Behavioural Processes 95 (2013) 90–99

Contents lists available at SciVerse ScienceDirect

Behavioural Processes

journal homepage: www.elsevier.com/locate/behavproc

# An adaptive drift-diffusion model of interval timing dynamics

Andre Luzardo[a], Elliot A. Ludvig[b], François Rivest[c,*]

[a] Champalimaud Neuroscience Programme, Portugal
[b] Princeton Neuroscience Institute and Department of Mechanical & Aerospace Engineering, Princeton University, USA
[c] Department of Mathematics and Computer Science, Royal Military College of Canada, Canada

## ARTICLE INFO

*Article history:*
Received 19 November 2012
Received in revised form 1 February 2013
Accepted 6 February 2013

*Keywords:*
Interval timing
Drift-diffusion processes
Cyclic schedules
Learning
Computational models
Pigeons

## ABSTRACT

Animals readily learn the timing between salient events. They can even adapt their timed responding to rapidly changing intervals, sometimes as quickly as a single trial. Recently, drift-diffusion models—widely used to model response times in decision making—have been extended with new learning rules that allow them to accommodate steady-state interval timing, including scalar timing and timescale invariance. These time-adaptive drift-diffusion models (TDDMs) work by accumulating evidence of elapsing time through their drift rate, thereby encoding the to-be-timed interval. One outstanding challenge for these models lies in the dynamics of interval timing—when the to-be-timed intervals are non-stationary. On these schedules, animals often fail to exhibit strict timescale invariance, as expected by the TDDMs and most other timing models. Here, we introduce a simple extension to these TDDMs, where the response threshold is a linear function of the observed event rate. This new model compares favorably against the basic TDDMs and the multiple-time-scale (MTS) habituation model when evaluated against three published datasets on timing dynamics in pigeons. Our results suggest that the threshold for triggering responding in interval timing changes as a function of recent intervals.

This article is part of a Special Issue entitled: SQAB 2012.

## 1. Introduction

Humans and other animals are exquisitely sensitive to the timing of rewards in a wide range of situations (for reviews, see Buhusi and Meck, 2005; Grondin, 2010). Most timing theories attempt to explain the steady-state behavior that emerges after extensive experience with the same interval duration (e.g., Gibbon, 1977; Killeen and Fetterman, 1988; Machado, 1997; Matell and Meck, 2004), as in the peak procedure (Catania, 1970; Roberts, 1981). Under these static conditions, timed responding usually exhibits timescale invariance: response-time distributions are strictly proportional to the timed interval. The dynamics of timing, however, including the initial acquisition period and changes in contingencies, have tended to receive less attention (but see Staddon and Higa, 1999; Balsam et al., 2002; Ludvig et al., 2012). In this paper, we introduce a new extension of time-adaptive drift-diffusion models (TDDMs), which deals with both the steady-state and dynamic features of interval timing.

In nature, rewards do not always appear with a fixed regularity. An environment where important time intervals are not stationary might select for animals that can adapt their time judgments rapidly. Studies using schedules with systematically varying intervals have largely supported this conjecture. For example, on cyclic-interval schedules, successive intervals vary according to a periodic function. Animals trained under these schedules adapt their post-reinforcement pause rapidly, as if tracking the preceding interval. Fig. 1 shows three example schedules where animals display this tracking: arithmetic ascending and descending (Innis and Staddon, 1971), sinusoidal (Higa et al., 1991) and square-wave sequences of intervals (Ludvig and Staddon, 2004). This tracking often occurs with a lag of only a single trial and sometimes even anticipates the next interval (e.g., Church and Lacourse, 1998; Ludvig and Staddon, 2005). This rapid temporal tracking presents a significant challenge for static timing models.

A second challenging feature of these cyclic-interval schedules is that the data often fail to exhibit the timescale invariance typical of interval timing on static schedules (e.g., Gallistel and Gibbon, 2000) though not universally (see Zeiler and Powell, 1994; Ludvig et al., 2008; Kehoe et al., 2009). With static schedules, for example, doubling the to-be-timed interval will usually result in timed responses that are both twice as long and twice as variable (e.g., Schneider, 1969; Gibbon, 1977). On dynamic schedules, in contrast, longer intervals often elicit proportionally shorter pauses (e.g., Innis and Staddon, 1971; Higa et al., 1991; Ludvig and Staddon, 2004,

---

* Corresponding author at: Department of Mathematics and Computer Science, Royal Military College of Canada, CP 17000, Succursale Forces, Kingston, ON, Canada K7K 7B4. Tel.: +1 613 541 6000x6232; fax: +1 613 541 6584.
*E-mail addresses:* francois.rivest@mail.mcgill.ca, francois.rivest@rmc.ca (F. Rivest).

0376-6357/$ – see front matter. Crown Copyright © 2013 Published by Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.beproc.2013.02.003

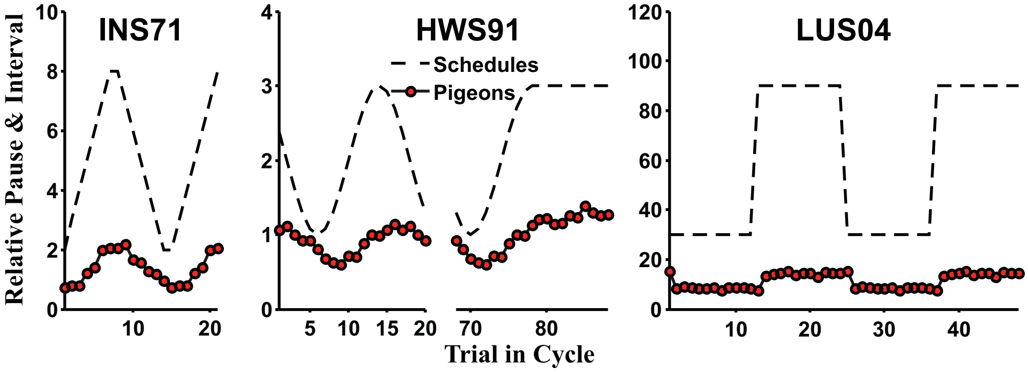

**Fig. 1.** Interval cycles and sample data from the three experiments modeled in this paper: INS71 (Innis and Staddon, 1971), HWS91 (Higa et al., 1991), and LUS04 (Ludvig and Staddon, 2004).

2005). The right panel of Fig. 1 illustrates how, on cyclic schedules, when the interval varies across a threefold range (from 30 to 90 s) the post-reinforcement pause less than doubles. This dependence of response timing on the absolute timescale poses a second challenge for timing theories.

One newly developed timing theory seems particularly well-suited to deal with these timing dynamics because of its ability to learn rapidly: the time-adaptive drift-diffusion models (TDDMs; Rivest and Bengio, 2011; Simen et al., 2011). Basic drift-diffusion models (DDMs) are the most widely used process model of real-time decision making (Ratcliff, 1978; Ratcliff and Rouder, 2000). In a DDM, decisions are made by accumulating noisy evidence (modeled as a random walk with drift) until a threshold is reached, eliciting a response. One strength of this class of models is that they assume variability within trials in several components of processing, making it possible to study those components in isolation and to make real-time predictions. These features make DDMs particularly well suited to model both measures of accuracy and response-time distributions (Ratcliff, 2002; Ratcliff and McKoon, 2008; Leite and Ratcliff, 2010). Moreover, strikingly, noisy ramping neural activity that matches the noisy evidence accumulation process in a DDM can be seen in certain cortical areas during decision-making tasks (Bogacz, 2007; Gold and Shadlen, 2007; Rushworth et al., 2011). As noted by Leon and Shadlen (2003), this activity shares some properties with interval timing and could be a potential neural substrate for timing in monkeys.

Recently, two similar formal extensions of this DDM framework to interval timing were introduced: the adaptive drift-diffusion process (ADDP, Rivest and Bengio, 2011) and the stochastic ramp and trigger (SRT, Simen et al., 2011). Though independently developed, the ADDP and the Level-4 SRT models are almost identical, and we refer to them by the joint name of TDDMs (time-adaptive DDMs). Both models extend the DDM by incorporating a learning rule that allows for rapid adaptation of the drift rate to the experienced time interval. These TDDMs can even adapt to a change in experienced time intervals as rapidly as a single trial. One potential limitation of these TDDMs, however, is that timescale invariance is built right into the algorithm. As a result, whether these models can account for dynamic timing data, which often lack such invariance, is unclear.

In the TDDMs, the memory of recent intervals is stored in the drift rate, which controls the slope of the signal over time. A change in the experienced intervals leads to a change in the drift rate (i.e., the accumulator slope). Interestingly, a similar change in the slope of the firing rate of neurons has been observed during timing tasks in the posterior thalamic region in rats (Komura et al., 2001) and multiple cortical areas in monkeys (Leon and Shadlen, 2003; Reutimann et al., 2004; Lebedev et al., 2008). The TDDMs have also been shown to be a reasonable approximation of more complex networks of spiking neurons (Simen et al., 2011).

In the timing literature, the only major model that explicitly contends with interval timing dynamics is the multiple-time-scale (MTS) habituation model (Staddon and Higa, 1999; Staddon et al., 2002). The MTS model uses a decaying memory trace to keep track of time. This memory trace is made up of a sum of decaying units that increases in activity when a reward is presented. A response threshold is updated at every trial to correspond to the level of decay in the memory trace at the moment the reward was last seen. This dynamic threshold enables MTS to reproduce the tracking behavior observed in cyclic schedules such as those shown in Fig. 1.

In this paper, we systematically compare the performance of the TDDMs and MTS on three different sets of dynamic timing data in pigeons from the published literature (Innis and Staddon, 1971; Higa et al., 1991; Ludvig and Staddon, 2004). These datasets use different cyclic-interval schedules and are fairly representative of the literature on the dynamics of interval timing. Because of the dependence on absolute timescales in the timing measures, MTS is able to better match the pigeon data than the basic TDDM models. Motivated by this shortcoming of the basic TDDM, we introduce a modified TDDM with a threshold that depends linearly on the slope of the drift rate (the LT-TDDM). This modified TDDM matches the timing data from pigeons as well as or better than MTS in most experiments.

## 2. Methods

### 2.1. Model specifications

#### 2.1.1. Time-adapting drift-diffusion model (TDDM)

A drift-diffusion process consists of a noisy signal randomly drifting over time but tending toward a particular direction as dictated by its drift rate. In the TDDM implementation, this signal $\varphi(t)$ starts at 0 at stimulus onset, and thus serves as the internal representation of elapsed time $t$ from stimulus onset. This drifting process is analogous to an accumulator that continuously integrates time at a drift rate $w$ with noise $\varepsilon(t)$:

$$\varphi(t) = \varphi(t-1) + w\Delta t + \varepsilon(t) \tag{1}$$

where $\Delta t$ is the time step used to simulate the continuous process and $\varepsilon(t)$ is Gaussian noise with mean 0 and variance $\sigma^2$ [$N(0,\sigma^2)$]. On reward, the signal $\varphi(t)$ is reset to 0. The signal has an upper absorbing boundary at 1, meaning that once $\varphi(t)$ reaches 1, it stays there until reset by a reward. To achieve Weber's law, whereby the variance in responding is proportional to the mean (Gibbon, 1977), the noise variance $\sigma^2$ must equal $\beta^2 w\Delta t$ (Rivest and Bengio,