneous state by some control signal. On a single trial, the stop-watch works in the following way: At trial onset all units are in the spontaneous state and as time (in the trial) progresses, more and more units will switch to the activated state. The total number of units in the activated state will influence the temporal aspects of responding. Through learning, the activation rates will be modified to be appropriate for the task at hand.

3 The abstract stop-watch

In this stop-watch version each unit has a probability p of switching from the spontaneous state to the activated state (the activation probability) within a small time-interval Δt (see Fig. 1(a)). The states themselves are not explicitly modelled. The activation probability p is assumed independent of how long time the unit has been in the spontaneous state. The units are in this sense memoryless or Markovian. Once a unit is activated, the probability of it switching back to the spontaneous state is much smaller than p and will be considered zero in the following. A small non-zero value of this probability would not change the results qualitatively. The state of the stop-watch at a particular point in time is fully characterized by the number of activated units at this time. A realization of one “trial” of the stop-watch is shown in Fig. 1(b). Note the random number of time steps between individual activations. The activation probability p determines the temporal durations the stop-watch can measure: if p is small it takes (on average) a long time until a certain number of units become activated, whereas if p is large, the time until the same number of units become activated will be smaller. This is illustrated in Fig. 1(c) where the state of the stop watch is shown as a function of time for two different values of p . Two things are noteworthy: first, by making

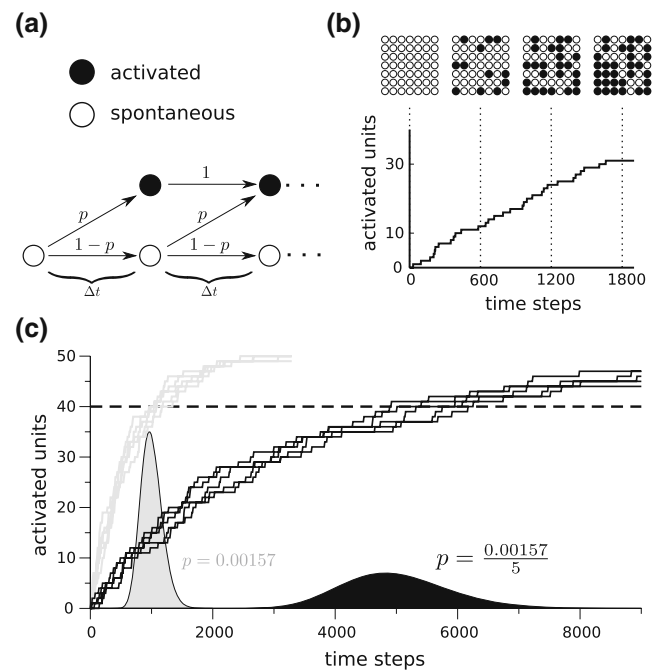


Fig. 1 Schematic illustration of the abstract stop watch. (a) The units of the stop watch have two possible states: spontaneous (*empty circles*); activated (*filled circles*). The switching from the spontaneous to the activated state occurs randomly according to a probability p in each time interval Δt . Once in the activated state, the probability of remaining in this state is close to one. (b) Time-evolution of a stop-watch consisting of 49 units (i.e. $M = 49$). A single trial with $p = 0.0005$ is shown. *Top*: four snapshots of the stop-watch at consecutively later times in the trial. At trial start all the units are in the spontaneous state and as time progresses more and more units become activated. *Bottom*: Detailed time course of the state of the stop-watch for the same trial. (c) Interval duration can be encoded in the activation probability. Five trials are shown for two different probabilities ($p = 0.00157$ in gray; $p = 0.00157/5$ in black) and $M = 50$. The histograms show the distribution of time points at which 80% of the units (indicated by the dashed line) are activated. The histograms were scaled (by the same factor) to be visible

the probability five times smaller, the time it takes until a given threshold is roughly five times longer (compare gray and black curves in Fig. 1(c)); second, the distribution of time points for 80% of units activated becomes wider as p becomes smaller (compare gray and black histograms in Fig. 1(c)). In the next section we will quantify the relation between p and interval duration and show that the inverse proportionality seen in the figure indeed holds true in general.

Since we assume that the activation probabilities are small we can describe each unit succinctly by an exponentially distributed random variable, U_i say, with a rate parameter that equals p (see Appendix). Thus,

U_i is the time at which unit i becomes activated. Using this formulation it is straight-forward to derive analytical expressions for various aspects of the stop-watch and moreover to demonstrate the general time-scale invariance of the stop-watch. This we do next.

3.1 The relationship between interval duration and activation probability

Time intervals of different durations will be estimated by changing the activation probability p of the single units. Hence, we need to understand how this probability relates to time. Since the activation times are assumed to be random, the time-evolution of the state of the stop-watch will (in general) not be the same on two different trials (see Fig. 1(c)). We therefore characterize the relation between p and time using the *expected time* until a certain fraction f of units have made a transition. Let $T_{f,M}$ denote the time it takes for $f \cdot M$ units to become activated, where M is the total number of units. Note that $T_{f,M}$ will differ from trial to trial: it is a random variable. Figure 1(c) shows two distributions of $T_{f,M}$ for $M = 50$ and $f = 0.8$ (the gray and black histograms). In Appendix we show that the expected value of $T_{f,M}$ is given by

$$E[T_{f,M}] = \frac{1}{p} \sum_{k=0}^{f \cdot M - 1} \frac{1}{M - k} \tag{1}$$

The importance of this expression is that it shows that time and probability are inversely proportional. Hence, if it takes, on average, 1 s for 80% of the units to go to the activated state (i.e. $f = 0.8$) when $p = p'$, it will take, on average, 10 s for 80% of the units to become activated when $p = p'/10$. An example of the inverse proportionality between probability and time was shown in Fig. 1(c), where decreasing p by a factor of five lead to an increase in $T_{f,M}$ by the same factor. The relationship between transition probability and time is illustrated in more detail in Fig. 2(a).

Next we consider the variability of the stop-watch. On any particular trial $T_{f,M}$ will typically differ from its expected value (given by Eq. (1)) and it is of interest to know *how much* it will typically differ, and how this depends on the parameters. In Appendix we show that the standard deviation of $T_{f,M}$ is given by

$$S[T_{f,M}] = \frac{1}{p} \left(\sum_{k=0}^{f \cdot M - 1} \frac{1}{(M - k)^2} \right)^{1/2} \tag{2}$$

Note that the standard deviation is also inversely proportional to p . Figure 2(a) shows the standard

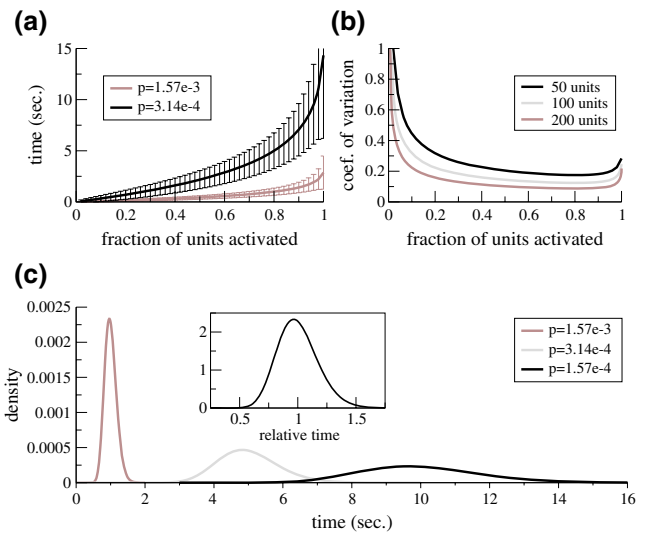


Fig. 2 The time it takes for a certain fraction of units to become activated. (a) Expected values of the time it takes for a certain fraction of units to be activated (Eq. (1)) for two different transition probabilities (solid lines). The error bars show two standard deviations around this value, calculated according to Eq. (2). Data shown for $M = 50$. (b) Coefficient of variation as a function of fraction of total units active for three different values of total number of units (M). (c) Density functions of T_{40} for three different values of p (0.00157, 0.000314, 0.000157) for a stop-watch with 50 units. The inset shows the three densities as functions of relative time, that is, after time is rescaled by the nominal intervals (1, 5, and 10 s)

deviation of $T_{f,M}$ as error bars, for two values of the activation probability.

From Eqs. (1) and (2) and Fig. 2(a) it is clear that both the expected value and the standard deviation of $T_{f,M}$ increase with f . Moreover, both expected value and standard deviation depend on the total number of units M as well as on the transition probability p . We next investigate the relative variability of $T_{f,M}$ using the coefficient of variation (CV). The CV is the standard deviation divided by the expected value and hence is a unit-less measure of the relative variability of a random variable (and it is often used in experimental investigations). Given Eqs. (1)–(2), we can express the CV of $T_{f,M}$ as

$$CV = \frac{S[T_{f,M}]}{E[T_{f,M}]} = \frac{\left(\sum_{k=0}^{f \cdot M - 1} \frac{1}{(M - k)^2} \right)^{1/2}}{\sum_{k=0}^{f \cdot M - 1} \frac{1}{M - k}} \tag{3}$$

Note that the CV is independent of p . Since we “encode” time intervals of different duration by changing p , Eq. (3) shows that with M and f fixed, the CV is independent of interval duration. This is one manifestation of the time-scale invariance that we will study

more generally below. In Fig. 2(b) the CV is shown as a function of f for three different values of M . Using Eq. (3) it is easy to show that the CV is always between 0 and 1 and is equal to one if and only if $f = 1/M$. Moreover, for a fixed f the CV is a decreasing function of M , in fact, for M large it will scale as $\sqrt{1/M}$ (see Appendix). This means that the precision of the stop-watch increases as the number of units increases. In Appendix we also demonstrate that for a fixed M , the CV has a minimum when $f \approx 0.8$ (which also can be seen in Fig. 2(b)).

Using Eq. (3) we can get a rough estimate of the minimum number of units in the stop-watch. For example, CVs around 0.3 have been reported in behavioral data (e.g. Gibbon 1977, Fig. 5) and to get such a CV using the abstract stop-watch requires at least 30 units. Of course, the CV in behavioral data can just serve to give a lower bound for the number of units since the function(s) mapping the state of the stop-watch onto an actual response is likely to add additional sources of variability.

Given that the expected value and the standard deviation of $T_{f,M}$ depend on p in the same way it is of interest to also see how the higher moments depend on p . We do this by looking at the probability density function of $T_{f,M}$. If we let $\phi = f \cdot M - 1$, the density is given by (see Appendix)

$$g_{f,M}(t) = p(M - \phi) \binom{M}{\phi} (1 - e^{-pt})^\phi (e^{-pt})^{(M-\phi)}. \quad (4)$$

That is, $Pr\{T_{f,M} \leq t'\} = \int_0^{t'} g_{f,M}(t) dt$. In Fig. 2(c) three densities, corresponding to three different activation rates, are shown. These rates were chosen according to Eq. (1) to correspond to mean time intervals of 1, 5 and 10 s. It is clear that the spread of the distributions increase together with the mean as Eq. (3) already indicated. When the three distributions are plotted as functions of relative time they are in-fact identical as is shown in the inset in Fig. 2(c). Thus the whole distribution of $T_{f,M}$ is scale invariant. This result can be derived directly from Eq. (4) but it will follow from the more general demonstration in the next section.

3.2 Time-scale invariance

In the previous section we showed that the distribution of $T_{f,M}$ is time-scale invariant. We will now demonstrate that this is a simple consequence of the activation-time distributions being exponential, together with the way we choose to encode the duration of different intervals. Indeed, we will show that the distributions of activation times, corresponding to

different interval durations, are identical when considered as functions of relative time. The implication of this is that the scale invariance of the stop-watch holds more generally, and moreover that any function of the stop-watch will be time-scale invariant as well.

Recall that the time it takes for a unit to become activated is assumed to have an exponential distribution. This implies that the probability that it will be activated before time t' , $F_p(t')$ say, is given by

$$F_p(t') = \int_0^{t'} p e^{-pt} dt = 1 - e^{-pt'}$$

To express time in units relative to some fixed duration T , we must change the independent variable to $\tau = t/T$. Since p is a rate (with the units of $1/t$) it must also be expressed in terms of relative time (τ). If we let p' denote this new rate parameter we get $p' = pT$. This gives

$$F_{p'}(\tau') = 1 - e^{-p'\tau'} = 1 - e^{-pT\tau'}$$

Now, if p is inversely proportional to T , $p = c/T$ say, this last expression becomes

$$F_p(\tau') = 1 - e^{-c\tau'}$$

This expression describes an exponential distribution with rate parameter c and is independent of p . This means that if time intervals of different durations are encoded by rate parameters that are inversely proportional to the nominal durations, the stop-watch units will be identical after rescaling time. In other words, the units are time-scale invariant. Since the stop-watch is just a set of such units, it follows that also the stop-watch is time-scale invariant. As an example, assume the stop-watch is used to time two intervals of duration T_1 and T_2 . Assume further that the corresponding activation rates are chosen according to Eq. (1) to be $p_1 = c/T_1$ and $p_2 = c/T_2$. Then if we express the probabilities as functions of relative time we get that $F_{p_1}(\tau) = F_{p_2}(\tau)$.

This argument rests on that the activation rates p are inversely proportional to the interval durations. This is a natural requirement, since we know from Eq. (1) that this is how the probabilities should be chosen to make $E(T_{f,M})$ coincide with the nominal time-interval durations.

An important consequence of this general scale invariance is the following: if the temporal aspects of responding are controlled by the stop-watch, responding will be scale invariant *independently* of the details of the control. Phrased differently, the scale invariance does not depend on the particular form of the “read-out” of the stop-watch. In Fig. 2(c) we used a simple threshold

to convert the state of the stop-watch to response onsets. The inset in this figure shows that the distribution of $T_{f,M}$ is scale invariant. The argument above explains why this is so. That scale invariance is a more general property is illustrated in Fig. 3 where another possible implementation of how a response behavior is triggered by the stop-watch is shown. For this figure we used a stop watch with 50 units and simulated data for three values of the activation rate p corresponding to interval durations of 1, 5 and 10 s. The rate of responding was taken as a sigmoidal function of the state of the stop-watch. In particular we used the following function to relate the response rate ($r(t)$) to the state of the stop-watch ($X(t)$):

$$r(t) = \frac{k}{1 + \exp(-\alpha(X(t) - \beta))} \quad (5)$$

The constants k , α and β were chosen to make the response curves look reasonable ($k = 4$, $\alpha = 0.25$, $\beta = 45$). Note that in this case response rate is changing with time, something often observed in behavioral data from interval timing tasks. To illustrate the scale invariance in this case, we plot, in Fig. 3(a), the mean (solid lines) and standard deviation (dotted lines) of the response rate as a function of time for the three time intervals. Figure 3(b) demonstrates that these functions are scale invariant, i.e. that they coincide when plotted as functions of relative time. However, according to the demonstration above, the scale invariance should apply to the whole distribution (which in this case is two-dimensional). The inset in Fig. 3(b) shows the response




Fig. 3 Another manifestation of time-scale invariance. **(a)** Output of stop-watch after a sigmoidal transformation (Eq. (5)) for three interval durations (*Brown* 1 s; *gray* 5 s and *black* 10 s) as a function of time. *Solid lines* show the mean “response rate” and *dotted lines* the standard deviation of the response rate. **(b)** Same data as in **(a)** after time has been rescaled by the interval duration. Data based on 20000 simulations of a stop watch having 50 units using a step size (Δt) of 1 ms. The activation probabilities (rates) for the three time intervals were: 1.5702×10^{-3} ; 3.1405×10^{-4} and 1.5702×10^{-4} and were calculated according to Eq. (1). The *inset* in **(b)** shows the relative frequency of the response rates at a relative time of 80% of the interval duration for the three intervals

rate histograms when 80% of the interval duration has expired. Note the high degree of overlap between the three histograms as predicted by the theory.

3.3 Interactions between units

In a biological implementation of the stop-watch the units may or may not interact with each other and it is therefore of interest to consider what type of interaction would be consistent with time-scale invariance. Interactions between the units can be modeled by making the activation rates depend on the state of the stop-watch. In the general case we can let the activation probability of unit i at time t be modulated by the state of some of the other units in the stop-watch. If this modulation is multiplicative, exact scale invariance will still hold. That is, if we denote the state of unit j at time t by $x_j(t)$ the following rule will, for reasonable choices of the functions f_i , preserve time-scale invariance:

$$p_i(t) = p \cdot f_i(x_1(t), x_2(t), \dots, x_M(t)). \quad (6)$$

Here p can be considered a baseline activation rate common to all units, and changing p will make the stop-watch “encode” different temporal durations just as in the model without interactions. According to this formulation (e.g. Eq. (6)), in the time-interval between the activation of the k -th and $k + 1$ -th units the activation rates are constant (but possibly different for different units). When the $k + 1$ -th unit becomes activated the activation rates changes according to Eq. (6). Since the activations times are exponentially distributed by assumption, and the exponential distribution is memoryless, the above modification of the activation rates does not affect the scale invariance. The argument made above (for the case of constant activation rates) still applies. Note that within this framework we can accommodate any type of interactions from all-to-all to sparse and random. In the [Electronic Supplementary Material](#) we exemplify this fact and state some analytical results for the all-to-all case.

4 A neuronal stop-watch

Here we will describe a generic mechanism by which the units in the abstract stop-watch can be approximated and implemented in the nervous system. There are several standard models of neuronal activity, both at the single cell and at the local network levels, that have parameter regimes for which this mechanism can be implemented. We give examples of two such implementations towards the end of the section.

The units of the abstract stop-watch have three critical characteristics that we want to keep: the units should be *bistable*, *memoryless*, and the activation rates should be *modifiable*. Bistability means that given a fixed configuration of a system, there are two distinguishable states that the system can be in for any longer amount of time. We will assume that the spontaneous state corresponds to an equilibrium point of the system. An example of a phenomenon that could be modeled as a system being at an equilibrium point is the membrane potential of a single neuron receiving a fixed input not strong enough to bring it to threshold. We do not make assumptions about the nature of the activated state (it could for example be an equilibrium point or a limit cycle) but we do assume that it is more stable than the spontaneous state. In the generic model described below (Eq. (7)) we will not model the dynamics of the activated state explicitly.

Memorylessness (also called the Markovian property) means in our case that the probability of switching from the spontaneous state to the activated state within a small time-interval should be independent of the time already spent in the spontaneous state. We saw in the previous section that memorylessness (i.e. the exponential distribution of activation times) was a crucial requirement to make the stop-watch scale invariant. We will implement memorylessness in the neuronal stop-watch by introducing a small amount of noise into the system. With the right balance between the noise amplitude and the stability of the equilibrium point, the unit can be made to stay for long times in the spontaneous state but to eventually leave this state and enter the activated state. For the units to be memoryless the time spent in the spontaneous state should be exponentially distributed. It is known that noise-induced escapes from an equilibrium point are in fact approximately exponential (see Appendix).

The abstract stop-watch is used to estimate intervals of different durations by changing the activation probabilities p of the units. We assume that the corresponding parameter in the neuronal stop-watch is the mean input to a unit and that this input can be modified between trials.

A simple model that approximates a bistable memoryless unit with modifiable activation rates is given by the following stochastic differential equation

$$dx = dt(\mu + \beta x^2) + \sqrt{dt}\sigma\xi(t), \quad x(0) = -\sqrt{|\mu|/\beta}. \quad (7)$$

This equation describes the time evolution of the state variable $x(t)$ (the activity level) which is initiated at $-\sqrt{|\mu|/\beta}$. Here $\xi(t)$ is the so-called white-noise process, i.e. a time series of temporally uncorrelated random

variables with a standard normal distribution and the real numbers $\beta > 0$, $\sigma > 0$, $\mu < 0$ are assumed to be constant during a realization (they are parameters, not variables). When considered without noise (i.e. $\sigma = 0$) the initial condition is a stable equilibrium point. All trajectories initiated to the left of $+\sqrt{|\mu|/\beta}$ (which is an unstable equilibrium point) will end up at this point. On the other hand, all trajectories starting to the right of $+\sqrt{|\mu|/\beta}$ will end up in the only other attracting “point”: $+\infty$. To understand the effect of a non-zero noise amplitude it is convenient to think about Eq. (7) as describing a diffusion of a “particle” on an “energy” landscape. The “energy” is given by the negative of the integral of the deterministic part of the right-hand side of Eq. (7). This scenario is depicted in Fig. 4(a). The “particle” starts at the local minimum of the “energy” (at $-\sqrt{|\mu|/\beta}$) and is pushed around by the noise until it eventually makes it past the local maximum (at $+\sqrt{|\mu|/\beta}$) from where on it will rapidly accelerate towards $+\infty$. The region around the stable equilibrium corresponds to the spontaneous state and the region to the right of the unstable equilibrium to the activated state. In Fig. 4(b) four realizations of Eq. (7) are shown, for two different values of μ . The activity stays close to the equilibrium points (~ -0.25 and ~ -0.3 respectively) for a long time until at some point it shoots up and rapidly accelerates towards $+\infty$. The time-point at which the activity diverges is used to model the transition time from the spontaneous state to the activated state (i.e. the activation time).




Fig. 4 Illustrations of the one-dimensional model. (a) Representation of the one-dimensional model in terms of an “energy” landscape. The black and gray curves show the energy of the one-dimensional model for two different values of the input (μ parameter in Eq. (7)). As the input increases the local minimum becomes more shallow and there is a corresponding increase in the probability per unit time of escape from the spontaneous state. (b) Four realizations of Eq. (7) are shown. The gray curves correspond to two trials with $\mu = -0.0117$ and the black curves to two trials with $\mu = -0.0178$. Other parameters were $\beta = 0.1901$, $\sigma = 0.06044$. The dotted line shows the threshold used to detect when a unit becomes activated

We assume that the mean input to the unit is modelled by the constant term in Eq. (7) (i.e. by μ). Changing this input will modify the activation probability of the unit where we take the activation probability to be the inverse of the mean activation time. The mean activation-time can not be expressed in terms of elementary functions (it is given by Eq. (13) below) but for the range of inputs we are considering there is a standard approximation that can be used (see Appendix):

$$p(\mu) \approx \frac{\sqrt{\beta|\mu|}}{\pi} \exp\left(\frac{-8|\mu|^{3/2}}{3\sqrt{\beta}\sigma^2}\right). \quad (8)$$

In Fig. 5 the relationship between μ and activation probability is shown (solid line). This relationship was obtained by solving Eq. (13) numerically. The approximation given by Eq. (8) is shown as dashed lines in the figure. For the values of input we will use ($\mu < -0.01$) this approximation is very good and as μ becomes more negative the error in the approximation goes to zero. In the “energy-landscape” representation of Eq. (7), the effect of changing the inputs (i.e. μ) corresponds to making the local minimum in the energy more or less shallow. Decreasing μ will make the well around the local minimum deeper and decrease the activation probability. Increasing μ will have the opposite effect. In Fig. 4(a) the energies corresponding to two different values of μ are shown illustrating this point.




Fig. 5 Activation probability as a function of input (μ in Eq. (7)). The two plots show the same data but with different scaling of the Y axes. Activation probability was taken as the inverse of the mean escape time from the spontaneous state. The mean escape time was obtained by solving Eq. (13) for different values of μ (solid lines). The approximation given by Eq. (8) is shown as dashed lines. The other parameters of the model were $\beta = 0.1901$, $\sigma = 0.06044$

Each time the model described by Eq. (7) is integrated it will take a random amount of time before the system enters the activated state. These activation times should be exponentially distributed for the units to be memoryless and hence behave as the units in the abstract model. Figure 6 show two activation time densities corresponding to two different levels of mean input. The two distributions are almost straight lines on the logarithmic plot indicating that they are close to exponential. However, for short times there is a clear deviation from exponential distributions as can be seen in the inset in Fig. 6. This deviation comes, at least in part, from the fact that it takes a finite amount of time to get from the initial condition at the stable equilibrium to the threshold. We investigated how this deviation from exponential activation-time distributions affect the time-scale invariance through numerical simulations. Each unit in the stop-watch was modelled by Eq. (7). We fixed the values of σ and β and changed the term corresponding to the input (i.e. μ). Simulations of a stop-watch consisting of 50 units were run for inputs (μ) corresponding to time intervals of 1,2,5,10 and 100 s. As “response onset” we used the time-point when 40 units have become activated. The results are shown in Fig. 7. The “response” distributions in this implementation are well described by the theoretical distribution given by Eq. (4) as can be seen in Fig. 7(a) and (b). For the shortest interval (1 s) there is a slight deviation from the theoretical distribution, the distribution resulting from using Eq. (7) is slightly steeper than the theoretical one (Fig. 7(a)). For the 10 s interval the correspondence between the simulated and




Fig. 6 Activation time densities of the one-dimensional model. Probability densities of activation times for the two values of μ used in Fig. 4(b). The inset shows the same data at a higher magnification. Parameters were as in Fig. 4(b). The density functions were obtained by solving Eq. (14) numerically (see Appendix)

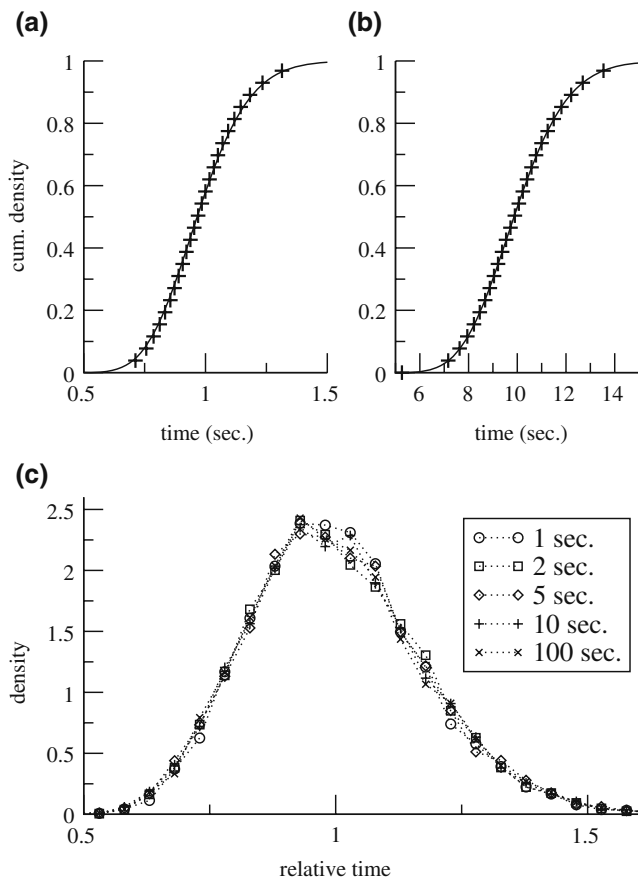


Fig. 7 “Response” distributions for the 1D model (Eq. (7)). (a)–(b) *Black crosses* show numerical estimates of the distribution of the time until 40 units have made a transition for two different mean durations (and values of input, μ). *Solid lines* shows the corresponding theoretical distributions calculated from Eq. (4). (a) $T_{40} = 1$ s; $\mu = -0.0117$. (b) $T_{40} = 10$ s; $\mu = -0.020$. (c) Estimated “response” densities shown as functions of relative time for five different values of interval durations (and μ). For 1 and 10 s the data are the same as that used in panel (a) and (b). For the other intervals the following values of μ were used: 2 s: -0.0146 ; 5 s: -0.0178 ; and 100 s: -0.0265 . All estimates are based on 8000 simulations for each value of μ . The transition probabilities for the five time intervals were calculated according to Eq. (1). The corresponding inputs to the model (Eq. (7)) were found by solving Eq. (13) numerically. The fixed parameters of Eq. (7) were: $\beta = 0.1901$, $\sigma = 0.06044$

theoretical distributions is very good (Fig. 7(b)). As a measure of how well the simulated distributions can be described by the theory we used the CV. For the theoretical distribution the CV is 0.175 independently of interval duration. For the simulated distributions the CVs were 0.168 ± 0.0006 , 0.173 ± 0.0006 , 0.174 ± 0.0006 , 0.174 ± 0.0006 , 0.175 ± 0.0006 , for 1,2,5,10 and 100 s respectively. The standard deviations (numbers after the \pm sign) were estimated by resampling with

replacements from the simulated distributions and are almost identical to standard deviations of the CVs resulting from sampling from a normal distribution. These small differences between the distributions are almost not noticeable when data are shown as functions of relative time (Fig. 7(c)). It is clear that this stop-watch version is time-scale invariant. Note that for durations of 5 s or longer the distributions are almost perfectly predicted by the theory. This is so because the distributions of transition times for the single units are almost exponential for 5 s and will converge to the exponential distribution as the duration increases. From this follows that a stop-watch build with units described by Eq. (7) can time arbitrary large time-intervals in a scale invariant manner. It is clear from the numerical results in Fig. 7 that this particular instantiation of the stop-watch can also time intervals at least as short as 1 s without deviating substantially from the scale invariance.

In the following sections we give two examples of model systems that have approximately the dynamics described by Eq. (7) and are possible implementations of the stop-watch in the nervous system.

4.1 Example 1: single neuron model

In this section we assume that the dynamics of the stop-watch units are described by the two-dimensional Morris-Lecar model (Morris and Lecar 1981; Rinzel and Ermentrout 1998; Izhikevich 2007). This is a model of a point neuron with two voltage dependent non-inactivating currents, one inwards and one outwards, as well as a leak current. In the original paper the inward current modelled was a calcium current and the outward current a potassium current (Morris and Lecar 1981). We will use the version of the model described and analyzed in Izhikevich (2007) in which it is referred to as the ‘persistent sodium plus potassium model’. Hence the inward current is taken to be a sodium current. However, we don’t intend to model any particular cell or system, rather the aim is to show how a model of a spiking neuron can serve as a unit in the stop-watch and moreover that the relevant dynamics (for the stop-watch) of this model is fully captured by Eq. (7) studied above.

The differential equations governing the model are

$$\begin{aligned}
 C\dot{V} &= I - g_l(V - E_l) - g_{Na}m_\infty(V)(V - E_{Na}) \\
 &\quad - g_kn(V - E_K) \\
 \dot{n} &= \frac{n_\infty(V) - n}{\tau_n}
 \end{aligned}
 \tag{9}$$

The steady-states of the so-called gating variables m_∞ , and n_∞ are given by

$$m_\infty(V) = \frac{1}{1 + \exp([m_{1/2} - V]/m_k)}$$

$$n_\infty(V) = \frac{1}{1 + \exp([n_{1/2} - V]/n_k)}$$

Here V denotes the membrane voltage and n is the gating variable of the potassium current. The definition of other terms and values of constants are given in [Appendix](#). We use this model in a regime where there is (in the absence of noise) a stable equilibrium (the spontaneous state) and a stable limit cycle (the activated, spiking, state). When a small amount of noise is added to the input current, the model can be made to spontaneously switch between these two states. Thus, we consider the input current I to be composed by two parts: a mean μ and a zero mean stochastic process $\xi(t)$, i.e. $I(t) = \mu + \sigma\xi(t)$. For the dynamics to show random switching from the spontaneous state to the activated state the system must be close to a configuration in which the spontaneous state becomes unstable. [Figure 8\(a\)](#) shows traces from three different “trials” with the model in such a configuration. In this figure the model parameters were the same, only the noise realizations differed. Note how the point of transition from the spontaneous state to the activated state differs between different realizations of the model.

When a dynamical system is close to a configurations where the qualitative features of the system change (e.g. the spontaneous state becomes unstable), it is often possible to reduce the dimensionality of the system (see for example [Carr 1981](#)). In [Appendix](#) we show that in the parameter regime we are interested in, the Morris-Lecar model (Eq. (9)) can be reduced to the one-dimensional model considered in the previous section (Eq. (7)). Indeed, this reduction explains the choice of parameter values we have used in applications of Eq. (7). [Figure 8\(b\)](#) illustrates this reduction by plotting one trial of both the Morris-Lecar model (black trace) and the reduced model (red trace) when driven by the same noisy inputs. As long as the Morris-Lecar model is in the spontaneous state the approximation by the reduced model is essentially perfect. Also the transition to the spiking state is very well approximated (see inset in [Fig. 8\(b\)](#)). This direct correspondence between the models implies that the distributions of activation times will be essentially the same. This means that a stop-watch where each unit is described by a Morris-Lecar model will be scale invariant over the same range as the one-dimensional model. To illustrate this

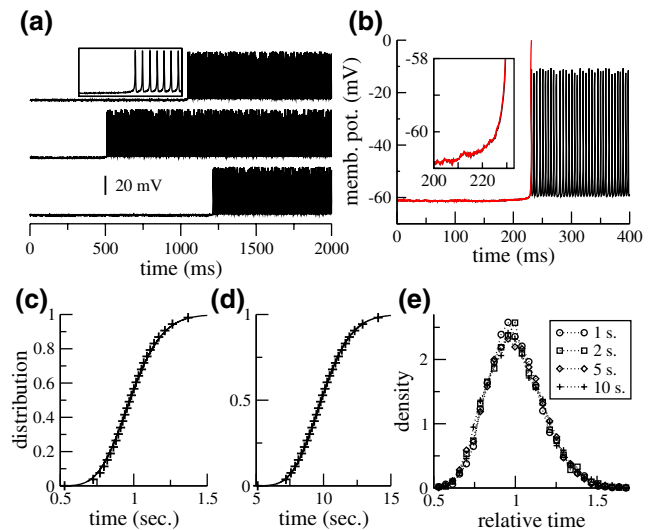


Fig. 8 Stop-watch with Morris-Lecar units. **(a)** Three “trials” of the Morris-Lecar model with the same mean input. The *inset* shows the transition from the spontaneous to spiking state of the top trace at higher magnification. **(b)** Comparison between the Morris-Lecar model and the reduced model (Eq. (7)). One realization of the Morris-Lecar model is shown in *black*. The same noise process was used to drive the reduced model and the result is shown in *red*. The output of the reduced model was truncated at zero. The *inset* shows the transition from the spontaneous state to the spiking state at higher magnification. **(c)** and **(d)** “Response” distributions from simulations of the the stop-watch using Morris-Lecar units for 1 and 10 s intervals. The *crosses* show the distribution estimated from simulations and the *solid line* the theoretical distribution computed from Eq. (4). The mean input in **(c)** was -4.50125 and in **(d)** -4.49296 . **(e)** Estimated “response” densities shown as functions of relative time for four different interval durations. The mean inputs for 2 and 5 s were: -4.49842 and -4.4952 respectively. Data in panel **(c)** to **(e)** were based on 4000 simulations for each value of the inputs. The value of the other parameters were held fixed and are given in [Appendix](#)

we run numerical simulations of a stop-watch having 50 Morris-Lecar units. The ‘read-out’ was taken as a simple threshold, detecting when at least 40 units are activated, i.e. have entered the spiking state. [Figure 8\(c\)](#) and **(d)** show the distributions of these times for two different values of input (corresponding to 1 and 10 s). The distributions are similar to the theoretical distributions (solid lines in the figure). For 1 s duration there is again a small (but significant) difference between the theoretical distribution and the distribution resulting from the Morris-Lecar stop-watch. [Figure 8\(e\)](#) shows that the estimated response densities coincide very well when displayed as functions of relative time. That is, also this version of the stop-watch is time-scale invariant. The CVs of the “response” distributions were 0.165, 0.171, 0.172 and 0.176 for 1,2,5, and 10 s respectively. The

standard deviations of these estimates were estimated to be approximately 0.002.

4.2 Example 2: micro-circuit model

In this section we assume that the dynamics of the stop-watch units are described by a model of a small network of recurrently connected excitatory neurons. We use the circuit model considered in Koulakov et al. (2002) where it was part of a model of a discrete integrator. Each unit consists of three excitatory neurons recurrently connected through NMDA-like synapses. Each neuron in the model has two compartments, one describing the soma and axon and the other describing the dendrites. The soma compartment have the standard spike generating conductances and the dendrites have a voltage dependent synaptic conductance with a slow time-scale (the NMDA-like conductance). The input current to the dendritic compartment contains two parts. One is due to the recurrent connections and the other is an external current. That is, $I = I_{\text{NMDA}} + I_{\text{ext}}$. Moreover we took I_{ext} to be composed of a constant term plus a zero mean stochastic part, i.e. $I_{\text{ext}}(t) = \mu + \sigma \xi(t)$, where the amplitude of the stochastic part (i.e. σ) was taken as 0.25 in all simulations of the model. The equations describing this model are given in the [Electronic Supplementary Material](#) but are in fact identical to the equations described in Koulakov et al. (2002) that may be consulted instead. Note that we are not including the weak connections between units that were used in Koulakov et al. (2002) as we do not need these to generate time-scale invariance. The recurrent connections in this model result in a positive feedback amplification and for the parameters values we use (the same as those used in Koulakov et al. 2002) there are two co-existing stable states: the spontaneous state where all cells are silent (not firing) and the activated state where all cells are firing at a (more or less) constant rate. Note that in this model the bistability is a network effect whereas in the Morris-Lecar example considered above the bistability was a single cell effect.

To investigate the scale invariance we run numerical simulations of the model for four different values of mean input current. With this model it is more cumbersome to derive an explicit relationship to Eq. (7) and therefore to determine the input current that corresponds to a particular mean “response” time. We estimated the “correct” mean inputs by simulations and aimed at staying close to the intervals used in the other implementations. Each unit is initiated in the spontaneous state and evolves until it has entered the spiking state. As the criterion for having entered the spiking

state we used that all three cells fired a spike within a time window of 100ms. The results are illustrated in Fig. 9. For the shortest interval (1 s) there is a slight deviation from the theoretical distribution (Fig. 9(a)) but with a 10 s interval, the distribution is practically indistinguishable from the theoretical one (Fig. 9(b)). That also this model is time-scale invariant is shown in Fig. 9(c) where estimated “response” densities are shown as functions of relative time. The CVs for the simulated “response” distributions were 0.160, 0.166, 0.170, 0.175 for 1,2,5, and 10 s respectively. The slightly worse fit to the theory for the shorter intervals as compared to the other instantiations of the stop-watch (compare the CVs to the ones stated above) can be explained as follows. For the parameters we used




Fig. 9 “Response” distributions for the micro-circuit model. (a)–(b) Numerically estimated “response” distributions (*crosses*) of the time until 40 units have made a transition for two different mean durations (and values of mean external input μ). *Solid lines* show the corresponding theoretical densities calculated from Eq. (4). (a) $T_{40} = 1$ s: $\mu = 4.324$. (b) $T_{40} = 11.1$ s: $\mu = 4.2985$. (c) Estimated “response” densities for four different values of interval duration (and μ) as functions of relative time. The inputs for 2 and 5 s were 4.3148 and 4.305 respectively. The noise amplitude was $\sigma = 0.25$ and values of all other parameter were fixed and taken from Koulakov et al. (2002) and are also given in the [Electronic Supplementary Material](#). The estimates are based on 8000 simulations for each value of the inputs

(identical to those in Koutrakov et al. 2002) it can happen that one cell fires a spike but the other two cells are relatively hyperpolarized and the excitation received is not enough to bring these to threshold. In other words, that one cell fires an action potential it is not always sufficient to make the system enter the activated state. This implies that the distribution of activation times of the units are slightly less exponential than in the Morris-Lecar case. However, Fig. 9(c) clearly shows that the micro-circuit version of the stop-watch is time-scale invariant to a good approximation.

5 Learning to time intervals

To use the stop-watch to control behavior an organism must be able to select the right activation probability for a given interval duration. (In the implementation given above this would correspond to selecting the right level of mean input.) However, organisms typically do not have direct access to the interval duration(s) involved in a particular experimental situation. Consequently, the right probability (or input) must be found through learning. From the point of view of the abstract stop-watch, learning would mean changing the activation probabilities according to some rule to make responding happen at an appropriate time. Here we consider the simple situation that the organism (i.e. model) only knows if it started responding too early (e.g. before food is available) or too late (e.g. after food is already available) on a particular trial. We will describe a simple learning rule that in this case preserves the scale invariance. Learning will introduce a new source of variance into the model that could potentially disrupt the scale invariance. In all the analyzes so far we assumed that the unit activation probability was a constant (for a particular interval duration). If this changes from trial to trial (due to learning) the overall variability of the stop-watch will be increased, and we must make sure that it increases in the right way.

Assume there are two target intervals of duration T_1 and T_2 that should be learned. Assume further that the corresponding target probabilities, p_1 and p_2 , are such that $E_{p_1}(T_{f.M}) = T_1$, and $E_{p_2}(T_{f.M}) = T_2$ for some fraction f . For each target interval duration the change in p due to learning should lead to a change in $E(T_{f.M})$ that is proportional to $E(T_{f.M})$. This would imply that the added variability would be proportional to the mean, as required by the time-scale invariance. Given the inverse relationship between interval duration and probability (Eq. (1)) it is straight forward to device such a learning rule. Let t_k denote the onset of responding on the k -th trial, and let $0 \leq \beta \leq 1$ be

a learning rate parameter. Then the following learning rule will have the desired effect:

$$p_{k+1} = \begin{cases} \frac{p_k}{1 + \beta} & \text{if } t_k < T \\ \frac{p_k}{1 - \beta} & \text{else} \end{cases} \quad (10)$$

In words: if responding starts too early, make p smaller by a factor $1/(1 + \beta)$, and if responding starts too late make p larger by a factor $1/(1 - \beta)$. Given the inverse relation between p and T this learning rule will change the current expected time interval (corresponding to the current p) in a proportional way. We evaluated this learning rule in numerical simulations of the stop-watch based on the one-dimensional dynamical model (Eq. (7)). In the simulations the stop-watch had to dynamically track changes in the interval duration between 1, 5 and 10 s. To implement the learning rule in this version of the stop-watch we changed the mean inputs (μ in Eq. (7)) on each trial so that the

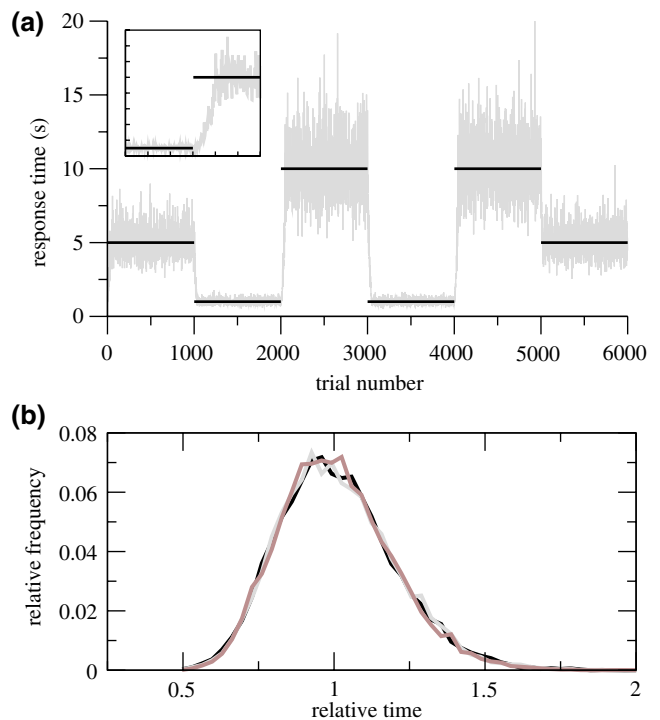


Fig. 10 Learning to time intervals. **(a)** Six-thousand “trials” of a stop watch consisting of 50 units each modelled as Eq. (7). For each trial, the time when 40 units became activated are plotted (gray curve). The target durations (black lines) changed between three levels (1, 5, and 10 s). Learning was implemented as described in the main text and $\beta = 0.05$. The inset shows the first transition from 1 to 10 s (i.e. centered at trial 2000) at higher magnification. Tick marks in the inset are: x-axis 50 trials; and y-axis 2 s. **(b)** Histograms of response times plotted as functions of relative time for the three interval durations used in **(a)** (brown 1 s, gray 5 s and black 10 s). Histograms were made from 10000 simulations for each interval duration

activation probabilities changes consistently with Eq. (10). This was achieved by using the approximation given by Eq. (8). Figure 10 shows that the learning rule indeed converges to the correct values and can thus be used to dynamically trace changing interval durations. The learning rate used in the figure, $\beta = 0.05$, was large enough for the system to learn a new interval duration in less than 50 trials (see inset in Fig. 10(a)) and small enough to not increase the variability too much when compared to the model without learning. Figure 10(b) shows that the time-scale invariance is completely preserved under this learning rule (and value of β).

6 Discussion

Time-scale invariance is a prominent property of behavioral data collected in a wide range of interval-timing tasks and in many different species (e.g. Gibbon 1977; Lejeune and Wearden 2006; Wearden and Lejeune 2008). We have studied a minimal model (the stop-watch) of a device that could be involved in the generation of such time-scale invariant behavior. The stop-watch consists of a number of bistable units and timing behavior is a function of the number of units in the activated state. Each unit starts in a spontaneous state and spends a random amount of time in this state before becoming activated. The distribution of these so-called activation-times, and the way this distribution depends on the duration of the time interval to be estimated, is the key to scale invariance. We first analyzed an abstract model where the activation times were modelled as exponentially distributed random variables. Using this formulation it is straightforward to demonstrate scale invariance and derive simple expressions for the mean (Eq. (1)), standard deviation (Eq. (2)), and density function (Eq. (4)) for the time until a certain fraction of units are activated. The scale invariance is an intrinsic property of the stop-watch and is therefore relatively insensitive to exactly how the state of the stop-watch is used to control behavior. This was illustrated by showing that two different “response behaviors” triggered by the stop-watch were both scale invariant (Figs. 2 and 3). We then described a generic mechanism by which bistable units with exponential activation times could potentially be obtained in the nervous system. In particular, we suggest noise-induced transitions from an equilibrium point (spontaneous state) to another more stable state (activated state) as a possible way to implement the abstract stop-watch (Eq. (7)). This mechanism is generic in the sense of being compatible with many different models of neural

activity. We illustrated this genericity by implementing the mechanism in two different models: one at the single cell level (Morris-Lecar model, Fig. 8), the other at the cortical micro-circuit level (Fig. 9). We also note that the suggested mechanism could be implemented in models both at finer scales, e.g. purkinje cell dendrite (Genet and Delord 2002) as well as coarser scales, e.g. large recurrently connected cortical networks (Amit and Brunel 1997; Hansel and Mato 2003). That is to say, these models also have parameter regimes where the relevant dynamics can be described by Eq. (7). Given that the physiological mechanisms underlying interval-timing are poorly understood at this point, focusing on a generic mechanism that accounts for behavioral data and is consistent with neurophysiology is certainly justified. More specific and detailed models will be called for once there is data with which to constrain such models.

In the following we will discuss certain aspects of the model and in particular how it can be extended in various ways. We will then relate our work to previous modelling work and experimental findings and speculate about how and where a stop-watch might be implemented in the nervous system.

6.1 Criticism and extensions

The suggested mechanism of interval timing relies on that one parameter in the model is changed as the nominal interval duration is altered. In the abstract model this parameter was the activation probability (or rate) and in the neuronal model it was the mean input to the units. Since it is well known that animals can learn to adapt the inputs to single neurons on the basis of the reinforcement history (e.g. Fetz 1969), changing the activation probability by changing the input is definitely biologically plausible. We also described a simple learning rule (Eq. (10)) through which the stop-watch can adapt the activation probability as the nominal interval durations are changed, without upsetting the scale invariance. However, this should primarily be seen as a demonstration of that the stop-watch *can* be used to track time-intervals dynamically. In an implementation aiming at fitting real data it is well possible that this learning rule needs to be augmented. To preserve scale invariance, the learning-induced change must be proportional to the activation probability (Eq. (10)). For the neuronal stop-watch this means that the relative change in the input will depend non-linearly on the present state of the system. The change must be smaller (in absolute terms) when the system is timing intervals of long duration than when

it is timing shorter intervals. One possible way that this could be implemented locally is by making the effect of an incoming spike depend on the state of the system. In standard conductance-based models of synaptic interactions this is exactly what happens: the current that flows into (or out of) the cell depends on the membrane potential. However, a detailed learning mechanism compatible with the physiology still has to be worked out.

We focused mainly on the simplest possible versions of the stop-watch. In most parts we assumed that the units are independent of each other and that once a unit is activated it will remain so for the rest of the trial. Both these assumptions can be relaxed without upsetting the scale invariance. One possible extension that we have studied is to consider that the units of the stop-watch are interconnected in a way so that the activation probability will depend on the number of units already activated. We demonstrated that the abstract model is scale invariant if the units interact multiplicatively according to Eq. (6). In the [Electronic Supplementary Material](#) we exemplify this type of connectivity and state some analytical results. A possible biological mechanism that could accommodate the interactions described by Eq. (6) is some form of synaptic plasticity. What is needed is that the strength of the synapses between the units are modulated by the interval duration. Such modulation would presumably take place on a time-scale of minutes to hours. In some behavioral experiments of interval timing this is indeed the time-scale required for animals to adjust their behavior to changes in the duration of the target interval (e.g. Ferster and Skinner 1957). Another, more short-term, implementation of the multiplicative scaling could come about through a balanced modulation of the background inputs. Such modulations can have a multiplicative effect on the output of neurons (Chance et al. 2002). In the [Electronic Supplementary Material](#) we also discuss an alternative, additive, way of connecting the units. In this case exact scale invariance does not hold anymore but as long as the interactions are weak, scale invariance is still a good approximation. Additive connectivity could for example be useful in order to make the state of the stop-watch increase as a more linear function of time and behave more like a linear integrator. Indeed models of neuronal integrators consisting of networks of bistable units have been suggested previously (Koulakov et al. 2002; Okamoto et al. 2007). We have considered that the activation probabilities are much larger than those of returning to the spontaneous state. There is some evidence supporting the validity of this assumption in the case where the stop-watch units would be bistable cells. Indeed, there

are cells in the enthorinal cortex that can be switched from a silent (spontaneous) state to a state of persistent spiking which can last for very long time periods (e.g. Egorov et al. 2002; Tahvildari et al. 2007). However, one could also consider the case where the units can return to the spontaneous state after they have become activated. Such systems can also be used to estimate interval durations in a similar manner to the stop-watch we have studied, but are not as tractable analytically nor are they necessarily more realistic. Note that in the examples we studied, units can actually switch back to the spontaneous state. However, for the parameters we used this is very unlikely to occur. There is still another aspect in which the model could be generalized. We have considered exponential activation times for the individual units. This leads to tractable calculations and a generic implementation. However, for time-scale invariance, exponential distributions are not strictly necessary. In fact, any distribution which is time-scale invariant would work equally well.

In the implementations of the stop-watch we assumed that each unit receives additive, uncorrelated, and normally distributed noise. However, a weak temporal correlation of the noise does not affect the scale invariance (see [Electronic Supplementary Material](#)), nor do reasonable deviations from normality. We note that for a neuron that receives input through a large number of temporally independent synaptic events, the total input can indeed be approximated by a constant mean plus normally distributed noise with a temporal correlation that depends on the synaptic time constants (e.g. Renart et al. 2003, Section 15.2.5). Noisy inputs that are correlated *between units* would not affect the time-scale invariance but could imply that Eqs. (1)–(3) do not hold. Indeed, the effect of correlated noise is similar to that of decreasing the total number of units. We have furthermore implicitly assumed that the mean input can be changed independently of the noise. However, this assumption is not crucial for the time-scale invariance. What is crucial is that the activation rates changes monotonously as a function of the mean input, something that is likely to hold for most systems. Indeed, if the input is modeled as a large number of Poisson spike trains (as above) with a common rate ν the “standard” approximation of the resulting current received by the post-synaptic cell is a normally distributed random variable where the mean and variance scales proportionally with ν (e.g. Renart et al. 2003, Section 15.2). This means that changing the mean input by changing the rate of the incoming spike trains would also change the variance of the white noise process. In this case μ in Eq. (7) would change proportionally to ν and the noise amplitude σ would change proportionally

to \sqrt{v} . Both of these changes would affect the transition times in the same direction. We also note that selective changes of the noise amplitude (without changes of the mean) could also be used as a mechanism to control the activation times of the units. This could for example be achieved by a simultaneous increase in both excitatory and inhibitory inputs.

In Section 3.2 we argued that the time-scale invariance was independent of the particular read-out mechanism. However, this is strictly true only under the assumption that the read-out mechanism have direct access to the true state of the stop-watch. If the outputs of the stop-watch units are noisy, the read-out mechanism would, in general, not be able to directly access the state of the stop-watch. Rather, it would have to infer this state from the noisy outputs. It is possible that the noise added in this inference process could make the ‘response timing’ deviate from scale invariance. In [Electronic Supplementary Material](#) we study one example of such noisy outputs. In particular, we take the two states of the stop-watch units to correspond to Poisson ‘spike trains’ with two different rates. We show that the ‘response times’ of a read-out unit that receives the sum of these outputs are still approximately time-scale invariant. This holds true even if the difference in the output rates is as small as 5 Hz. From this we can conclude that noisy outputs from the stop-watch do not necessarily destroy the scale invariance.

In summary, the times-scale invariance is not critically dependent on any of the assumptions we have made about the units but would for example also hold for interacting units receiving temporally and spatially correlated inputs.

6.2 Relation to previous modeling and experimental work

The idea of using abstract bistable units with random activation times to estimate the duration of time intervals was first introduced by Miall (1993). Miller and Wang (2006) have suggested an implementation of this original idea which is in several aspects similar to the model we present. In particular, their model uses units with exponentially distributed activation rates to achieve scale-invariant behavior and the biological implementation they have in mind is also of the noise-driven saddle-node type. We have studied such models from a more general point of view, and demonstrated that time-scale invariance is a generic feature of such models. Our work is in this sense a generalization and extension of the model proposed by Miller and Wang.

Fukai and Okamoto and co-workers have investigated different realizations of the ideas above

(Okamoto and Fukai 2001; Kitano et al. 2003; Okamoto et al. 2007). Their models are different from our in architecture, in the way time and in particular intervals of different durations are encoded and in assumptions related to the underlying neurobiology. Their more formal models cannot be (or have not been) analyzed by elementary methods and hence the mechanism behind scale invariance is not clear. Further the relations between their more formal and more neurobiological models are not explicit and hence the range of time-intervals where scale invariance is supported in the more neurobiological models is not known.

Time is an important variable in most experiments with behaving animals but only relatively few studies have directly looked at the neural correlates of interval timing (Niki and Watanabe 1979; Kojima et al. 1981; Matell et al. 2003; Leon and Shadlen 2003; Roux et al. 2003; Sakurai et al. 2004; Kalenscher et al. 2006; Oshio et al. 2006; Renoult et al. 2006; Chiba et al. 2008; Lebedev et al. 2008). A common finding in many of these studies is that of specific neural activity that precedes the response. Such activity has been reported in monkey (Niki and Watanabe 1979; Kojima et al. 1981) and pigeon prefrontal cortex (Kalenscher et al. 2006) as well as in motor and premotor cortices (Lebedev et al. 2008). This activity is often monotonously increasing or decreasing with time (so-called climbing activity), with the rate of increase (or decrease) being dependent on the time-interval duration (Kalenscher et al. 2006; Lebedev et al. 2008). Such duration dependent climbing activity has also been found in other tasks involving delay periods of different durations, for example in monkey prefrontal (Kojima and Goldman-Rakic 1982; Brody et al. 2003), and inferotemporal (Reutimann et al. 2004) cortices, as well as rat thalamus (Komura et al. 2001). Another common finding in timing tasks is neurons that have a, more-or-less symmetric, bump of activity with the peak coinciding with a time point of potential response initiation (monkey motor (Roux et al. 2003; Renoult et al. 2006) and parietal (Leon and Shadlen 2003) cortices; rat striatum (Matell et al. 2003)). This activity is similar to the climbing activity in that it builds up to a maximum that more-or-less coincides with a potential response. Moreover, Renoult et al. (2006) show that this activity can be scale invariant. The correlation between interval duration and neuronal activity that these studies have shown could indicate that these neurons are directly involved in the temporal control of behavior. However, this interpretation is complicated by the fact that monkeys, at least, can perform some interval timing tasks with parts of the prefrontal cortex lesioned (Manning 1973; Rosenkilde et al. 1981).

The stop-watch we have described is in perfect agreement with these experimental studies: the state of the stop-watch is monotonously increasing with time with a rate that depends on the interval duration (e.g. Fig. 1(c)). To account for the experimental findings we just need to postulate that the neurons with climbing activity are receiving inputs from the stop-watch, potentially located elsewhere. This gives a parsimonious explanation of the fact that similar type of climbing activity is found in different brain areas and that this activity might not be necessary for all interval timing tasks. Alternatively the climbing activity could be generated intrinsically by single cells (Durstewitz 2003) or by the local network (Reutimann et al. 2004). These alternative explanations are of course not mutually exclusive as more than one timing mechanism could be at play at the same time. We note however that it is not clear if these models of climbing activity support scale invariance over a wide range of time intervals.

6.3 Where in the brain could a stop-watch be implemented

We will now discuss some evidence indicating that the basal ganglia could be a place where the stop-watch is implemented. Although there are several different brain regions known to be involved in various aspects of timing (see Ivry 1996; Mauk and Buonomano 2004 for reviews) there is converging evidence indicating that the basal ganglia are crucially involved in ‘interval-timing’ tasks. In particular, using both pharmacological and lesions studies in animals and functional imaging in humans, Meck and co-workers have demonstrated a critical role of the basal ganglia in interval timing (see Meck 1996; Buhusi and Meck 2005; Meck et al. 2008 for review). In particular, an intact striatum seems to be necessary for rats to exhibit timing behavior (Meck 2006). That the basal ganglia are involved in interval timing fits well with that this is a set of brain regions that has been relatively well preserved during evolution (Smeets et al. 2000). Interestingly, the membrane potentials of cells in both ventral striatum (e.g. Wilson 1993) and dorsal (O’Donnell and Grace 1995) striatum are known to be bistable. The switching between the states seems to be random (Stern et al. 1997, Fig. 3), and the mean duration spent in each state is input dependent (Wilson and Kawaguchi 1996, Fig. 3). This indicates that striatum is one location where the stop-watch could potentially be implemented. Moreover, dopamine can modulate the excitability of the two activity states of striatal neurons (Hernandez-Lopez et al.

1997; West and Grace 2002) and presumably change the probability of switching between them (Nicola et al. 2000). Since dopamine is strongly implicated in the reward mechanisms of the nervous system (Schultz 1998) this dopamine dependence suggests a mechanism through which the activation rates could be modified in the stop-watch, and hence enable the timing of intervals of different duration. This argument is speculative but can be taken as a demonstration of that the neural mechanisms needed to implement the stop-watch seem to be available in a region heavily implicated in interval timing. The stop-watch can therefore be said to be biologically plausible. Whether something like the stop-watch is really used in interval timing must of course be investigated experimentally.

The basal ganglia is a key player in another model of interval timing: the striatal beat frequency model (Matell and Meck 2004). According to this model, the striatal cells act as coincidence detectors of oscillatory inputs from the frontal cortex. By making the cortical cells oscillate with slightly different frequencies long time intervals (much longer than the period of the oscillators) can be encoded by using the phase difference of different oscillators (also an idea originally suggested by Chris Miall (1989)). Through dopamine-induced plasticity, the duration of different intervals can be learned. This model is consistent with a large body of experimental findings. However, direct experimental evidence at the neuronal level, which is consistent with this model but not others, is lacking. It would also be important to verify that the striatal beat-frequency model can be implemented in a biophysically realistic model. The stop-watch and the striatal beat frequency model are similar in that they both propose a device (perhaps located in the striatum) that uses inputs coming from other brain regions to encode time intervals. However, the mechanisms through which this is done are very different. The stop-watch makes a strong prediction with respect to the activation times of the individual units: these must be random from trial to trial. The striatal beat frequency model, on the other hand, predicts that the between-trial variability in responding should be small. This could perhaps indicate an experimental way of discriminating between the two models.

In conclusion, we have suggested a simple model of a device that could be used to produce time-scale invariant behavior and shown how this model could be implemented in a range of neuronal models. We believe that the combination of being both simple and generic makes the suggested model an interesting possibility to consider for a plausible neuronal interval timer.

Acknowledgements We thank Matthew Matell for comments on a previous version of this manuscript. This work was supported by a grant BFU2007-61710 from the Spanish Government and by the grant FP6IST027198 from European Commission. AL acknowledges support from the Ramon y Cajal program. RA acknowledges support from the program Beatriu de Pinós from the Generalitat de Catalunya.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

Appendix

The abstract stop-watch

The abstract stop-watch consists of M units that can be in two states: spontaneous and activated. The units switch from the spontaneous to the activated state at random times (called activation times). In the discrete-time case the switching is governed by a fixed probability p of switching per unit time and the activation times will have a geometric distribution with parameter p . In the continuous-time case the activation times are exponentially distributed with rate parameter p . For small p the exponential distribution is an excellent approximation of the geometric distribution (e.g. Feller 1968). The primary use of the discrete-time version is to make the stop-watch more intuitive (Fig. 1). All analysis is done with the continuous-time version. In this case the number of units in the spontaneous state as a function of time is a stochastic process known as a linear death process (see for example Taylor and Karlin 1998). Even if this is a well understood stochastic process the results we are using do not seem to be available in standard sources and we therefore give the derivations here.

The distribution of the number of units in a certain state

Since the time point each unit switches to the activated state is an exponential random variable (denoted U_i) it follows that the probability that it switches in the interval from 0 to t is given by $1 - e^{-pt}$. Given that the units are identical and independent the probability of having x units in the activated state at time t is given by the following binomial distribution

$$Pr\{X(t) = x\} \stackrel{\text{def}}{=} a_x(t) = \binom{M}{x} (1 - e^{-pt})^x (e^{-pt})^{(M-x)}. \tag{11}$$

This follows from considering the state of each unit the outcome of a Bernoulli trial with probability $1 - e^{-pt}$ of “success” and e^{-pt} of “failure”. For an alternative derivation see Taylor and Karlin (1998).

The distribution of $T_{f \cdot M}$

Let $T_{f \cdot M}$ be the time the process enters the state of $f \cdot M$ units activated. That is

$$T_{f \cdot M} = \min_t \{X(t) = f \cdot M\}.$$

We will now derive the probability density function of $T_{f \cdot M}$. First note that the number of units in the activated state is a monotonously increasing function of time. Furthermore, the probability that two units will enter the activated state at exactly the same time is zero. This implies that to enter the state characterized by having $f \cdot M$ units activated is only possible through the state $f \cdot M - 1$. Next we will approximate the probability of entering state $f \cdot M$ in the interval from t to $t + \Delta t$ where Δt is a time interval small enough that we can ignore the event of having more than one transition during the interval. Then it follows that the state at time t must be $f \cdot M - 1 \stackrel{\text{def}}{=} \phi$ and this has probability $a_\phi(t)$ of occurring (from Eq. (11)). The probability of having exactly one of the remaining units switching to the activated state during the next Δt is given by

$$c(\Delta t) = (M - \phi)(1 - e^{-p\Delta t})e^{-p\Delta t(M-f \cdot M)}.$$

Note that c does not depend on t . Taylor expanding this expression in Δt gives

$$c = (M - \phi)p\Delta t + \text{h.o.t.},$$

where h.o.t. denotes terms of order $(\Delta t)^2$ or higher. Then we can approximate the probability that the system enters state $f \cdot M$ before time T by

$$Pr\{T_{f \cdot M} \leq T\} \approx \sum_{i=1}^K a_\phi(t_i)[p(M - \phi)\Delta t + \text{h.o.t.}],$$

where K is the number of intervals of length Δt that fits in $[0, T]$ and t_i is the mid-point of interval i . In the limit of $K \rightarrow \infty$ (and $\Delta t \rightarrow 0$) we get

$$Pr\{T_{f \cdot M} \leq T\} = p(M - \phi) \int_0^T a_\phi(t) dt. \tag{12}$$

This is an equality since the higher order terms disappear in the limit and the approximations made to start with become exact. Thus, Eq. (12) implies that the probability density function of $T_{f,M}$ is given by

$$g_{f,M}(t) = p(M - \phi) \binom{M}{\phi} (1 - e^{-pt})^\phi (e^{-pt})^{(M-\phi)},$$

which is identical to Eq. (4) above.

First two moments of $T_{f,M}$

The mean and the variance of $T_{f,M}$ can in principle be obtained from integrating the density function Eq. (4). This, however is tedious and we can proceed through a more intuitive route. Let Y_k be the interval between the occurrence of the k th and the $(k + 1)$ th switch to the activated state. This interval will have the distribution of the minimum of $M - k$ exponential random variables all having the same rate parameter. This distribution is in fact also an exponential but with rate parameter $p(M - k)$. (To see this, consider the joint probability of that no unit has become activated at time t and use that the activation times are exponentially distributed and independent.) Moreover, for different k s the variables Y_k are independent. The expected value of $T_{f,M}$ is just the sum of the expected values of Y_k for $k = 0, 1, \dots, f \cdot M - 1$, or in other words

$$E(T_{f,M}) = \frac{1}{p} \sum_{k=0}^{f \cdot M - 1} \frac{1}{M - k}.$$

The same reasoning gives the expression for the variance of $T_{f,M}$:

$$V(T_{f,M}) = \frac{1}{p^2} \sum_{k=0}^{f \cdot M - 1} \frac{1}{(M - k)^2},$$

from which we get the standard deviation by taking the square root.

Optimal value of f

In this section we want to determine the fraction of units f in the activated state so that the CV is as small as possible. The CV is given by

$$CV = \frac{\sqrt{\frac{1}{M^2} \sum_{k=0}^{M \cdot f - 1} \frac{1}{(1 - k/M)^2}}}{\frac{1}{M} \sum_{k=0}^{M \cdot f - 1} \frac{1}{1 - k/M}}.$$

If we introduce a new variable $x = k/M$ we can approximate the above sums by integrals. This gives

$$CV \approx \frac{\sqrt{\frac{1}{M} \int_0^y \frac{dx}{(1-x)^2}}}{\int_0^y \frac{dx}{1-x}} \stackrel{\text{def}}{=} c(y),$$

where y is a variable that ranges from 0 to 1. This expression shows that when M is large (so that the approximation is good) the CV scales as $\sqrt{1/M}$.

Next, we want to find the minimum of this function (i.e. $c(y)$) in the interval $[0, 1)$. To do this we take the derivative with respect to y and set the resulting expression equal to zero. This gives

$$\frac{y}{\log(1-y)} + 0.5 = 0.$$

Solving this equation numerically we get that $y = 0.7968$. To approximate the sums with integrals, we need that M is reasonably large. However, we have verified numerically that the minima of the CV is attained close to 0.8 for small M as well.

The neuronal stop-watch

In the neuronal version of the stop-watch we assume that each unit is described by a bistable system where the spontaneous state corresponds to an equilibrium point. We further assume that the stability of this state can be controlled by changing the input to the system and moreover that there is a level of input beyond which this state is no longer stable. A generic model of this situation is then given by Eq. (7). Indeed, by changing the variable to $y = \beta x$ in this equation we get

$$dy = dt(\beta\mu + y^2) + \sqrt{dt}\beta\sigma\xi(t)$$

which is the so-called normal-form of a saddle-node bifurcation driven by white noise (for an introduction to bifurcation theory and normal-forms see Strogatz (1994)). This model is generic in the sense that it approximates *any* system of the above type where the change of stability is associated with one zero eigenvalue (see Guckenheimer and Holmes 1983, Section 3.4 for a more exact statement). What this means in non-technical terms is that given a system with the above characteristics, Eq. (7) can be used to describe the dynamics of this system close to the point where the spontaneous state is no longer stable. Since this is exactly the regime we are interested in, Eq. (7) is indeed the correct model to use. In Fig. 8 we showed that the Morris-Lecar model is very well approximated by this equation.

Numerical simulations

To integrate Eq. (7), as well as the Morris-Lecar and micro-circuit models with noisy inputs we use Heun’s method (Greiner et al. 1988), with step size 0.02 for Eq. (7) and 0.002 for the Morris-Lecar and micro-circuit models.

In all numerical solutions of Eq. (7) we start the solution at the stable equilibrium and use a threshold $x > 2$ to detect when the system has escaped from the spontaneous state. Using another (larger) threshold value would not affect the results since once the system has left the spontaneous state it will very rapidly diverge to $+\infty$.

The first passage times

The first passage time is a random variable that describes the time when the state-variable in Eq. (7) reaches a given level for the first time. There are no closed-form expressions in terms of elementary functions for the passage time distribution nor for its expected value. If we let $U(x) = -\mu x - (\beta/3)x^3$ be the “energy function” then the mean first passage time as a function of initial condition x_0 is given by (Pontryagin et al. 1933/1989, Eq. (24)):

$$m(x_0) = \frac{2}{\sigma^2} \int_{x_0}^b \exp\left\{\frac{-2U(y)}{\sigma^2}\right\} \int_{-\infty}^y \exp\left\{\frac{2U(z)}{\sigma^2}\right\} dz dy. \tag{13}$$

(See Karlin and Taylor (1981, Section 15.3) for a more recent account). In this expression b denotes the threshold level. We solved this equation numerically with $x_0 = -\sqrt{|\mu|}/\beta$ and $b = 2$ using Romberg integration (Press et al. 1992). We used the inverse of the mean first passage time as the activation probability for the neuronal stop-watch.

The approximation given by Eq. (8) was first derived in Kramers (1940), a recent and accessible account is given in Van Kampen (2007, Section XIII.2).

To get the distribution of passage times we solved numerically the Fokker-Planck equation (also called Kolmogorov forward equation) associated with Eq. (7). This equation reads:

$$\frac{\partial P(x, t)}{\partial t} = -\frac{\partial}{\partial x}([\mu + \beta x^2]P(x, t)) + \frac{\sigma^2}{2} \frac{\partial^2 P(x, t)}{\partial x^2}. \tag{14}$$

We used the so-called backward time centered space discretization method to solve this equation and used a step-size of 0.02 for both time and space variables. The

first passage density is given by the so-called probability flux

$$J(t, x) = (\mu + \beta x^2)P(t, x) - \frac{\sigma^2}{2} \frac{\partial P(t, x)}{\partial x},$$

at the right boundary (Van Kampen 2007). To solve for this numerically we evaluated the flux at a point close to but not quite at the boundary since the discontinuity at the boundary often leads to small oscillations in the solution there.

An intuitive argument for why the first passage time distributions will be exponential as $|\mu|$ increases is given next. As $|\mu|$ increases (or σ decreases) the deterministic part of Eq. (7) will have increasingly stronger influence over the dynamics. This means that if the system is initiated somewhere to the right of the unstable equilibrium it will after a short amount of time be in a quasi-equilibrium state. That is, P in Eq. (14) which will be a narrowly peaked function centered at the stable equilibrium point. It is not a true equilibrium distribution because eventually all the probability will escape from the domain. For large values of $|\mu|$ this quasi-equilibrium will be established much faster than the mean escape time. This implies that the probability of escaping becomes independent of initial condition. Since the process is Markovian by assumption, this independence of initial condition implies that the escape times are also Markovian (since the system is uniquely characterized by the state variable). If the escape times are Markovian the distribution must be exponential.

The Morris-Lecar model

We follow the formulation of the Morris-Lecar model given in Izhikevich (2007). In this work it is referred to as the ‘persistent sodium plus potassium model’ with a high-threshold I_K . The parameters are the same as those stated in Izhikevich’s book (in the legend to Fig. 4.1) but we state them here as well for easy reference: $E_L = -80$, $E_{Na} = 50$, $E_K = -90$, $g_L = 8$, $g_{Na} = 20$, $g_K = 10$, $\tau_n = 0.165$, $m_{1/2} = -20$, $n_{1/2} = -25$, $m_k = 15$, $n_k = 5$. Here the E s are equilibrium potentials, the g s conductances and τ_n a time constant. The noise amplitude of the input current was $\sigma^2 = 0.06^2$.

For these values of the parameters the system (with $\sigma = 0$) undergoes a saddle-node bifurcation for an input current of $I_b = 4.5129$ for which $V = V_b = -60.933$. In the vicinity of this point the system can be approximated by the following one-dimensional equation

$$\frac{dV}{dt} = 1.0073(I - I_b) + 0.1901(V - V_b)^2.$$

If we add white noise with variance σ^2 to the input current in the Morris-Lecar model we get the following reduced model

$$\frac{dV}{dt} = 1.0073(I - I_b) + 0.1901(V - V_b)^2 + 1.0073\sigma\xi(t)$$

which is formally equivalent to Eq. (7) in Section 4. The reduction of the Morris-Lecar model to the form of Eq. (7) was made through a so-called center manifold reduction for a saddle node bifurcation, a standard procedure (e.g. Kuznetsov 2004). In fact, a simplified reduction of this very model is performed in Izhikevich (2007, Section 6.1.1) to which the reader is referred for more details. Note however that the exact correspondence between the models can only be obtained through the center manifold reduction (or something equivalent) of the full system.

References

- Amit, D. J., & Brunel, N. (1997). Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cerebral Cortex*, 7, 237–252.
- Brody, C. D., Hernandez, A., Zainos, A., & Romo, R. (2003). Timing and neural encoding of somatosensory parametric working memory in macaque prefrontal cortex. *Cerebral Cortex*, 13, 1196–1207.
- Buhusi, C. V., & Meck, W. H. (2005). What makes us tick? Functional and neural mechanisms of interval timing. *Nature Reviews. Neuroscience*, 6, 755–765.
- Carr, J. (1981). *Applications of centre manifold theory*. New York: Springer.
- Chance, F. S., Abbott, L. F., & Reyes, A. D. (2002). Gain modulation from background synaptic input. *Neuron*, 35, 773–872.
- Chiba, A., Oshio, K. I., & Inase, M. (2008). Striatal neurons encoded temporal information in duration discrimination task. *Experimental Brain Research*, 186, 671–676.
- Durstewitz, D. (2003). Self-organizing neural integrator predicts interval times through climbing activity. *Journal of Neuroscience*, 23, 5342–5353.
- Egorov, A. V., Hamam, B. N., Fransén, E., Hasselmo, M. E., & Alonso, A. A. (2002). Graded persistent activity in entorhinal cortex neurons. *Nature*, 420, 173–178.
- Feller, W. (1968). *An introduction to probability theory and its applications* (3rd ed.). New York: Wiley.
- Ferster, C. B., & Skinner, B. F. (1957). *Schedules of reinforcement*. Acton: Copley.
- Fetz, E. E. (1969). Operant conditioning of cortical unit activity. *Science*, 163, 955–957.
- Fiala, J. C., Grossberg, S., & Bullock, D. (1996). Metabotropic glutamate receptor activation in cerebellar purkinje cells as substrate for adaptive timing of the classically conditioned eye-blink response. *Journal of Neuroscience*, 16, 3760–3774.
- Genet, S., & Delord, B. (2002). A biophysical model of nonlinear dynamics underlying plateau potentials and calcium spikes in purkinje cell dendrites. *Journal of Neurophysiology*, 88, 2430–2444.
- Gibbon, J. (1977). Scalar expectancy-theory and webers law in animal timing. *Psychological Review*, 84, 279–325.
- Gibbon, J., Church, R. M., & Meck, W. H. (1984). Scalar timing in memory. *Annals of the New York Academy of Sciences*, 423, 52–77.
- Gibbon, J., Malapani, C., Dale, C. L., & Gallistel, C. R. (1997). Toward a neurobiology of temporal cognition: Advances and challenges. *Current Opinion in Neurobiology*, 7, 170–184.
- Greiner, A., Strittmatter, W., & Honerkamp, J. (1988). Numerical-integration of stochastic differential-equations. *Journal of Statistical Physics*, 51, 95–108.
- Grossberg, S., & Schmajuk, N. A. (1989). Neural dynamics of adaptive timing and temporal discrimination during associative learning. *Neural Networks*, 2, 79–102.
- Guckenheimer, J., & Holmes, P. (1983). *Nonlinear oscillations, dynamical systems, and bifurcation of vector fields*. New York: Springer.
- Hansel, D., & Mato, G. (2003). Asynchronous states and the emergence of synchrony in large networks of interacting excitatory and inhibitory neurons. *Neural Computation*, 15, 1–56.
- Hernandez-Lopez, S., Bargas, J., Surmeier, D. J., Reyes, A., & Galarraga, E. (1997). D-1 receptor activation enhances evoked discharge in neostriatal medium spiny neurons by modulating an l-type Ca²⁺ conductance. *Journal of Neuroscience*, 17, 3334–3342.
- Ivry, R. B. (1996). The representation of temporal information in perception and motor control. *Current Opinion in Neurobiology*, 6, 851–857.
- Izhikevich, E. M. (2007). *Dynamical systems in neuroscience*. Cambridge: MIT.
- Kalenscher, T., Ohmann, T., Windmann, S., Freund, N., & Gunturkun, O. (2006). Single forebrain neurons represent interval timing and reward amount during response scheduling. *European Journal of Neuroscience*, 24, 2923–2931.
- Karlin, S., & Taylor, H. M. (1981). *A second course in stochastic processes* (1st ed.). New York: Academic.
- Killeen, P. R., & Fetterman, J. G. (1988). A behavioral-theory of timing. *Psychological Review*, 95, 274–295.
- Kitano, K., Okamoto, H., & Fukai, T. (2003). Time representing cortical activities: Two models inspired by prefrontal persistent activity. *Biological Cybernetics*, 88, 387–394.
- Kojima, S., & Goldman-Rakic, P. S. (1982). Delay-related activity of prefrontal neurons in rhesus-monkeys performing delayed-response. *Brain Research*, 248, 43–49.
- Kojima, S., Matsumura, M., & Kubota, K. (1981). Prefrontal neuron activity during delayed-response performance without imperative go signals in the monkey. *Experimental Neurology*, 74, 396–407.
- 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, A. A., Raghavachari, S., Kepecs, A., & Lisman, J. E. (2002). Model for a robust neural integrator. *Nature Neuroscience*, 5, 775–782.
- Kramers, H. A. (1940). Brownian motion in a field of force and the diffusion model of chemical reactions. *Physica*, 7, 284–304.
- Kuznetsov, Y. A. (2004). *Elements of applied bifurcation theory* (3rd ed.). New York: Springer.
- Lebedev, M. A., O’Doherty, J. E., & Nicolelis, M. A. L. (2008). Decoding of temporal intervals from cortical ensemble activity. *Journal of Neurophysiology*, 99, 166–186.
- Lejeune, H., & Wearden, J. H. (2006). Scalar properties in animal timing: Conformity and violations. *Quarterly Journal of Experimental Psychology*, 59, 1875–1908.

- Leon, M. I., & Shadlen, M. N. (2003). Representation of time by neurons in the posterior parietal cortex of the macaque. *Neuron*, *38*, 317–327.
- Machado, A. (1997). Learning the temporal dynamics of behavior. *Psychological Review*, *104*, 241–265.
- Manning, F. J. (1973). Performance under temporal schedules by monkeys with partial ablations of prefrontal cortex. *Physiology & Behavior*, *11*, 563–569.
- Matell, M. S., & Meck, W. H. (2004). Cortico-striatal circuits and interval timing: Coincidence detection of oscillatory processes. *Cognitive Brain Research*, *21*, 139–170.
- Matell, M. S., Meck, W. H., & Nicolelis, M. A. L. (2003). Interval timing and the encoding of signal duration by ensembles of cortical and striatal neurons. *Behavioral Neuroscience*, *117*, 760–773.
- Mauk, M. D., & Buonomano, D. V. (2004). The neural basis of temporal processing. *Annual Review of Neuroscience*, *27*, 307–340.
- Meck, W. H. (1996). Neuropharmacology of timing and time perception. *Cognitive Brain Research*, *3*, 227–242.
- Meck, W. H. (2006). Neuroanatomical localization of an internal clock: A functional link between mesolimbic, nigrostriatal, and mesocortical dopaminergic systems. *Brain Research*, *1109*, 93–107.
- Meck, W. H., Penney, T. B., & Pouthas, V. (2008). Cortico-striatal representation of time in animals and humans. *Current Opinion in Neurobiology*, *18*, 145–152.
- Miall, R. C. (1989). The storage of time intervals using oscillating neurons. *Neural Computation*, *1*, 359–371.
- Miall, R. C. (1993). Neural networks and the representation of time. *Psychologica Belgica*, *33*, 255–269.
- Miller, P., & Wang, X. J. (2006). Stability of discrete memory states to stochastic fluctuations in neuronal systems. *Chaos*, *16*, 026,109.
- Morris, C., & Lecar, H. (1981). Voltage oscillations in the barnacle giant muscle-fiber. *Biophysics Journal*, *35*, 193–213.
- Nicola, S. M., Surmeier, D. T., & Malenka, R. C. (2000). Dopaminergic modulation of neuronal excitability in the striatum and nucleus accumbens. *Annual Review of Neuroscience*, *23*, 185–215.
- Niki, H., & Watanabe, M. (1979). Prefrontal and cingulate unit-activity during timing behavior in the monkey. *Brain Research*, *171*, 213–224.
- O'Donnell, P., & Grace, A. A. (1995). Synaptic-interactions among excitatory afferents to nucleus-accumbens neurons—Hippocampal gating of prefrontal cortical input. *Journal of Neuroscience*, *15*, 3622–3639.
- Okamoto, H., & Fukai, T. (2001). Neural mechanism for a cognitive timer. *Physical Review Letters*, *86*, 3919–3922.
- Okamoto, H., Isomura, Y., Takada, M., & Fukai, T. (2007). Temporal integration by stochastic recurrent network dynamics with bimodal neurons. *Journal of Neurophysiology*, *97*, 3859–3867.
- Oshio, K. I., Chiba, A., & Inase, M. (2006). Delay period activity of monkey prefrontal neurones during duration-discrimination task. *European Journal of Neuroscience*, *23*, 2779–2790.
- Pontryagin, L., Andronov, A., & Vitt, A. (1933/1989). Appendix. On the statistical treatment of dynamical systems. In F. Moss, & P. V. E. McClintock (Eds.), *Noise in nonlinear dynamical systems* (pp. 329–348). Cambridge: Cambridge University Press.
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. (1992). *Numerical recipes in C*. Cambridge: Cambridge University Press.
- Renart, A., Brunel, N., & Wang, X. J. (2003). Mean-field theory of irregularly spiking neuronal populations and working memory in recurrent cortical networks. In J. Feng (Ed.), *Computational neuroscience: A comprehensive approach* (pp. 431–490). Boca Raton: Chapman&Hall/CRC.
- Renoult, L., Roux, S., & Riehle, A. (2006). Time is a rubber-band: Neuronal activity in monkey motor cortex in relation to time estimation. *European Journal of Neuroscience*, *23*, 3098–3108.
- Reutimann, J., Yakovlev, V., Fusi, S., & Senn, W. (2004). Climbing neuronal activity as an event-based cortical representation of time. *Journal of Neuroscience*, *24*, 3295–3303.
- Rinzel, J., & Ermentrout, B. (1998). Analysis of neural excitability and oscillations. In C. Koch, & I. Segev (Eds.), *Methods in neuronal modeling* (pp. 251–291). Cambridge: MIT.
- Rosenkilde, C. E., Rosvold, H. E., & Mishkin, M. (1981). Time discrimination with positional responses after selective prefrontal lesions in monkeys. *Brain Research*, *210*, 129–144.
- Roux, S., Coulmance, M., & Riehle, A. (2003). Context-related representation of timing processes in monkey motor cortex. *European Journal of Neuroscience*, *18*, 1011–1016.
- Sakurai, Y., Takahashi, S., & Inoue, M. (2004). Stimulus duration in working memory is represented by neuronal activity in the monkey prefrontal cortex. *European Journal of Neuroscience*, *20*, 1069–1080.
- Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, *80*, 1–27.
- Smeets, W. J. A. J., Marin, O., & Gonzalez, A. (2000). Evolution of the basal ganglia: New perspectives through a comparative approach. *Journal of Anatomy*, *196*, 501–517.
- Staddon, J. E. R., & Cerutti, D. T. (2003). Operant conditioning. *Annual Review of Psychology*, *54*, 115–144.
- Staddon, J. E. R., & Higa, J. J. (1999). Time and memory: Towards a pacemaker-free theory of interval timing. *Journal of the Experimental Analysis of Behavior*, *71*, 215–251.
- Stern, E. A., Kincaid, A. E., & Wilson, C. J. (1997). Spontaneous subthreshold membrane potential fluctuations and action potential variability of rat corticostriatal and striatal neurons *in vivo*. *Journal of Neurophysiology*, *77*, 1697–1715.
- Strogatz, S. H. (1994). *Nonlinear dynamics and chaos*. Boulder: Westview.
- Tahvildari, B., Fransén, E., Alonso, A. A., & Hasselmo, M. E. (2007). Switching between “on” and “off” states of persistent activity in lateral entorhinal layer III neurons. *Hippocampus*, *17*, 257–263.
- Taylor, H. M., & Karlin, S. (1998). *An introduction to stochastic modeling* (3rd ed.). London: Academic.
- Van Kampen, N. G. (2007). *Stochastic processes in physics and chemistry* (3rd ed.). Amsterdam: Elsevier.
- Wearden, J. H., & Lejeune, H. (2008). Scalar properties in human timing: Conformity and violations. *Quarterly Journal of Experimental Psychology*, *61*, 569–587.
- West, A. R., & Grace, A. A. (2002). Opposite influences of endogenous dopamine d-1 and d-2 receptor activation on activity states and electrophysiological properties of striatal neurons: Studies combining *in vivo* intracellular recordings and reverse microdialysis. *Journal of Neuroscience*, *22*, 294–304.
- Wilson, C. J. (1993). The generation of natural firing patterns in neostriatal neurons. *Chemical Signalling In The Basal Ganglia*, *99*, 277–297.
- Wilson, C. J., & Kawaguchi, Y. (1996). The origins of two-state spontaneous membrane potential fluctuations of neostriatal spiny neurons. *Journal of Neuroscience*, *16*, 2397–2410.