City Research Online

City, University of London Institutional Repository

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

A Drift-Diffusion Model of Interval Timing in the Peak Procedure

André Luzardo[a,*], François Rivest[b], Eduardo Alonso[a], Elliot A. Ludvig[c]

[a]*Department of Computer Science, City University London, London, United Kingdom*
[b]*Mathematics and Computer Science, Center for Neuroscience Studies, Royal Military College of Canada - Queen's University, Kingston, Ontario, Canada*
[c]*Department of Psychology, University of Warwick, Coventry, United Kingdom*

Abstract

Drift-diffusion models (DDMs) are a popular framework for explaining response times in decision-making tasks. Recently, the DDM architecture has been used to model interval timing. The Time-adaptive DDM (TDDM) is a physiologically plausible mechanism that adapts in real-time to different time intervals while preserving timescale invariance. One key open question is how the TDDM could deal with situations where reward is omitted, as in the peak procedure—a benchmark in the timing literature. When reward is omitted, there is a consistent pattern of correlations between the times at which animals start and stop responding. Here we develop a formulation of the TDDM's stationary properties that allows for the derivation of such correlations analytically. Using this simplified formulation we show that a TDDM with two thresholds–one to mark the start of responding and another the stop–can reproduce the same pattern of correlations observed in the data, as long as the start threshold is allowed to be noisy. We confirm this by running simulations with the standard TDDM formulation and show that the simplified formulation approximates well the full model under steady-state conditions. Moreover, we show that this simplified version of the TDDM is formally equivalent to Scalar Expectancy Theory (SET) under stationary behaviours, the most prominent theory of interval timing. This equivalence establishes the TDDM as a more complete drift-diffusion based theory with SET as a special case under steady-state conditions.

Keywords: interval timing, peak procedure, computational models, drift-diffusion model, scalar expectancy theory

1. Introduction

Learning the time between important events is a fundamental feature of cognition. Humans and other animals can readily learn the timing of upcoming

---

*Corresponding author
Email address: `Andre.Luzardo.1@city.ac.uk` (André Luzardo)

rewards and adapt their behavior accordingly [1, 2]. A range of psychological
and computational theories have been proposed for interval timing [3, 4, 5, 6,
7, 8, 9] which succeed at capturing the broad outlines of timing behaviour, but
they often flounder when dealing with the statistics of the micro-structure of
real-time responding.

Particularly vexing for timing models are the behavioural patterns when
predictably-timed rewards are occasionally omitted, as in the peak procedure
[10]. This peak procedure is likely the most popular interval timing task. Al-
though major timing models such as Scalar Expectancy Theory (SET)[3], Be-
havioral Theory of Timing (BeT)[4], Learning to Time (LeT)[11, 8] and Multiple
Time Scales (MTS)[5] can reproduce the global averaged behaviour in this task,
very few models have been able to account for the pattern of behaviour observed
in individual trials. The notable exception is SET, which provides good quan-
titative fits to animal data [12] and remains the theory of choice for explaining
static timing phenomena.

Recently, a series of studies have adapted the classic drift-diffusion model
(DDM) used to explain the dynamics of real-time decision-making in behaviour
and the brain [13, 14, 15] to interval timing [9, 16, 17, 18]. The Time-adaptive
DDM (TDDM) explains timing as the result of a noisy drift-diffusion process
with an adaptive drift rate, which is adjusted based on the time interval ob-
served. The TDDM has a plausible neural implementation, in that the for-
malism is also a mathematical approximation of the net effect of excitation
and inhibition in the activity of a pool of neurons [18]. The model builds on
earlier theories, such as SET, by adopting a different, more complete mathemat-
ical formulation that allows modeling the trial-by-trial dynamics of timing (i.e.,
the learning), while still explaining core properties of interval timing (such as
timescale invariance). The general modular architecture (accumulator, memory
storage and decision rules) is preserved, however, raising interesting questions
as to the exact formal relationship between the TDDM and SET.

The TDDM has been successfully applied to some key features of interval
timing. Most notably, it can account for the scalar property, a ubiquitous
feature of timing data where the distribution of response times scales with the
interval being timed. The model has also been shown to learn quickly and to
reproduce the behaviour observed in fixed-interval schedules [17], the bisection
procedure [19], and tasks where time intervals are changing either randomly [18]
or cyclically [16], but has been only cursorily applied to the aggregate data in
the peak procedure thus far [9].

One of the main advantages of a mathematical model is the capacity to de-
rive precise quantitative predictions from as few assumptions as possible. In
this respect, the TDDM is particularly well placed among other timing models.
As previously demonstrated [9], Weber's law, which in the context of interval
timing manifests itself as a constant coefficient of variation (CV), follows from
adjustments made to the parameters of the inverse Gaussian distribution pre-
dicted by the model. In contrast, SET's main theoretical component—a Poisson
pacemaker—cannot by itself produce a constant CV. The usual solution is to
add an assumption that the noise in the memory for the remembered intervals

2

is so large as to overcome the noise in the pacemaker. This solves the problem but at the cost of adding an extra assumption and doing away with the Poisson pacemaker.

In this paper, we show that the TDDM can account for both the global averaged response curve in the peak procedure—and reproduce the statistics of behaviour in individual trials. We demonstrate this first analytically, through a new simplified formulation of TDDM's stationary properties, which we then show is equivalent to a constrained version of SET. The analytical results from the simplified model are then validated through simulations with the complete TDDM formulation.

These results extend the range of phenomena for which the TDDM can account and suggest that the Poisson pacemaker postulated by SET—but not actually used—may be substituted by the result of an opponent Poisson process [9]. Furthermore, and in light of previous successes, these results suggest there might be a single comprehensive drift-diffusion-based theory of decision making and timing, which could cover both the steady-state properties as well as the learning dynamics.

The paper next reviews the studies that have examined the patterns of correlations in single-trial analyses of the peak procedure. We then revisit the TDDM and develop a simplified stochastic model approximating TDDM's stationary properties. Given that formulation, the simplest possible extension of TDDM to support start and stop behaviours is analytically derived. This simplified formulation is shown to be equivalent to a constrained version of SET with two thresholds [12], and shown to be a good approximation of the full TDDM through simulations. Finally, some predictions are made with the full TDDM about possible sequential effects in the peak procedure.

1.1. Peak Procedure

In the peak procedure subjects are first trained on a fixed-interval (FI) schedule where the first response after an interval has elapsed since the appearance of a stimulus produces a reward (see diagram on the left in Fig. 1). When behavior on FI trials has stabilized, peak trials are then introduced. These peak trials are interspersed randomly between the normal FI trials, last 3 or 4 times longer, and are not rewarded. When peak trials are first introduced during training, animals start responding as usual before the FI time and then continue responding throughout the whole (long) interval. With sufficient experience with the peak trials, a different pattern emerges where responding eventually ceases or lowers in frequency soon after the expected reward time [20]. This pattern of starts and stops that appears after sufficient training is the focus of our analysis here.

The panel on the right in Fig. 1 shows an example of how, when response rate is averaged over these peak trials even for a single individual, there is an apparent smooth, symmetrical rise and fall in responding around the time of reinforcement. When individual trials are analyzed, however, a more abrupt shift in response rate is often observed. On many trials, animals start with a low response rate, switch to a high response rate, and then go back down
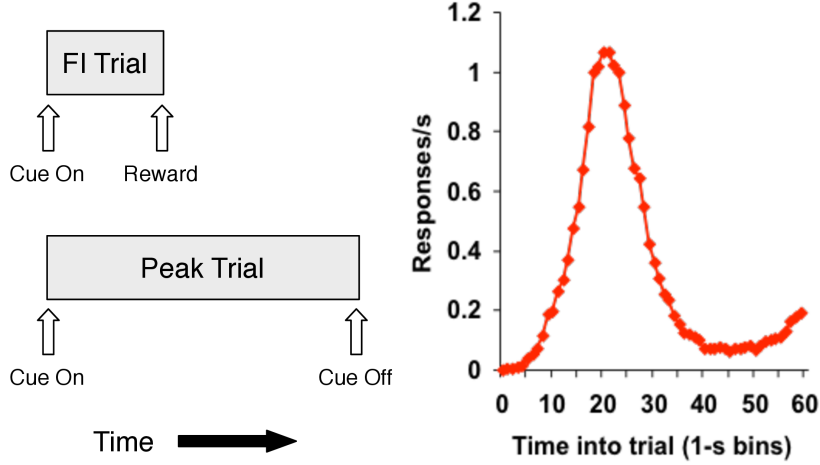
3

Figure 1: Peak procedure. Left: Schematic of fixed-interval (FI) and peak trials. Right: Response rate averaged over peak trials for one rat subject (data from [21]). Reward was ordinarily available after 20 seconds.

again after the usual interval has elapsed and no reward has arrived [22, 12, 23]. This three-state system (low-high-low) can be characterized by its two transition points: the start (switch from low to high) and stop (switch from high to low) times. In addition, the middle time and duration of the high-frequency bout can be calculated from the start and stop times.

A detailed analysis of these variables may shed light on the internal mechanisms of interval timing and provide constraints on current timing models. Table 1 collates the results from the major studies in the literature that have examined the statistics of these four variables. When possible we have separated the data by FI duration and, in only one case, also by reward magnitude. We did this because in a few cases the correlations were reported to be significantly different as a function of FI duration [23] and reward magnitude [24], although this was not the norm. Of particular note are the coefficients of variation of each variable and their correlations. Note the strong similarity in the correlation patterns across 4 different species. The key results are as follows:

1. Positive correlation between start (S_1) and stop (S_2): $\rho(S_1, S_2) > 0$;

2. Negative correlation between start and duration (D): $\rho(S_1, D) < 0$;

3. Positive correlation between duration and middle (M): $\rho(D, M) > 0$;

4. Coefficient of variation (CV) for the start larger than for stop: $\mathrm{CV_{start}} > \mathrm{CV_{stop}}$.

4

The correlation results above mean that, in general, start times that occur early/late into the trial are usually followed by early/late stops. In contrast, the duration of the period of high frequency responding tends to decrease with late starts and increase with early starts. Also, the coefficient of variation ($CV = \sigma/\mu$) is larger for starts than for stops.

The four properties above will be referred to as the *stationary* features of the peak procedure, to differentiate from the *dynamic* features related to learning when to stop when peak trials are first introduced. Existing timing theories do contend with some of this static data. SET can capture the first three features through the variance in its clock and memory processes [31] as can the $TD(\lambda)$ algorithm coupled with a neural net function approximator [32]. The Multiple Time Scales model [5] has the capacity to generate start and stop responses, but it is unclear whether that model can fit the full range of these timing features. Finally, BeT can be made to explain the correlation patterns if it is modified to include a variable transitional period when behavior switches (Gibbon and Church [31] called this modification the *quasi-serial model*, whilst Killeen and Fetterman [33] called it *augmented BeT*). As Church, Meck and Gibbon [12] showed, however, this model is incompatible with correlations observed between subjects which SET can explain. It remains an open question whether other timing theories can accommodate these features.

Given the formal similarity between BeT and the TDDM [9], this raises significant doubt as to whether the TDDM can be made to accommodate these trial-by-trial properties of the peak procedure. Below, the main formulation of the TDDM is briefly reviewed and then a new simplified version of TDDM's stationary properties is introduced which allows for the derivation of the four static properties above analytically. The parameters thus obtained with the simplified version are used with the full TDDM formulation; these simulations show that the results match appropriately.

2. Theory

2.1. The Time-adaptive Drift-diffusion Model (TDDM)

The drift-diffusion foundation of the TDDM is given by the stochastic differential equation:

$$dx = A \cdot dt + m \cdot \sqrt{A} \cdot dW \tag{1}$$

where $A > 0$ is the drift or slope and $m > 0$ is a constant that determines the amount of Gaussian white noise dW. Note that limiting A to be positive and scaling the noise by $\sqrt{A}$ are particular to the TDDM, not to DDMs in general.

In decision models, $x(t)$ represents the current level of evidence toward a conclusion. When $x(t)$ reaches a threshold, that indicates that enough evidence has been accumulated in favor of a corresponding conclusion. For example, if the process starts at $x(0) = 0$, two thresholds at $z^+ = 1$ and $z^- = -1$ could represent two different conclusions (in decision models, A can be negative).

Noisy evidence (noise dW and input evidence A in equation (1)) are accumulated by pushing x toward one conclusion or the other.

Similarly, in the TDDM the accumulation process gathers noisy evidence that time has elapsed with $A > 0$ and $x \geq 0$ [17, 16]. A single threshold $z > 0$ is used to mark the expected level of x at the time T of a salient event, a reward for example, such that x crosses z at time $t = T$ on average. The time it takes for x to reach z (since $A > 0$, this will eventually occur) represents the psychological or subjective estimate for a given time interval. Moreover, because the noise dW on x at every time step is Gaussian, the distribution of times t it takes for x to reach z, which is the inverse, is given by the inverse Gaussian distribution [9]:

$$p\left(t; \frac{z}{A}, \frac{z^2}{m^2 A}\right) = \frac{z}{m\sqrt{At^3 2\pi}} \exp\left(\frac{-(At-z)^2}{2Am^2 t}\right) \tag{2}$$

with mean $\mu = z/A$ and variance $\sigma^2 = m^2 z/A^2$. In the TDDM, different time intervals are learned by adapting the drift rate A while holding the threshold z fixed, such that $A \to z/T$. Therefore, A can be seen as the rate at which the time interval is elapsing. It is the time-adaptability of A combined with the $\sqrt{A}$ factor in equation (1) that differentiates TDDM from DDM, and that gives TDDM its timescale invariance property [9]. If the events to be timed are unit-size rewards, than A is equivalent to a reward rate.

The coefficient of variation of the time estimate produced by the TDDM is solely a function of the noise m and threshold z, and can be derived directly from the distribution (2):

$$\mathrm{CV}_{\mathrm{TDDM}} = \frac{m}{\sqrt{z}}. \tag{3}$$

Finally, equation (1) can be approximated numerically as

$$X(t + \Delta t) = X(t) + A \cdot \Delta t + m \cdot \sqrt{A \cdot \Delta t} \cdot \mathcal{N}(0, 1). \tag{4}$$

The time-adaptive property of TDDM comes from slope A taking the form of an exponential moving harmonic average of the observed intervals (Theorem 2 in[17]):

$$A_{n+1} = (1 - \alpha)A_n + \alpha \frac{1}{\hat{I}_n} \tag{5}$$

where $\hat{I}_n$ is the estimated time interval observed on the nth trial and $0 < \alpha \leq 1$ is the learning rate. Without loss of generality, if $z = 1$, then at the time of reward ($t = T$), the value of $X(t)$ (which must then be positive) produces a positive estimate $\hat{I}_n = X_n(T)/A_n$ of the time interval T. The application of the learning rule, when the cue indicating the end of the interval is observed, thus changes the slope A toward the inverse of this new perceived interval $\hat{I}_n$ such that:

$$A_{n+1} = (1 - \alpha)A_n + \alpha \frac{A_n}{X_n(T)} \tag{6}$$

6

Given each trial begins with $X(0) = 0$, it can be seen that intervals perceived to be longer than T (with $X_n(T) > 1$) will decrease the rate A toward the inverse of $\hat{I}_n > 0$, while intervals perceived shorter than T (with $X_n(T) < 1$) will increase it. We refer the interested reader to [17] and [9] for a full exposition of the learning rules at the real-time (time-step) level. In the case of the peak procedure, where the intervals are always of the same duration, equations (4) and (6) are the only ones needed.

Note that in the peak procedure, we assume the learning rule is only applied when the reward occurs. As a result, there are no updates to the rate A on peak trials. In line with this assumption, some experiments have reported that peak trials have little or no impact on the start and stop times [20, 21], but others have reported a leftward shift of the response curve following peak trials [12] (but no relationship with peak-trial duration). This latter finding suggests some influence of peak trials on later trials, but this effect would not seem to be mediated by an update of the memory for time as encoded in the drift rate A. Any update due to the long peak trials should actually induce later responding and a right-ward shift of the response curve (opposite to that observed). Instead, this finding probably reflects an incomplete reset of the accumulator following peak trials.

2.2. Application of the TDDM to the peak procedure

As suggested by [9], perhaps the simplest way to account for the abrupt switch back to a low response rate in individual trials of the peak procedure is to use two thresholds $z_1 < z_2$. The first threshold z_1 marks the time to start the high response rate, and the second threshold z_2 marks the time to stop. This modification would immediately produce the difference in the CV for starts and stops (feature #4) because by (3)

$$\text{CV}_{\text{start}} = \frac{m}{\sqrt{z_1}} > \frac{m}{\sqrt{z_2}} = \text{CV}_{\text{stop}}.$$

Note that feature #4 rules out the possibility of having two consecutive single-threshold timers (one determining the start time, and one determining the stop time), because starting the second timer when the first one reaches a threshold would force $\text{CV}_{\text{stop}} \geq \text{CV}_{\text{start}}$. Similarly, having two parallel single-threshold timers would break feature #1 by producing $\rho(S_1, S_2) = 0$, unless there is really only one timer, or unless some extra common factor is added. Thus, the single-timer two-threshold approach seems the most sensible option available. The mechanism for the emergence of such a second threshold during training remains unknown, but for the current purpose of explaining the stationary features of the peak procedure, we need only to assume its existence.

Due to TDDM's learning ability, however, it is not immediately clear what other modifications would be needed in order for the model to exhibit the four stationary properties of the peak data. Many variations on the model could possibly fit the data, but for reasons of parsimony we aimed to find a suitable variation with the fewest modifications possible. To aid this search in model

7

space, we built a simplified version of the TDDM without noise in the diffusion process (i.e. setting $m = 0$ in equation (4)) and without applying the slope-adaptation rules (equation (6)) directly. Instead, we propose to approximate the stationary distribution of the slope A as a result of these two interacting components.

2.3. TDDM: A simplified formulation

The learning rule for adapting A in equation (6) can be seen as a Robbins-Monro algorithm [34], which estimates the reward rate (assuming fixed unit-size rewards) for the stimulus on non-peak trials. The slope A_n is therefore an improving estimate of the reward rate given the noisy measure $\hat{I}_n = X(T)/A_n$ of the true time interval T provided by the accumulator at the time of reward. This algorithm, under some conditions, converges by the law of large numbers to a normal distribution centered on the reward rate. Under steady-state conditions, we can therefore approximate the full model by replacing the noise m in the accumulation or memory process and the application of the learning rule altogether by sampling the slope A_n at the beginning of each trial from a Gaussian distribution $\mathcal{N}(z/T, \sigma^2)$ clipped above 0.

Sampling the slope A produces a range of values across trials. On trials when the slope is high, both start and stop times are low; on trials when the slope is low, both start and stop times are high. As a result, there is a positive correlation across trials (feature #1, see Fig. 2). This sampling, however, is insufficient to generate a negative correlation between the start and the duration (feature #2).

As shown on the simplified TDDM diagram in Fig. 2, a simple way to account for the negative correlation while maintaining $CV_{start} > CV_{stop}$ (feature #4) is to add noise to the start threshold z_1. At the beginning of each trial, a threshold $Z_{1,n}$ is sampled from a probability distribution with a given mean and variance (for simplicity we used a uniform distribution $\mathcal{U}(a, b)$ where a and b are estimated from the data but in principle other distributions could have been used instead). Note that this increases the CV_{start} to which equation (3) does not apply anymore, while maintaining feature #4. On trials when the sampled threshold is low, the start time will be low and the duration high, while the opposite will happen when the threshold is high, giving rise to a negative correlation between the start time and the high response rate duration (feature #2, see Fig. 2).

Note that adding noise to the stop threshold instead of the start threshold would not be sufficient to fit the data. It would increase the CV_{stop}, which may then require additional noise in the start threshold to maintain feature #4 ($CV_{start} > CV_{stop}$). Thus, although adding noise to the stop threshold is also possible, adding noise only to the start threshold is simpler (and as we will show, appears to be sufficient).

A major advantage of this simplified formulation is that it allows a direct analytic derivation of estimates for the model parameters in the full TDDM which can reproduce static features 1-4 of the peak procedure, without the need for time-consuming optimization routines. The simplified model will first
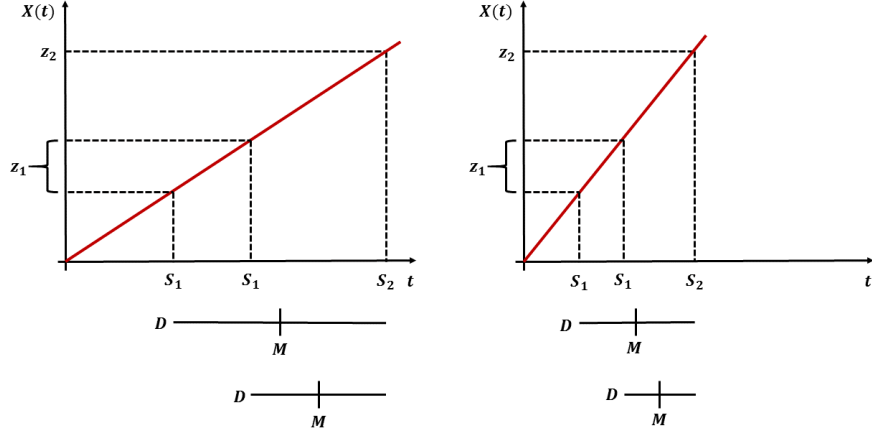
8

Figure 2: Diagram illustrating how the correlations seen in the data can be derived from the simplified model. The panels show two hypothetical trials: the left with a low slope A and the right with a high slope A. This variable slope induces a positive correlation between S_1 and S_2, and between D and M. To highlight the effect of the noise in threshold z_1, two corresponding values S_1 are shown (minimum and maximum) on each panel, with their corresponding durations and middle points. This noisy threshold induces a negative correlation between S_1 and D.

be used to directly generate predictions for the correlations. Secondly, it will be used to estimate most parameters of the full model directly from the same CVs. The full model will then be simulated to produce predictions for these same correlations, and to asses the degree of agreement between the data, the simplified model, and the full model. Finally, the simplified model will also be mapped to SET. Thus, the simplified model will provide a link between the CVs and the correlations, as well as between the full model and SET under stationary conditions.

3. Results

3.1. Simplified TDDM

Let S_1, S_2, M and D be the times of start time, stop time, middle and duration respectively with start and stop thresholds Z_1 and z_2 respectively. Then according to the simplified model:

$$S_1 = \frac{Z_1}{A}, \qquad\qquad M = \frac{S_1 + S_2}{2} = \frac{Z_1 + z_2}{2A},$$
$$S_2 = \frac{z_2}{A}, \qquad\qquad D = S_2 - S_1 = \frac{z_2 - Z_1}{A}.$$

9

Let Z_1 and $1/A$ be independent random variables with means $E(Z_1)$, $E(1/A)$ and variances $\sigma^2_{Z_1}$, $\sigma^2_{1/A}$ respectively. It is then possible to write all the statistics of the peak procedure in terms of the squared coefficients of variation of Z_1 and $1/A$, $\gamma^2_{Z_1}$ and $\gamma^2_{1/A}$ respectively, and the ratio of the thresholds $\theta = z_2/\mathrm{E}(Z_1)$ (see supplementary material for the complete mathematical derivations of these statistics). The squared coefficients of variation are:

$$\gamma^2_{S_1} = \gamma^2_{1/A} + \gamma^2_{Z_1} \cdot \gamma^2_{1/A} + \gamma^2_{Z_1}, \tag{7}$$

$$\gamma^2_{S_2} = \gamma^2_{1/A}, \tag{8}$$

$$\gamma^2_D = \frac{\gamma^2_{Z_1}(1 + \gamma^2_{1/A})}{(\theta - 1)^2} + \gamma^2_{1/A}, \tag{9}$$

$$\gamma^2_M = \frac{\gamma^2_{Z_1}(1 + \gamma^2_{1/A})}{(\theta + 1)^2} + \gamma^2_{1/A}. \tag{10}$$

The correlations are:

$$\rho(S_1, S_2) = \frac{1}{\sqrt{\gamma^2_{Z_1} + \frac{\gamma^2_{Z_1}}{\cdot\gamma^2_{1/A}} + 1}}, \tag{11}$$

$$\rho(S_1, D) = \frac{\gamma^2_{1/A}(\theta - 1 - \gamma^2_{Z_1}) - \gamma^2_{Z_1}}{\sqrt{\gamma^2_{1/A}(\gamma^2_{Z_1} + 1) + \gamma^2_{Z_1}} \cdot \sqrt{\gamma^2_{1/A}[(\theta - 1)^2 + \gamma^2_{Z_1}] + \gamma^2_{Z_1}}}, \tag{12}$$

$$\rho(D, M) = \frac{\gamma^2_{1/A}(\theta^2 - 1 - \gamma^2_{Z_1}) - \gamma^2_{Z_1}}{\sqrt{\gamma^2_{1/A}[(\theta + 1)^2 + \gamma^2_{Z_1}] + \gamma^2_{Z_1}} \cdot \sqrt{\gamma^2_{1/A}[(\theta - 1)^2 + \gamma^2_{Z_1}] + \gamma^2_{Z_1}}}. \tag{13}$$

We can demonstrate that a constant start threshold z_1, as opposed to a stochastic threshold, is qualitatively inconsistent with the data. If z_1 is constant, then $\gamma^2_{z_1} = 0$, which when substituted into equations (11) to (13) yields positive correlations for all three cases, contrary to the data:

$$\rho(S_1, S_2) = \frac{1}{\sqrt{0 + 0 + 1}} = 1,$$

$$\rho(S_1, D) = \frac{\gamma^2_{1/A}(\theta - 1)}{\sqrt{\gamma^2_{1/A}} \cdot \sqrt{\gamma^2_{1/A}(\theta - 1)^2}} = \frac{\gamma^2_{1/A}(\theta - 1)}{\gamma^2_{1/A}(\theta - 1)} = 1,$$

$$\rho(D, M) = \frac{\gamma^2_{1/A}(\theta^2 - 1)}{\sqrt{\gamma^2_{1/A}(\theta + 1)^2} \cdot \sqrt{\gamma^2_{1/A}(\theta - 1)^2}} = 1.$$

By inverting some of equations (7)-(13) we can estimate the model parameters from the data. We have a choice of which equations to use, the coefficients

10

of variation or the correlations. It is simpler to use the coefficients of variation (7)-(10):

$$\gamma_{1/A} = \gamma_{S_2} \tag{14}$$

$$\gamma_{Z_1} = \sqrt{\frac{\gamma_{S_1}^2 - \gamma_{S_2}^2}{1 + \gamma_{S_2}^2}} \tag{15}$$

$$\theta = \begin{cases} 1 + \sqrt{\frac{\gamma_{S_1}^2 - \gamma_{S_2}^2}{\gamma_D^2 - \gamma_{S_2}^2}} & \text{if using (9)} \\ -1 + \sqrt{\frac{\gamma_{S_1}^2 - \gamma_{S_2}^2}{\gamma_M^2 - \gamma_{S_2}^2}} & \text{if using (10)} \end{cases} \tag{16}$$

The equations above show that, with the simplified TDDM, three pieces of information from the data, namely the coefficients of variation of start, stop and either middle or duration, are sufficient to determine the three correlations in Table 1.

Although not necessary to find the correlations, one could simulate (or sample) the simplified TDDM using the same three CVs to determine its five parameters, which are the mean start threshold $E(z_1)$, stop threshold z_2, mean accumulator slope $E(A)$, start threshold variance $\sigma_{Z_1}^2$, and slope variance σ_A^2. In setting these parameters we have assumed without losing generality that the accumulator aimed at hitting a threshold $z = 1$ at the time of reinforcement, and that the two thresholds to start and stop responding are set so as to surround the reinforcement threshold: $(E(Z_1) + z_2)/2 = 1$. This assumption implies that the midpoint M of high frequency responding is exactly in the middle of start and stop, something that is roughly true (see Table 1 in [12]). Using $\theta = z_2/E(Z_1)$, it follows that $z_2 = \theta \cdot E(Z_1)$ and $E(Z_1) = 2/(1+\theta)$. The five simplified TDDM parameters, can be expressed directly in terms of data, namely the three CVs ($\text{CV}_{\text{start}} = \gamma_{S_1}$, $\text{CV}_{\text{stop}} = \gamma_{S_2}$, and $\text{CV}_{\text{dur}} = \gamma_D$), as follows:

$$E(Z_1) = \frac{2}{1 + \theta} = \frac{2}{2 + \sqrt{\frac{\gamma_{S_1}^2 - \gamma_{S_2}^2}{\gamma_D^2 - \gamma_{S_2}^2}}}, \tag{17}$$

$$\sigma_{Z_1}^2 = (\gamma_{Z_1} \cdot E(Z_1))^2 = \left(\frac{\gamma_{S_1}^2 - \gamma_{S_2}^2}{1 + \gamma_{S_2}^2}\right) E(Z_1)^2, \tag{18}$$

$$z_2 = E(Z_1) \cdot \theta = E(Z_1) \cdot \left(1 + \sqrt{\frac{\gamma_{S_1}^2 - \gamma_{S_2}^2}{\gamma_D^2 - \gamma_{S_2}^2}}\right), \tag{19}$$

$$E(A) = \frac{1}{T}, \tag{20}$$

$$\begin{aligned} \sigma_A^2 &= \sigma_{1/A}^2 \cdot E(A)^4 \\ &= E(1/A)^2 \cdot \gamma_{1/A}^2 \cdot E(A)^4 \\ &= 1/E(A)^2 \cdot \gamma_{1/A}^2 \cdot E(A)^4 = (\gamma_{S2} \cdot E(A))^2. \end{aligned} \tag{21}$$

11

In deriving σ_A^2 we made use of the Taylor approximation $\sigma_{1/A}^2 \approx E(A)^{-4} \cdot \sigma_A^2$ (see Appendix D). In the case of θ, if both CV_{mid} and CV_{dur} are available from the data we have a choice between which one to use (see equation (16)). We have used θ derived from CV_{dur} in all models because this value provided a better agreement with the data, probably because it conveys information about the anti-correlated noise between the start and stop which is filtered out in the middle point. Notice that the correlation equations (11) to (16) do not take as input the interstimulus interval T of the FI trials, which is coded in the variable $1/A$; making the correlations completely independent of the FI timescale in this model. The estimates for the five simplified TDDM parameters are calculated with equations (17) to (21) based on the three CVs from the data.

The correlations derived from the simplified TDDM can be found under the column **Simp.** in Table 2. Because the model takes as input the coefficients of variation for start, stop and duration, we were not able to model the two studies [22, 25] that did not report these variables. Out of the 19 experiments left, only two violated constraints imposed by the model. The data with FI 5 seconds in [23] has $\gamma_D^2 = \gamma_{S2}^2$ which causes the denominator in the radical in equation (19) (stop threshold) to be zero. The data in [24] for FI 17 seconds and LOW reward has $\gamma_{S2}^2 > \gamma_{S1}^2$ which generates imaginary values for the start (equation (17)) and stop (equation (19)) thresholds. Therefore we did not report the correlations produced by the models for these two experiments.

In some cases, one or more correlations were not reported in the original paper (marked as n/a in Table 2), so in these cases results are provided as predictions only. Considering all 44 correlations available out of the 17 experiments with the appropriate CVs pattern, the model's correlations deviate by 0.09 on average from the correlations in the data, with the worst case being $\rho(S_1, S_2)$ in [24] for FI 10H (see also Figure 3, first bar of each plot). In five cases (out of 51) the model generates the wrong correlation sign, but those values are very close to zero (≤ 0.03, marked by an * in Table 2), and two of them are not available in the animal data to be verified. Such very small, wrong-signed, correlations can also be observed in [24] for FI 17 seconds and LOW for $\rho(S_1, D)$. In general, such small correlations are not robust enough to be clearly considered of a specific sign. Overall, using only three pieces of information from the data the simplified model reproduces the pattern of correlations generally well and generates values similar to the data in most cases as shown on the left-most bar of each plot in Figure 3.

3.2. Full TDDM

The full model (equations (4) and (6)) was then simulated assuming learning occurs only on rewarded trials (and $X_n(0) = 0$ at the beginning of each trial). Peak and rewarded trials were randomly generated following the proportions used in the original animal experiments and using the same interval lengths (although the model is completely independent of the interval length). The parameters were set using the simplified model equations, with the DDM noise parameter set to $m = \gamma_{1/A}$ by equation (3) and the assumption that the learning
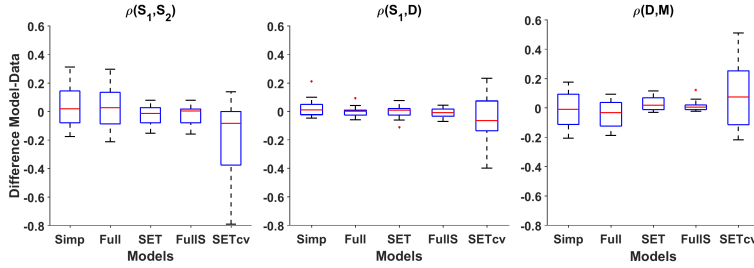
12

Figure 3: Box-plot of the differences in correlations between the animal data and each model for all 16 experiments to which the simplified TDDM can be fit. Each plot represents, from left to right: the simplified TDDM (Simp), the full TDDM using the same parameters values (Full), SET fit to the correlations (SET), the full TDDM using parameters from SET (FullS), and SET fit using only the same 3 CVs as the simplified model (SETcv). In short, the simplified model and SET can both provide good estimates of their equivalent TDDM on stationary conditions. Though SET gets lower error on correlations than the simplified TDDM, it does so at the cost of error on CVs. SET, however, is not as good as the simplified model at predicting the correlations from CVs only (last column of each sub-plot).

³³⁵ rule aims at threshold $z = 1$ when the reward occurs. The slope was initialized with the exact value $A = 1/T$. The only remaining free parameter was the learning rate α. For each experiment, and each $\alpha \in \{0, 0.05, 0.10, ..., 1.00\}$, we ran 30 simulations of 1000 trials each, to gather means and standard deviations for each measure of interest, namely the 4 CVs and the 3 correlations.

³⁴⁰ The best matching correlation triplets derived from the full TDDM can be found under the column **Full** in Table 2. Considering all 44 correlations available out of the 17 experiments with the appropriate CVs pattern, the model's correlations deviate only by 0.08 on average from the correlations in the data, with the worst case being again $\rho(S_1, S_2)$ in [24] for FI 10H. The simulations ³⁴⁵ did not generate the wrong correlation sign on average, except for 1 experiment from [30], where the model predicts $\rho(D, M) \approx -0.02$, but the real value is unknown. As shown in Table 2, the correlations are quite similar, across all results (animal, simp. TDDM, and full TDDM). Despite only one wrong correlation sign (as opposed to five for the simplified model) out of 51, the full TDDM gives ³⁵⁰ very similar results to the simplified TDDM on this task. As shown in Figure 3 (two first bars of each plot), the distribution of errors is very similar for both models. Moreover, on average, the absolute difference between a correlation in the simplified model and the full model is 0.04. These results show that the simplified model is a good approximation of the full model under stationary ³⁵⁵ conditions. In addition, the extra learning parameter in the full model is shown to be of little use in explaining the current datasets on the intra-trial structure of responding on the peak procedure. This outcome is as expected because the extra learning rate parameter in the TDDM should play an important role mostly when learning new intervals, but is expected to have little impact on ³⁶⁰ stationary behaviours.

To better evaluate the quality of these fits, we examined the absolute error

13

on the CVs between the original data and the CVs produced by the simulations of the full TDDM. The average absolute CV error was 0.08 with the worst case being 0.17 in [24] for FI 10H, CV_{dur}. This discrepancy, albeit small, is most likely due to the setting of $m = \gamma_{1/A}$, as $\gamma_{1/A}$ represents a mixture of perceptual noise m and the learning rate α. Whether there could be an analytic formula to set those parameters (m and α) remains an open question. The full model has an extra degree of freedom in the learning rate α, but to properly estimate this balance between m and α would require additional information about the learning dynamics, before steady-state occurs. In the simplified model, $\gamma_{1/A}$ merges m and α parameters from the full model, and this issue does not arise. Nonetheless, we can still estimate the learning rate, which across all 17 experiments was consistently high and ranged between 0.45 and 0.85 (with a median of 0.70). Moreover, in about 9 of the 16 experiments for which α has been optimized (for [28], we used the median 0.70 since no correlations were available), a small learning rate could adversely impact the sign of $\rho(D, M)$. A high TDDM learning rate was also found to best fit a different timing task in [18], suggesting that a high learning rate may be necessary, even when learning is not directly observable.

To further investigate this hypothesis, we measured the correlation between the perceived time estimate of the model (provided by $1/X_n(T)$) on reward trials preceding peak trials and the stop time of the subsequent peak trials. The average correlation across all 17 experiments was 0.61, as expected from such a high learning rate (e.g., for $\alpha = 0.7$, the last 2 trials account for 91% of the memorized interval). Unfortunately, it might be challenging to measure an animal's perceived time estimate for a single trial to see such trial-to-trial correlations in behaviour. Another possibility would thus be to slightly change the rewarded interval on those pre-peak reward trials to generate enough measurable changes in responding on the subsequent peak trials. We found that varying those intervals using 1.33 times the CV was sufficient to produce a correlation between one reward trial's intervals and the following peak trials stop times for all 17 experiments (generating a correlation $> .47$ across all experiments with a mean of 0.61). For example, if the FI is 20s and the $CV_{stop} = 0.15$, then using pre-peak trial intervals of 16s and 24s should be enough to detect a learning effect. This prediction remains to be tested in animals.

Finally, as an illustration that the TDDM can produce similar average curves seen in the data, Fig. 4 shows an example of how the model predicts the start and stop times on each trial and how averaging these times produces the peak in response rate around the time of reinforcement.

3.3. Relationship with SET

An important further observation is that the simplified TDDM presented above is formally equivalent to a constrained version of the two-threshold version of SET developed for the peak procedure in [12].

SET incorporates 4 components: a clock (or accumulator), a memory of reinforcement times, a source of noise, and a comparator. Originally, the accumulator worked by accumulating noisy Poisson pulses, but this within-trial noise
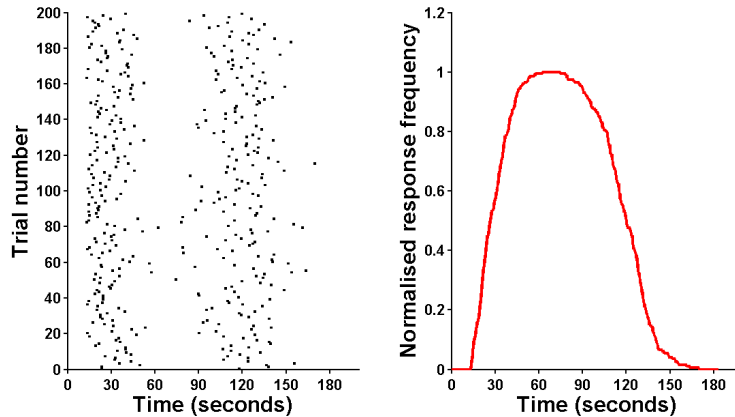
Figure 4: TDDM simulation using parameters derived from data in [35]. Left plot: start and stop times for each peak trial. Right plot: averaging over all trials produces the peak of response frequency seen in data.

was removed in more modern versions (either because it was unnecessary or because it would break timescale invariance [36, 5]). Thus, SET has a noise-free accumulator n representing the *perfect* estimate of elapsed time in the current trial. The memory is the memory of reinforcement times. It is represented as a distribution determined by an expected value and a coefficient of variation providing across-trial variability and timescale invariance. At the beginning of each trial, a sample n^* is generated from the memory and used for comparison with the accumulated running time n. There is finally a comparison threshold b such that the animal starts responding when the ratio $n/(n^* \cdot k^*) > b$, where k^* is a multiplicative noise source, which could easily be interpreted as related to clock speed (k^* is usually introduced before saving the time to memory in SET, but this aspect is not important here).

For the peak procedure, a second threshold was introduced for the stop time in [12]. Because an extra source of noise was necessary to explain the trial statistics, two stochastic thresholds were used, $l = 1 - b_1$ and and $u = 1 + b_2$ (each having a mean and variance). Thus, the model starts responding a short time before the remembered interval and stops a short time after. Note that the single threshold in the original version of SET was considered constant. Adding noise to the start threshold is necessary to get the pattern of inter-trial correlations discussed here, but the simplified TDDM results suggest that noise in the second threshold is not necessary. The expected value of the time memory is assumed to come from experience and is therefore fixed by the task. The remaining parameter of the model is the CV of that memory γ (k^* is not mentioned in [12]).

The main difference between SET and the simplified TDDM is terminological (a symbol mapping is provided in Table 3). In the original full TDDM

15

there are two sources of noise generating trial-to-trial variability because the start threshold z_1 was regarded as a constant. The first source comes from the accumulator, in the term $m \cdot \sqrt{A} \cdot dW$ in equation (1). The drift rate $A = 1/T$ encodes the interval memory, and the $\sqrt{A}$ factor generates timescale invariance on reading rate A. The second source of noise comes from trial-to-trial variability in A due to the learning rule in equation (6). These two sources of variability are merged together in the simplified TDDM in the distribution of A_n which uses the trained interval as its expected value. Thus A_n in the simplified TDDM is equivalent to SET's n^* memory, except that A_n represents rates while n^* represents durations.

In adapting the TDDM to the peak procedure data, we found that it was analytically sufficient to add noise only into the start threshold (noise in the stop threshold can be added, but must remain smaller than in the start threshold). Thus, while in [12] SET uses two noisy thresholds, the simplified TDDM shows that a noisy start threshold and a fixed stop threshold are sufficient to explain the data. Moreover, SET allows for full fitting of the two threshold positions (in terms of proportion of time interval). Because the CVs themselves cannot give threshold locations, but only some information about their relative time positions, we use the fact that the peak response is near the FI to estimate a threshold ratio θ and locate the thresholds on both sides of it. Thus, the TDDM has only 3 parameters directly estimated from specific measurements (3 CVs) whereas SET requires a search process to optimize its 5 parameters. Beyond that, the thresholds in SET and the simplified TDDM are the same.

The key equations of the simplified TDDM (Equations (7) to (13)) are equivalent to the two-threshold SET equations in Fig. 13 in [12]. Table 3 further shows the correspondence between SET and the simplified TDDM in terms of their parameters. As an example, we show how SET's formula for the coefficient of variation of the middle time (cv(m) in Fig. 13 in [12] is equivalent to γ_M in the Simplifed TDDM):

$$\gamma_+^2 = \frac{\sigma_u^2 + \sigma_l^2}{(U+L)^2} = \frac{\sigma_{z_2}^2 + \sigma_{Z_1}^2}{(z_2 + E(Z_1))^2} = \frac{\sigma_{Z_1}^2}{(\theta E(Z_1) + E(Z_1))^2} = \frac{\sigma_{Z_1}^2}{E(Z_1)^2 (\theta + 1)^2}$$

$$= \frac{\gamma_{z_1}^2}{(\theta + 1)^2}.$$

$$\mathrm{cv}^2(\mathrm{m}) = \gamma^2 (1 + \gamma_+^2 (1 + \gamma^{-2}))$$

$$= \gamma^2 + \gamma_+^2 \gamma^2 + \gamma_+^2 = \gamma_{1/A}^2 + \frac{\gamma_{Z_1}^2}{(\theta + 1)^2} \gamma_{1/A}^2 + \frac{\gamma_{Z_1}^2}{(\theta + 1)^2}$$

$$= \frac{\gamma_{Z_1}^2 (1 + \gamma_{1/A}^2)}{(\theta + 1)^2} + \gamma_{1/A}^2 = \gamma_M^2,$$

because $\sigma_{z_2}^2 = 0$ as we assume here that z_2 is constant. The other six equations in Fig. 13 in [12] can be transformed into the notation from the simplified TDDM in a similar manner.

To quantitatively compare SET and the simplified TDDM, we implemented SET and, using a grid search, found the parameters that best fit the 3 correla-

tions for each experiment. In the grid search, the mean memory time was the FI, the CV could vary from 0.05 to 0.25, the largest CV, by steps of .05, the two thresholds could vary from .1 to .9 by steps of .1 on each side of the FI independently, and their standard deviations could vary from 0.00 to $0.75 - \mathrm{CV}_{\text{param value}}$ and $0.25 - \mathrm{CV}_{\text{param value}}$ for the start and stop threshold respectively, where $\mathrm{CV}_{\text{param value}}$ was the value of the CV parameter for a given grid point. In the simulations, if the memory time or start threshold where negative or if the stop threshold was smaller than start threshold, these were then resampled. If a given set of parameters generated 10% or more of such invalid values, that set of parameters was skipped.) A total of 30 simulations of 1000 trials each were run for every grid point (as in the full TDDM). The mean absolute error on all available correlations for all 17 experiments TDDM could fit was 0.05 (as opposed to 0.09 for the simplified model), and there was no sign error (comparing the first and third bars on each plot of Figure 3).

The absolute difference in parameters value between the two models is provided in the last column of Table 3. The noise levels were very close to each other, including the stop threshold. The main difference was in the extra freedom SET has to position the start and stop thresholds relative to the FI. This error reduction, however, came at a cost. SET has 2 extra parameters and the fitting used all 7 pieces of information available, compared to the simplified model that only used 3. Whereas the simplified model used the exact CVs, SET produced substantial error on each of them (0.04 on average, for each CV). Thus, the reduced error on the correlations was counteracted by increased error on the CVs. When fitting SET using only the same 3 CVs as the simplified TDDM, SET's absolute error on correlations increased from 0.05 to 0.19 (more than double the simplified model, see the last bars of each plot on Figure 3 as compared to the third bars). Thus, SET has a little less predictive power, but the discrepancy between the SET and TDDM thresholds would be better analysed by looking at the start and stop temporal positions relative to the FI in future studies.

To further assess whether SET (like the simplified model) could be used as an estimate of the full TDDM under stationary behaviours, we augmented the full TDDM to support two thresholds with specific means and variances. We then ran the same simulations as before, but feeding the SET parameters directly into this augmented TDDM as estimates (using $m = \gamma$ and optimizing α to match the correlation as before). The average absolute error for this augmented full TDDM average was 0.04 (as opposed to 0.05 for SET alone, comparing the third and fourth bars on each plot of Figure 3), and the average absolute difference in correlation between SET and its TDDM equivalent was 0.03 (as opposed to 0.04 for the simplified TDDM and its full equivalent). The learning rate was also similar to the previous TDDM simulations with $\alpha \geq 0.55$ for all experiments with a median of 0.90, thus confirming the prediction about sequential effects made in the previous section. In short, the TDDM, although well-known for modelling the learning dynamics, can be well approximated by SET (with the same parameters) under stationary conditions.

We have thus established the precise mapping between the simplified TDDM

and SET, as well as having showed that they can both serve to provide estimates for the parameters for their TDDM equivalent under stationary conditions. The simplified TDDM model provides a way to show that adding noise in the start threshold and having a fixed stop threshold is sufficient for the full TDDM to fully account for the micro-structure of peak-interval data, and to link SET to the TDDM under stationary conditions without removing the TDDM's ability to also reproduce trial-to-trial time learning dynamics [18, 16]. Moreover, the formulae developed to estimate the TDDM's parameters can also be turned into formulae to estimate SET's parameters directly from the CVs, without the need to run thousands of simulations to find the best possible parameters (as in [12]). While SET can also be used as an estimate of the full TDDM for stationary conditions, a major advantage of the full TDDM over SET is that it is not limited to stationary behaviours, as it also covers time-adaptive phenomena through its learning rule.

4. Discussion

This paper extends the TDDM to deal with the micro-structure of responding on the peak procedure, when rewards are occasionally omitted. A simplified version of the TDDM was introduced, which allows for analytic derivation of model predictions for steady-state behaviour. This simplified version is also shown to be equivalent to a constrained version of Scalar Expectancy Theory, making SET a special case of the broader drift-diffusion framework.

This paper also provides, to the best of our knowledge, the most comprehensive review of single-trial analyses in the peak procedure, which reveals a consistent pattern of correlations and CVs observable in the trial micro-structure across 19 of 21 data sets. The TDDM and its simplified model were both able to provide good fits for most correlations available as well as to make realistic predictions for those which were not, such as in [28]. The two cases in which they both failed to fit the data—FI 5 in [23] and FIs 17L in [24]—were in direct violation of the model's constraints imposed by equations 15 and 16 (the data had either $\gamma_{S_2}^2 > \gamma_D^2$ or $\gamma_{S_2}^2 > \gamma_{S_1}^2$). In a few cases the simplified model had the wrong sign, but those predictions were too close to zero for the sign to be clear and in some of those cases, the experimental correlations were not available.

The TDDM, in its original formulation, could be expanded in many different ways to fit the peak procedure data. For example, multiple accumulators, in parallel, or in sequence, could be used to determine when to start and to stop, with each potentially having different noise factors, thresholds, and learning rates. Given that using distinct timers for the start and stop could easily be eliminated, we developed an approximation or simplified model of the TDDM using a single timer for steady-state or stationary behaviours. We then used it to find the simplest modifications that could account for the peak procedure statistics. By adding to this simplified TDDM an extra threshold to account for the stop time on peak trials and noise to the start time threshold, we were able to generate predictions that fit the data well. Simulations using the full TDDM formulation corroborated the simplified TDDM results and also validated the

18

simplified version as a good approximation of the full TDDM's steady-state behaviours. Thus, in future studies, this simplified version can be useful in formulating hypothesis and generating parameters that can later be tested with the full TDDM.

The paper also demonstrates that this simplified TDDM is formally equivalent to a specific version of SET on many tasks covered by SET. Therefore, a direct comparison between these two models could be made. Although mathematically equivalent, there are possible differences in interpretation. For example, SET postulates a Poisson pacemaker but this feature does not influence timing noise, because, within SET the error in reading and writing in memory overcomes the Poisson variability. In the TDDM, the Poisson noise resides in the opponent-process underlying the noise m in the DDM [18]. That DDM noise can be considered both the error in reading the slope value at every time step [17] as well as perceptual error generated by this accumulation of noise. This error, combined with the learning process, seems sufficient to explain the inter-trial variability covered in SET by the memory error. Finally, though SET works in steady-state tasks such as the peak procedure, the theory has no possible mechanism to account for any learning dynamics. While we showed that SET can be used to estimate the full TDDM (or its parameters), the full TDDM on the other hand, can adapt its drift rate to also account for behaviour on dynamic tasks, as previously demonstrated in the case of random and cyclic schedules of reinforcement [18, 16]. This adaptive element may also predict hitherto unreported effects. Depending on the values chosen for α and m, the full TDDM could show sequential trial-to-trial dependencies by appropriate manipulation of the FI on reward trials preceding peak trials. A deeper account of such sequential effects would require data from the beginning of training to set α and m more precisely, but some testable predictions were made with the model assuming α stays high (≥ 0.5).

While the question of whether there is one central clock or a bunch of distributed timers in the brain is still debated [37], SET usually posits a central clock, whereas the TDDM does not require such a central pacemaker and could very well suppose different accumulators for each interval being timed. The data from [23] also contains an experiment looking at the timing of multiple durations, which could pose a challenge to such a multiple timers idea. In one experiment, 3 distinct durations were used, allowing evaluation of the start and stop times for the different timed durations in the same peak trial. Interestingly, they found that stop times for the different durations were correlated within trials, suggesting a common clock noise (in SET), which would go against having the 3 different intervals in 3 different integrator rates for the TDDM. There is, however, also evidence that diffuse neurotransmitters, such as dopamine [38], can influence time perception. Thus, in a multiple-duration procedure, dopamine might diffusely influence all the timers in parallel. In the TDDM, this could be modelled by adding a common k factor in front of A in equation (1), affecting all timers speed equally. Such a parameter could also be considered equivalent to SET's k^* parameter and could be added to the simplified model as well, keeping both models similar (but [23] only scratched the surface on this point).

19

It would be necessary to develop formulae to isolate the variability of k across trials from the other sources of noise, possibly from the stop time correlation in the multiple duration trials. This is beyond the scope of this paper.

The TDDM also provides a degree of neural plausibility that most other timing models lack. For example, Simen et al. [18] showed how the TDDM could be represented at a lower level by a pool of Poisson firing neurons, even at long multi-second timescales [18, 39]. In addition, neurons with climbing activity which adapt their climbing rate to the necessary delay (exactly like TDDM) have already been recorded in numerous brain areas over short timescales, such as the thalamus (adapting between 1- to 2-s intervals) [40], the posterior parietal cortex (adapting between 300ms and 800ms intervals) [41], and in the inferotemporal cortex (adapting between 4 and 8 seconds intervals) [42].

The TDDM, however, can reproduce adaptive behavioural data on timescales up to 90 s and technically should work on static data for any timescale —not only the shorter timescales above. Recent work by Mello et al. [43] fills this gap in the timescale of the neural substrate. They found evidence for an adaptive neural representation in the striatum, which changes for intervals varying from 12 to 60 s. In rats performing a timing task, neurons from a pool of striatal neurons fired in sequence with each neuron roughly representing a different elapsed time. When the time interval was changed, this neural representation adapted to the interval duration, with the same neuron representing a similar relative time, but different absolute time. Such a dynamic adjustment of the time scale is a fundamental element of the TDDM, but lies outside the scope of SET's static representation.

The adaptive nature of the neural representation also does not accord with most learning models of timing which assume a fixed and predetermined time-code [5, 6, 7, 8]. An important outstanding question is how a population of spiking neurons such as the one found [43] could implement the learning rule as defined in equation (5) from [18, 17]. Nonetheless, the TDDM is perhaps unique among timing models to incorporate an adaptive representation of elapsed time, in accordance with the neural data, and to still fit the behavioural data under static conditions. Therefore, the TDDM seems well positioned as a simple high-level model that also accurately express lower-level neural representations.

One particularly appealing feature of the TDDM is the breadth of situations to which similar DDM formalisms can be applied. Whereas SET is a dedicated timing model, the TDDM is a model based on some of the same principles of the more general drift-diffusion class of models, which are well established in decision making and encounter wide physiological support [44, 14]. Though the computations are not strictly identical, in that the drift rate is adaptively scaled in the TDDM, the fact that a similar drift-diffusion process may underlie such diverse areas as interval timing and decision making suggests this process may be, to borrow a concept from theoretical neuroscience [45], a canonical cognitive computation.

References

[1] C. V. Buhusi, W. H. Meck, What makes us tick? Functional and neural mechanisms of interval timing., Nature Reviews Neuroscience 6 (10) (2005) 755–765.
URL http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed{&}id=16163383{&}retmode=ref{&}cmd=prlinkspapers2://publication/doi/10.1038/nrn1764

[2] S. Grondin, Timing and time perception: a review of recent behavioral and neuroscience findings and theoretical directions., Attention, perception & psychophysics 72 (3) (2010) 561–582.
URL http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed{&}id=20348562{&}retmode=ref{&}cmd=prlinkspapers2://publication/doi/10.3758/APP.72.3.561

[3] J. Gibbon, Scalar expectancy theory and Weber's law in animal timing, Psychological Review 84 (3) (1977) 279–325.

[4] P. R. Killeen, J. G. Fetterman, A behavioral theory of timing., Psychological Review 95 (2) (1988) 274–295.
URL http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed{&}id=3375401{&}retmode=ref{&}cmd=prlinkspapers2://publication/uuid/8D608E96-488C-470F-AD23-7F0A65FEF58C

[5] J. E. R. Staddon, J. J. Higa, Time and memory: towards a pacemaker-free theory of interval timing., Journal of the experimental analysis of behavior 71 (2) (1999) 215–251.
URL http://www.ncbi.nlm.nih.gov/pmc/articles/PMC1284701/papers2://publication/uuid/BDD2DD5B-324A-4E12-B7A6-B8A5964C5643

[6] S. Grossberg, N. A. Schmajuk, Neural dynamics of adaptive timing and temporal discrimination during associative learning, Neural Networks 2 (2) (1989) 79–102.
URL http://www.sciencedirect.com/science/article/pii/0893608089900269papers2://publication/doi/doi:10.1016/0893-6080(89)90026-9

[7] E. A. Ludvig, R. S. Sutton, E. J. Kehoe, Stimulus representation and the timing of reward-prediction errors in models of the dopamine system., Neural computation 20 (12) (2008) 3034–3054.
URL http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed{&}id=18624657{&}retmode=ref{&}cmd=prlinkspapers2://publication/doi/10.1162/neco.2008.11-07-654

[8] A. Machado, M. T. Malheiro, W. Erlhagen, Learning to Time: a perspective., Journal of the experimental analysis of behavior 92 (3) (2009)

21

423–58. doi:10.1901/jeab.2009.92-423.
URL   http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.
fcgi?dbfrom=pubmed{&}id=20514171{&}retmode=ref{&}cmd=
prlinkspapers2://publication/doi/10.1901/jeab.2009.
92-423http://www.pubmedcentral.nih.gov/articlerender.fcgi?
artid=2771665{&}tool=pmcentrez{&}rendertype=abstra

[9] P. Simen, F. Rivest, E. A. Ludvig, F. Balci, P. Killeen, Timescale Invari-
ance in the Pacemaker-Accumulator Family of Timing Models, Timing &
Time Perception 1 (2) (2013) 159–188. doi:10.1163/22134468-00002018.
URL http://wrap.warwick.ac.uk/58690/http://booksandjournals.
brillonline.com/content/journals/10.1163/22134468-00002018

[10] S. Roberts, Isolation of an internal clock., Journal of Experimental Psy-
chology: Animal Behavior Processes 7 (3) (1981) 242–268. doi:10.1037/
0097-7403.7.3.242.
URL http://doi.apa.org/getdoi.cfm?doi=10.1037/0097-7403.7.3.
242

[11] A. Machado, Learning the temporal dynamics of behavior., Psychological
Review 104 (2) (1997) 241–265.
URL   http://eutils.ncbi.nlm.nih.gov/entrez/eutils/
elink.fcgi?dbfrom=pubmed{&}id=9127582{&}retmode=
ref{&}cmd=prlinkspapers2://publication/uuid/
087AAE03-52E2-498A-8A18-92A70CE44EF5

[12] R. M. Church, W. H. Meck, J. Gibbon, Application of scalar timing
theory to individual trials., Journal of Experimental Psychology: Animal
Behavior Processes 20 (2) (1994) 135–155.
URL   http://eutils.ncbi.nlm.nih.gov/entrez/eutils/
elink.fcgi?dbfrom=pubmed{&}id=8189184{&}retmode=
ref{&}cmd=prlinkspapers2://publication/uuid/
5983C704-7A71-4F1A-89E9-4DACE7E601FE

[13] R. Ratcliff, A theory of memory retrieval., Psychological Review 85 (2)
(1978) 59.
URL   http://psycnet.apa.org/journals/rev/85/2/59/papers2:
//publication/uuid/C23DBAF7-0C7E-434A-8A3C-D9F0945328B8

[14] J. I. Gold, M. N. Shadlen, The neural basis of decision making., Annu.
Rev. Neurosci. 30 (2007) 535–574.
URL   http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.
fcgi?dbfrom=pubmed{&}id=17600525{&}retmode=ref{&}cmd=
prlinkspapers2://publication/doi/10.1146/annurev.neuro.29.
051605.113038

[15] A. Voss, M. Nagler, V. Lerche, Diffusion models in experimental psychol-
ogy: a practical introduction., Experimental psychology 60 (6) (2013) 385–

402. `doi:10.1027/1618-3169/a000218`.
URL `http://www.ncbi.nlm.nih.gov/pubmed/23895923`

[16] A. Luzardo, E. A. Ludvig, F. Rivest, An adaptive drift-diffusion model of interval timing dynamics, Behavioural Processes 95 (2013) 90–99. `doi:10.1016/j.beproc.2013.02.003`.
URL `http://www.ncbi.nlm.nih.gov/pubmed/23428705http://linkinghub.elsevier.com/retrieve/pii/S0376635713000247`

[17] F. Rivest, Y. Bengio, Adaptive Drift-Diffusion Process to Learn Time Intervals, Arxiv preprint arXiv:1103.2382.
URL `http://arxiv.org/abs/1103.2382`

[18] P. Simen, F. Balci, L. de Souza, J. D. Cohen, P. Holmes, A model of interval timing by neural integration., The Journal of neuroscience : the official journal of the Society for Neuroscience 31 (25) (2011) 9238–9253. `doi:10.1523/JNEUROSCI.3121-10.2011`.

[19] F. Balcı, P. Simen, Decision processes in temporal discrimination., Acta psychologica`doi:10.1016/j.actpsy.2014.03.005`.
URL `http://www.sciencedirect.com/science/article/pii/S0001691814000808`

[20] F. Balci, C. R. Gallistel, B. D. Allen, K. M. Frank, J. M. Gibson, D. Brunner, Acquisition of peak responding: What is learned?, Behavioural Processes 80 (1) (2009) 67–75.
URL `http://linkinghub.elsevier.com/retrieve/pii/S0376635708002222papers2://publication/doi/10.1016/j.beproc.2008.09.010`

[21] E. A. Ludvig, K. Conover, P. Shizgal, The effects of reinforcer magnitude on timing in rats., Journal of the experimental analysis of behavior 87 (2) (2007) 201–218.
URL `http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed{&}id=17465312{&}retmode=ref{&}cmd=prlinkspapers2://publication/uuid/882D9138-04C7-4DE2-BD43-0397C8BBC4E9`

[22] J. Gibbon, R. M. Church, Representation of time, Cognition 37 (1) (1990) 23–54.
URL `http://www.sciencedirect.com/science/article/pii/001002779090017Epapers2://publication/uuid/9AF8B274-28F2-43B1-9C15-F40779F668BA`

[23] C. R. Gallistel, A. King, R. McDonald, Sources of variability and systematic error in mouse timing behavior., Journal of experimental psychology. Animal behavior processes 30 (1) (2004) 3–16. `doi:10.1037/0097-7403.30.1.3`.
URL `http://www.ncbi.nlm.nih.gov/pubmed/14709111`

[24] F. Balcı, M. Wiener, B. Cavdaroğlu, H. Branch Coslett, Epistasis effects of dopamine genes on interval timing and reward magnitude in humans., Neuropsychologia 51 (2) (2013) 293–308. doi:10.1016/j.neuropsychologia.2012.08.002.
URL http://www.sciencedirect.com/science/article/pii/S0028393212003417

[25] C. V. Buhusi, W. H. Meck, Relativity theory and time perception: single or multiple clocks?, PloS one 4 (7) (2009) e6268. doi:10.1371/journal.pone.0006268.
URL http://dx.plos.org/10.1371/journal.pone.0006268papers2://publication/uuid/0B0B0A09-9C7F-4602-9708-A642322EC667http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2707607{&}tool=pmcentrez{&}rendertype=abstract

[26] K. Cheng, R. Westwood, J. D. Crystal, Memory variance in the peak procedure of timing in pigeons., Journal of Experimental Psychology: Animal Behavior Processes 19 (1993) 68–76. doi:10.1037/0097-7403.19.1.68.

[27] B. C. Rakitin, J. Gibbon, T. B. Penney, C. Malapani, S. C. Hinton, W. H. Meck, Scalar expectancy theory and peak-interval timing in humans., Journal of Experimental Psychology: Animal Behavior Processes 24 (1) (1998) 15–33. doi:10.1037/0097-7403.24.1.15.
URL http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed{&}id=9438963{&}retmode=ref{&}cmd=prlinkspapers2://publication/uuid/89530A59-6A09-4439-B83F-61EB28CEB739

[28] M. S. Matell, M. Bateson, W. H. Meck, Single-trials analyses demonstrate that increases in clock speed contribute to the methamphetamine-induced horizontal shifts in peak-interval timing functions., Psychopharmacology 188 (2) (2006) 201–12. doi:10.1007/s00213-006-0489-x.
URL http://www.ncbi.nlm.nih.gov/pubmed/16937099

[29] M. S. Matell, G. S. Portugal, Impulsive responding on the peak-interval procedure., Behavioural processes 74 (2) (2007) 198–208. doi:10.1016/j.beproc.2006.08.009.
URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1931419{&}tool=pmcentrez{&}rendertype=abstract

[30] F. Sanabria, E. A. Thrailkill, P. R. Killeen, Timing with opportunity cost: concurrent schedules of reinforcement improve peak timing., Learning & behavior 37 (3) (2009) 217–29. doi:10.3758/LB.37.3.217.

[31] J. Gibbon, R. M. Church, Comparison of variance and covariance patterns in parallel and serial theories of timing., Journal of the experimental analysis of behavior 57 (3) (1992) 393–406.
URL http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.

<div style="text-align: right">805</div>

```
fcgi?dbfrom=pubmed{&}id=1602270{&}retmode=ref{&}cmd=
prlinkspapers2://publication/doi/10.1901/jeab.1992.57-393
```

[32] J. L. Shapiro, J. Wearden, Reinforcement Learning and Time Perception – a Model of Animal Experiments, in: Advances in Neural Information Processing Systems, 2002, pp. 115–122.
URL `http://papers.nips.cc/paper/2000-reinforcement-learning-and-time-perception-a-model`

[33] P. R. Killeen, J. G. Fetterman, The behavioral theory of timing: transition analyses., Journal of the experimental analysis of behavior 59 (2) (1993) 411–22. `doi:10.1901/jeab.1993.59-411`.
URL `http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1322052{&}tool=pmcentrez{&}rendertype=abstract`

[34] M. Duflo, Random Iterative Models, Springer Science & Business Media, 2013.
URL `https://books.google.com/books?id=Kgf8CAAAQBAJ{&}pgis=1`

[35] K. Cheng, R. Westwood, Analysis of single trials in pigeons' timing performance., Journal of Experimental Psychology: Animal Behavior Processes 19 (1) (1993) 56–67. `doi:10.1037/0097-7403.19.1.56`.
URL `http://doi.apa.org/getdoi.cfm?doi=10.1037/0097-7403.19.1.56`

[36] J. Gibbon, R. M. Church, Sources of variance in an information processing theory of timing, in: H. L. Roitblat, H. S. Terrace, T. G. Bever (Eds.), Animal Cognition, Erlbaum, Hillsdale, NJ, 1984, Ch. 26, pp. 465–488.

[37] R. B. Ivry, J. E. Schlerf, Dedicated and intrinsic models of time perception, Trends in Cognitive Sciences 12 (7) (2008) 273–280. `doi:10.1016/j.tics.2008.04.002`.

[38] C. V. Buhusi, Dopaminergic Mechanisms of Interval Timing and Attention, in: W. H. Meck (Ed.), Functional and neural mechanisms of interval timining, CRC Press, Boca Raton, FL, US, 2003, pp. 317–338. `doi:10.1201/9780203009574.ch12`.

[39] P. Simen, F. Balci, L. Desouza, J. D. Cohen, P. Holmes, Interval timing by long-range temporal integration., Frontiers in integrative neuroscience 5 (2011) 28. `doi:10.3389/fnint.2011.00028`.
URL `http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3130150{&}tool=pmcentrez{&}rendertype=abstract`

[40] Y. Komura, R. Tamura, T. Uwano, H. Nishijo, K. Kaga, T. Ono, Retrospective and prospective coding for predicted reward in the sensory thalamus, Nature 412 (6846) (2001) 546–549.
URL `papers2://publication/uuid/0C735082-F2F4-4EEC-A87D-D09523893268`

[41] M. I. Leon, M. N. Shadlen, Representation of time by neurons in the posterior parietal cortex of the macaque., Neuron 38 (2) (2003) 317–327.
URL http://eutils.ncbi.nlm.nih.gov/entrez/eutils/
elink.fcgi?dbfrom=pubmed{&}id=12718864{&}retmode=
ref{&}cmd=prlinkspapers2://publication/uuid/
1247779A-368C-46DE-8CAE-E39A4A2E24A2

[42] J. Reutimann, V. Yakovlev, S. Fusi, W. Senn, Climbing neuronal activity as an event-based cortical representation of time., The Journal of neuroscience : the official journal of the Society for Neuroscience 24 (13) (2004) 3295–3303.
URL http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.
fcgi?dbfrom=pubmed{&}id=15056709{&}retmode=ref{&}cmd=
prlinkspapers2://publication/doi/10.1523/JNEUROSCI.4098-03.
2004

[43] G. Mello, S. Soares, J. Paton, A Scalable Population Code for Time in the Striatum, Current Biology 25 (9) (2015) 1113–1122.
doi:10.1016/j.cub.2015.02.036.
URL http://www.sciencedirect.com/science/article/pii/
S0960982215002055

[44] R. Bogacz, Optimal decision-making theories: linking neurobiology with behaviour., Trends in Cognitive Sciences 11 (3) (2007) 118–125.
URL http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.
fcgi?dbfrom=pubmed{&}id=17276130{&}retmode=ref{&}cmd=
prlinkspapers2://publication/doi/10.1016/j.tics.2006.12.006

[45] M. Carandini, D. J. Heeger, Normalization as a canonical neural computation., Nature reviews. Neuroscience 13 (1) (2012) 51–62.
doi:10.1038/nrn3136.
URL http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.
fcgi?dbfrom=pubmed{&}id=22108672{&}retmode=ref{&}cmd=
prlinkspapers2://publication/doi/10.1038/nrn3136http:
//www.pubmedcentral.nih.gov/articlerender.fcgi?artid=
3273486{&}tool=pmcentrez{&}rendertype=abstract

Appendix A. Variances and covariances

In the simplified TDDM, the random variables are the threshold Z_1 and the inverse of the slope $1/A$ (we show in S4 how to obtain the mean and variance of A from the distribution of $1/A$). Then the times of start, stop, middle and duration can all be written in terms of these random variables as follows:

26

$$S_1 = \frac{Z_1}{A} \qquad\qquad M = \frac{S_1 + S_2}{2} = \frac{Z_1 + z_2}{2A}$$

$$S_2 = \frac{z_2}{A} \qquad\qquad D = S_2 - S_1 = \frac{z_2 - Z_1}{A}$$

To derive the variances we use the relation $\mathrm{Var}(Y_1 Y_2) = \mathrm{E}(Y_1)^2 \cdot \mathrm{Var}(Y_2) + \mathrm{E}(Y_2)^2 \cdot \mathrm{Var}(Y_1) + \mathrm{Var}(Y_1) \cdot \mathrm{Var}(Y_2)$ for two independent random variables Y_1 and Y_2:

$$\mathrm{Var}(S_1) = \mathrm{Var}\left(Z_1 \cdot \frac{1}{A}\right)$$

$$= \mathrm{E}(Z_1)^2 \cdot \mathrm{Var}\left(\frac{1}{A}\right) + \mathrm{E}\left(\frac{1}{A}\right)^2 \cdot \mathrm{Var}(Z_1) + \mathrm{Var}(Z_1) \cdot \mathrm{Var}\left(\frac{1}{A}\right).$$

$$\mathrm{Var}(S_2) = \mathrm{Var}\left(z_2 \cdot \frac{1}{A}\right) = z_2^2 \cdot \mathrm{Var}\left(\frac{1}{A}\right).$$

$$\mathrm{Var}(D) = \mathrm{Var}\left(\frac{1}{A} \cdot (z_2 - Z_1)\right)$$

$$= \mathrm{Var}\left(\frac{1}{A}\right) \cdot \mathrm{Var}(Z_1) + \mathrm{E}\left(\frac{1}{A}\right)^2 \cdot \mathrm{Var}(Z_1) + (z_2 - \mathrm{E}(Z_1))^2 \cdot \mathrm{Var}\left(\frac{1}{A}\right).$$

$$\mathrm{Var}(M) = \frac{1}{4} \cdot \mathrm{Var}\left(\frac{1}{A} \cdot (Z_1 + z_2)\right)$$

$$= \frac{1}{4}\left[\mathrm{Var}\left(\frac{1}{A}\right) \cdot \mathrm{Var}(Z_1 + z_2) + \mathrm{E}\left(\frac{1}{A}\right)^2 \cdot \mathrm{Var}(Z_1 + z_2) + \mathrm{E}(Z_1 + z_2)^2 \cdot \mathrm{Var}\left(\frac{1}{A}\right)\right]$$

$$= \frac{1}{4}\left[\mathrm{Var}\left(\frac{1}{A}\right) \cdot \mathrm{Var}(Z_1) + \mathrm{E}\left(\frac{1}{A}\right)^2 \cdot \mathrm{Var}(Z_1) + (z_2 + \mathrm{E}(Z_1))^2 \cdot \mathrm{Var}\left(\frac{1}{A}\right)\right].$$

To find the covariances we use the definition $\mathrm{Cov}(Y_1, Y_2) = \mathrm{E}(Y_1 Y_2) -$

27

$\mathrm{E}(Y_1)\mathrm{E}(Y_2)$:

$$\mathrm{Cov}(S_1, S_2) = \mathrm{E}\left(\frac{Z_1}{A} \cdot \frac{z_2}{A}\right) - \mathrm{E}\left(\frac{Z_1}{A}\right) \cdot \mathrm{E}\left(\frac{z_2}{A}\right)$$

$$= z_2 \cdot \mathrm{E}\left(\frac{1}{A^2}\right) \cdot \mathrm{E}(Z_1) - z_2 \cdot \mathrm{E}\left(\frac{1}{A}\right)^2 \cdot \mathrm{E}(Z_1)$$

$$= z_2 \cdot \mathrm{E}(Z_1) \cdot \mathrm{Var}\left(\frac{1}{A}\right).$$

$$\mathrm{Cov}(S_1, D) = \mathrm{E}(S_1 \cdot (S_2 - S_1)) - \mathrm{E}(S_1) \cdot \mathrm{E}(S_2 - S_1)$$

$$= \mathrm{E}(S_1 \cdot S_2) - \mathrm{E}(S_1) \cdot \mathrm{E}(S_2) - \mathrm{E}(S_1^2) + \mathrm{E}(S_1)^2$$

$$= \mathrm{Cov}(S_1, S_2) - \mathrm{Var}(S_1).$$

$$\mathrm{Cov}(D, M) = \mathrm{E}\left(\frac{(S_2 - S_1) \cdot (S_1 + S_2)}{2}\right) - \mathrm{E}(S_2 - S_1) \cdot \mathrm{E}\left(\frac{S_1 + S_2}{2}\right)$$

$$= \frac{\mathrm{E}(S_2^2) - \mathrm{E}(S_1^2) - \mathrm{E}(S_2)^2 + \mathrm{E}(S_1)^2}{2}$$

$$= \frac{\mathrm{Var}(S_2) - \mathrm{Var}(S_1)}{2}.$$

Appendix B. Coefficients of variation

To derive the squared coefficients of variation we use the definition $\gamma_Y^2 = \mathrm{Var}(Y)/\mathrm{E}(Y)^2$ for a random variable Y:

$$\gamma_{S_1}^2 = \frac{\mathrm{Var}(S_1)}{\mathrm{E}(S_1)^2} = \frac{\mathrm{E}(Z_1)^2 \cdot \mathrm{Var}\left(\frac{1}{A}\right) + \mathrm{E}(\frac{1}{A})^2 \cdot \mathrm{Var}(Z_1) + \mathrm{Var}(Z_1) \cdot \mathrm{Var}\left(\frac{1}{A}\right)}{\mathrm{E}(Z_1)^2 \cdot \mathrm{E}\left(\frac{1}{A}\right)^2}$$

$$= \gamma_{1/A}^2 + \gamma_{Z_1}^2 \cdot \gamma_{1/A}^2 + \gamma_{Z_1}^2.$$

$$\gamma_{S_2}^2 = \frac{\mathrm{Var}(S_2)}{\mathrm{E}(S_2)^2} = \frac{z_2^2 \cdot \mathrm{Var}\left(\frac{1}{A}\right)}{z_2^2 \cdot \mathrm{E}\left(\frac{1}{A}\right)^2} = \gamma_{1/A}^2.$$

$$\gamma_D^2 = \frac{\mathrm{Var}\left(\frac{1}{A}\right) \cdot \mathrm{Var}(Z_1) + \mathrm{E}\left(\frac{1}{A}\right)^2 \cdot \mathrm{Var}(Z_1) + (z_2 - \mathrm{E}(Z_1))^2 \cdot \mathrm{Var}\left(\frac{1}{A}\right)}{\mathrm{E}(\frac{z_2 - Z_1}{A})^2}$$

$$= \frac{\mathrm{Var}\left(\frac{1}{A}\right) \cdot \mathrm{Var}(Z_1) + \mathrm{E}\left(\frac{1}{A}\right)^2 \cdot \mathrm{Var}(Z_1) + (z_2 - \mathrm{E}(Z_1))^2 \cdot \mathrm{Var}\left(\frac{1}{A}\right)}{\mathrm{E}\left(\frac{1}{A}\right)^2 \cdot (z_2 - \mathrm{E}(Z_1))^2}$$

$$= \frac{\gamma_{Z_1}^2(1 + \gamma_{1/A}^2)}{(\theta - 1)^2} + \gamma_{1/A}^2.$$

$$\gamma_M^2 = \frac{\frac{1}{4}[\mathrm{Var}\left(\frac{1}{A}\right) \cdot \mathrm{Var}(Z_1) + \mathrm{E}\left(\frac{1}{A}\right)^2 \cdot \mathrm{Var}(Z_1) + (z_2 + \mathrm{E}(Z_1))^2 \cdot \mathrm{Var}\left(\frac{1}{A}\right)]}{\frac{1}{4}\mathrm{E}(\frac{Z_1 + z_2}{A})^2}$$

$$= \frac{\mathrm{Var}\left(\frac{1}{A}\right) \cdot \mathrm{Var}(Z_1) + \mathrm{E}\left(\frac{1}{A}\right)^2 \cdot \mathrm{Var}(Z_1) + (z_2 + \mathrm{E}(Z_1))^2 \cdot \mathrm{Var}\left(\frac{1}{A}\right)}{\mathrm{E}\left(\frac{1}{A}\right)^2 \cdot (z_2 + \mathrm{E}(Z_1))^2}$$

$$= \frac{\gamma_{Z_1}^2(1 + \gamma_{1/A}^2)}{(\theta + 1)^2} + \gamma_{1/A}^2.$$

Appendix C. Correlations

To find the correlations we use the definition $\rho(Y_1, Y_2) = \frac{\text{Cov}(Y_1, Y_2)}{\sqrt{\text{Var}(Y_1)}\sqrt{\text{Var}(Y_2)}}$:

$$\rho(S_1, S_2) = \frac{\text{Cov}(S_1, S_2)}{\sqrt{\text{Var}(S_1)}\sqrt{\text{Var}(S_2)}}$$

$$= \frac{\text{E}(Z_1)\sqrt{\text{Var}\left(\frac{1}{A}\right)}}{\sqrt{\text{Var}\left(\frac{1}{A}\right) \cdot \text{Var}(Z_1) + \text{Var}(Z_1)\text{E}\left(\frac{1}{A}\right)^2 + \text{E}(Z_1)^2 \cdot \text{Var}\left(\frac{1}{A}\right)}}$$

$$= \sqrt{\frac{\text{E}(Z_1)^2 \cdot \text{Var}\left(\frac{1}{A}\right)}{\text{Var}\left(\frac{1}{A}\right) \cdot \text{Var}(Z_1) + \text{Var}(Z_1)\text{E}\left(\frac{1}{A}\right)^2 + \text{E}(Z_1)^2 \cdot \text{Var}\left(\frac{1}{A}\right)}}$$

$$= \frac{1}{\sqrt{\gamma_{Z_1}^2 + \frac{\gamma_{Z_1}^2}{\cdot\gamma_{1/A}^2} + 1}}.$$

For $\rho(S_1, D) = \frac{\text{Cov}(S_1, D)}{\sqrt{\text{Var}(S_1)}\sqrt{\text{Var}(D)}}$ we proceed first by dividing both numerator and denominator by $\text{E}(Z_1)^2 \cdot \text{E}\left(\frac{1}{A}\right)^2$:

$$\frac{\text{Cov}(S_1, D)}{\text{E}(Z_1)^2 \cdot \text{E}\left(\frac{1}{A}\right)^2} = \frac{z_2\text{E}(Z_1)\text{Var}\left(\frac{1}{A}\right) - \text{E}(Z_1)^2\text{Var}\left(\frac{1}{A}\right) - \text{E}\left(\frac{1}{A}\right)^2\text{Var}(Z_1) - \text{Var}(Z_1)\text{Var}\left(\frac{1}{A}\right)}{\text{E}(Z_1)^2 \cdot \text{E}\left(\frac{1}{A}\right)^2}$$

$$= \gamma_{1/A}^2(\theta - \gamma_{Z_1}^2 - 1) - \gamma_{Z_1}^2.$$

$$\frac{\sqrt{\text{Var}(S_1)}\sqrt{\text{Var}(D)}}{\text{E}(Z_1)^2 \cdot \text{E}\left(\frac{1}{A}\right)^2} = \sqrt{\frac{\text{Var}(S_1)}{\text{E}(Z_1)^2 \cdot \text{E}\left(\frac{1}{A}\right)^2}} \cdot \sqrt{\frac{\text{Var}(D)}{\text{E}(Z_1)^2 \cdot \text{E}\left(\frac{1}{A}\right)^2}}$$

$$= \sqrt{\gamma_{Z_1}^2 + \gamma_{1/A}^2\gamma_{Z_1}^2 + \gamma_{1/A}^2} \cdot \sqrt{\gamma_{1/A}^2(\theta - 1)^2 + \gamma_{1/A}^2\gamma_{Z_1}^2 + \gamma_{Z_1}^2}.$$

Putting everything back together we have

$$\rho(S_1, D) = \frac{\gamma_{1/A}^2(\theta - 1 - \gamma_{Z_1}^2) - \gamma_{Z_1}^2}{\sqrt{\gamma_{1/A}^2(\gamma_{Z_1}^2 + 1) + \gamma_{Z_1}^2} \cdot \sqrt{\gamma_{1/A}^2[(\theta - 1)^2 + \gamma_{Z_1}^2] + \gamma_{Z_1}^2}}.$$

Proceeding in the same way as above for $\rho(D, M) = \frac{\text{Cov}(D, M)}{\sqrt{\text{Var}(D)}\sqrt{\text{Var}(M)}}$ we derive

$$\rho(D, M) = \frac{\gamma_{1/A}^2(\theta^2 - 1 - \gamma_{Z_1}^2) - \gamma_{Z_1}^2}{\sqrt{\gamma_{1/A}^2[(\theta + 1)^2 + \gamma_{Z_1}^2] + \gamma_{Z_1}^2} \cdot \sqrt{\gamma_{1/A}^2[(\theta - 1)^2 + \gamma_{Z_1}^2] + \gamma_{Z_1}^2}}.$$

Appendix D. Note on the reciprocal normal distribution $1/A$

We make the following claim: the combined effect of the Gaussian noise in the TDDM and the adaptation rules cause the slope A to be a normally distributed random variable. We justify this claim on the basis that, because of noise, the slope never converges absolutely to one value that is used on every trial. Instead, it fluctuates around an average value, being corrected up when the process crosses the threshold too late, or down when it crosses too early.

We have shown in the paper that it is possible to use the data to estimate the coefficient of variation of the reciprocal of the slope $1/A$, namely $\gamma_{1/A} = \gamma_{S_2}$ (equation (14) in the main text). However, to run the simulations we need estimates of the mean and variance of A. The question is how to obtain such estimates from the reciprocal distribution $1/A$.

If, as we claim, A is normal, then the distribution of $1/A$ is undefined since

$$\mathrm{E}[1/A] = \frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^{\infty} a^{-1} e^{\frac{-(a-\mu)^2}{2\sigma^2}}\, da$$

does not converge. A possible solution is to estimate using Taylor series. Let A be a random variable with $\mathrm{E}[A] = \mu$. Estimating $g(\mu)$ we get:

$$g(A) = g(\mu) + g'(\mu)(A - \mu) + \ldots$$

Then we can say that approximately

$$\mathrm{E}[g(A)] \approx g(\mu),$$
$$\mathrm{Var}[g(A)] \approx g'(\mu)^2 \mathrm{Var}[A].$$

Finally, if $g(A) = 1/A$ then

$$\mathrm{E}\left[\frac{1}{A}\right] \approx \frac{1}{\mu},$$
$$\mathrm{Var}\left[\frac{1}{A}\right] \approx \left(\frac{1}{\mu}\right)^4 \mathrm{Var}[A].$$

To generate the parameter A for the simulation we sample from a normal distribution with

$$\mathrm{E}[A] = \mu,$$
$$\mathrm{Var}[A] = \mu^4 \mathrm{Var}\left[\frac{1}{A}\right] = \mu^2 \gamma_{1/A}^2$$

since $\mathrm{Var}[1/A] = \mathrm{E}[1/A]^2 \gamma_{1/A}^2 = \mu^{-2}\gamma_{1/A}^2$.

31

Table 1: Review of all the studies with published data on individual trial analysis in the peak procedure. Results are separated by FI duration when possible; otherwise averages are reported [22, 25]. [24] has separate conditions with low (L) and high (R) rewards. Note that in some studies one or more correlations were not reported and that [22] and [25] did not report CVs (marked as n/a).

Experiment	FI	Species	$\rho(S_1, S_2)$	$\rho(S_1, D)$	$\rho(D, M)$	CV_{start}	CV_{stop}	CV_{mid}	CV_{dur}
[22]	mix	pigeon	0.34	-0.34	0.38	n/a	n/a	n/a	n/a
[26]	12.5	pigeon	0.38	-0.36	0.31	0.41	0.17	0.19	0.24
[12]	15	rat	0.32	-0.28	0.51	0.55	0.20	0.22	0.26
[12]	30	rat	0.32	-0.42	0.27	0.55	0.18	0.20	0.25
[12]	60	rat	0.30	-0.31	0.46	0.45	0.22	0.22	0.31
[27]	8	human	0.78	-0.26	0.04	0.14	0.11	0.12	0.24
[27]	12	human	0.91	-0.02	0.20	0.15	0.12	0.13	0.20
[27]	21	human	0.91	-0.10	0.10	0.15	0.13	0.14	0.26
[23]	5	mouse	0.45	-0.14	0.55	0.48	0.16	0.18	0.16
[23]	15	mouse	0.40	n/a	0.43	0.71	0.16	0.20	0.19
[23]	45	mouse	0.49	-0.49	0.01	0.46	0.17	0.21	0.28
[28]	30	rat	n/a	n/a	n/a	0.17	0.12	0.11	0.22
[29]	15	rat	0.22	-0.51	0.19	0.35	0.14	0.15	0.26
[25]	mix	rat	0.63	-0.43	n/a	n/a	n/a	n/a	n/a
[30]	15A	pigeon	0.51	-0.55	n/a	0.65	0.25	0.32	0.46
[30]	15B	pigeon	0.36	-0.51	n/a	0.49	0.22	0.25	0.42
[30]	60	pigeon	0.50	-0.41	n/a	0.66	0.30	0.36	0.51
[24]	10L	human	0.79	-0.15	0.17	0.10	0.09	0.09	0.26
[24]	10H	human	0.66	-0.18	0.25	0.10	0.10	0.09	0.39
[24]	17L	human	0.79	0.01	0.34	0.10	0.10	0.09	0.33
[24]	17H	human	0.63	-0.20	0.19	0.12	0.11	0.10	0.48

Table 2: Correlations for all experiments. Each experiment's correlation is followed by the approximation made using the CVs by the simplified model and then by the full model simulation results (α optimized). Notice that [23] with FI 5 and [24] with FI 17L cannot be fitted as their CVs do not follow the model's constraints. Correlations not following the general trend are marked by an *.

Experiment	FI	$\rho(S_1,S_2)$	Simp.	Full	$\rho(S_1,D)$	Simp.	Full	$\rho(D,M)$	Simp.	Full
[26]	12.5	0.38	0.41	0.42	-0.36	-0.35	-0.36	0.31	0.30	0.28
[12]	15	0.32	0.36	0.39	-0.28	-0.32	-0.28	0.51	0.40	0.42
[12]	30	0.32	0.33	0.34	-0.42	-0.42	-0.42	0.27	0.28	0.27
[12]	60	0.30	0.49	0.47	-0.31	-0.27	-0.31	0.46	0.34	0.30
[27]	8	0.78	0.79	0.75	-0.26	-0.19	-0.28	0.04	0.16	0.08
[27]	12	0.91	0.80	0.78	-0.02	-0.00	-0.08	0.20	0.35	0.29
[27]	21	0.91	0.87	0.83	-0.10	0.00*	-0.13	0.10	0.28	0.19
[23]	5	0.45	—	—	-0.14	—	—	0.55	—	—
[23]	15	0.40	0.23	0.19	n/a	-0.34	-0.39	0.43	0.52	0.47
[23]	45	0.49	0.37	0.35	-0.49	-0.51	-0.55	0.01	0.08	0.04
[28]	30	n/a	0.71	0.68	n/a	-0.21	-0.26	n/a	0.21	0.17
[29]	15	0.22	0.40	0.44	-0.51	-0.56	-0.48	0.19	-0.02*	0.08
[30]	15A	0.51	0.38	0.42	-0.55	-0.57	-0.51	n/a	-0.02*	0.05
[30]	15B	0.36	0.45	0.44	-0.51	-0.53	-0.54	n/a	-0.00*	-0.02*
[30]	60	0.50	0.45	0.49	-0.41	-0.45	-0.40	n/a	0.11	0.16
[24]	10L	0.79	0.90	0.89	-0.15	-0.10	-0.14	0.17	0.13	0.10
[24]	10H	0.66	0.97	0.96	-0.18	0.03*	-0.08	0.25	0.15	0.06
[24]	17L	0.79	—	—	0.01*	—	—	0.34	—	—
[24]	17H	0.63	0.92	0.91	-0.20	-0.18	-0.20	0.19	0.03	0.01

Table 3: Equivalence between the main variables in the simplified TDDM and 2-threshold version of SET ([12]). Note that the simplified TDDM uses rate instead of time ($A = 1/T$). T is determined by the task and is only necessary to reproduce the response curve, but is not needed for calculating CVs and the correlations analysis because of timescale invariance. Finally, SET allows for two independent noisy thresholds whereas, in the TDDM, the expected values of both thresholds are fixed from θ and the variance of the second threshold is 0. The last column gives the parameters absolute difference averaged across all 17 experiments the simplified model could fit.

SET	Simplified TDDM	Avg. Abs. Diff.
$E(n^*) = T^*$	$E(A_n) = 1/T$	0.00
γ	$\gamma_{1/A}$	0.03
$E(l) = E(1 - b_1)$	$E(Z_1) = 2/(1 + \theta)$	0.10
σ_{b_1}	σ_{Z_1}	0.03
$E(u) = E(1 + b_2)$	$z_2 = 2\theta/(1 + \theta)$	0.19
σ_{b_2}	0	0.04