Verdonck, S., Loossens, T. and Philiastides, M. G. (2020) The Leaky Integrating Threshold and its impact on evidence accumulation models of choice RT. *Psychological Review*, (doi: 10.1037/rev0000258).

http://eprints.gla.ac.uk/220386/

Deposited on 13 July 2020

The Leaky Integrating Threshold and its impact on evidence accumulation models of choice RT

Stijn Verdonck(`stijn.verdonck@kuleuven.be`)
KU Leuven

Tim Loossens
KU Leuven

Marios G. Philiastides
Institute of Neuroscience and Psychology, University of Glasgow

Abstract

A common assumption in choice response time modelling is that after evidence accumulation reaches a certain decision threshold, the choice is categorically communicated to the motor system that then executes the response. However, neurophysiological findings suggest that motor preparation partly overlaps with evidence accumulation, and is not independent from stimulus difficulty level. We propose to model this entanglement by changing the nature of the decision criterion from a simple threshold to an actual process. More specifically, we propose a secondary, motor preparation related, leaky accumulation process that takes the accumulated evidence of the original decision process as a continuous input, and triggers the actual response when it reaches its own threshold. We analytically develop this Leaky Integrating Threshold (LIT), applying it to a simple constant drift diffusion model, and show how its parameters can be estimated with the D*M method. Reanalyzing three different datasets, the LIT extension is shown to outperform a standard drift diffusion model using multiple statistical approaches. Further, the LIT leak parameter is shown to be better at explaining the speed/accuracy trade-off manipulation than the commonly used boundary separation parameter. These improvements can also be verified using traditional diffusion model analyses, for which the LIT predicts the violation of several common selective parameter influence assumptions. These predictions are consistent with what is found in the data and with what is reported experimentally in the literature. Crucially, this work offers a new benchmark against which to compare neural data to offer neurobiological validation for the proposed processes.

Keywords: LIT, choice RT, diffusion model, motor preparation, threshold, boundary

The study of the neurobiological and computational underpinnings of perceptual decision making has received considerable attention across species and levels of description (Gold and Heekeren, 2013; Gold and Shadlen, 2007; Heekeren et al., 2008; Sajda et al., 2009). For example, several non-human primate electrophysiology studies have revealed patterns of single-unit activity that exhibit an integrative decision mechanism in areas of the parietal and prefrontal cortex (Ding and Gold, 2012; Kim and Shadlen, 1999; Meister et al., 2013; Shadlen and Newsome, 2001). Specifically, firing rates of a subset of neurons in these areas build up over time with a rate proportional to the amount of sensory evidence (i.e., difficulty of the task) and eventually converge to a common firing level (decision boundary) as animals commit to a choice. In humans, macroscopic measurements of neural activity using M/EEG (O'Connell et al., 2012; Philiastides et al., 2014a; Pisauro et al., 2017; Polania et al., 2014; Ratcliff et al., 2009; Wyart et al., 2012) and fMRI experiments (Heekeren et al., 2004; Kayser et al., 2010; Noppeney et al., 2010; Ploran et al., 2011) have revealed comparable activation patterns in similar brain regions.

The longstanding assumption in this accumulation-to-bound framework has been that the motor system is engaged to indicate the eventual choice, only after an internal decision threshold is reached. More recent evidence, however, appears to contradict this strict temporal dichotomy between decisional and motor processes and suggests that the motor system is already engaged during the decision making, not just after. Research using electromyography of hand muscles, identified partially executed (covert) responses when different decision alternatives competed for representation (McBride et al., 2018; Spieser et al., 2017). These covert responses (Rochet et al., 2014), suggest that activity in the motor system is being updated continuously as the decision process unfolds (Servant et al., 2015). In line with interpretation, recent work on volition argued that internally generated actions depend on the leaky integration of endogenous neuronal noise (Haggard, 2019; Schurger et al., 2012).

Similarly, response-selective activity in MEG studies (extracted by exploiting the well-known contralateral motor bias) was shown to exhibit choice-predictive activity that built up gradually during a direction discrimination task and reflected the temporal integral of the activity arriving from motion-selective sensory areas (Donner et al., 2009,0). Separate fMRI and transcranial magnetic stimulation work offered additional support that brain areas involved in evidence accumulation are tightly coupled with the motor system, such that evolving decisions are also reflected in motor activity when the relevant decisions are mapped onto actions (Filimon et al., 2013; Klein-Flugge and Bestmann, 2012). Intriguingly, recent multi-unit recordings in animals demonstrated that while neural representations in sensory and association cortices were sufficient to perform a simple discrimination task, inactivation of a downstream pre-motor area, led to gross behavioral impairments (Wu et al., 2019).

Despite this evidence, the precise functional role of the motor system in decision-making remains unclear. More specifically, it remains unknown whether the motor apparatus merely receives "echoes" of the relevant decision variables or is more actively involved in driving the decision itself. Notwithstanding, these recent findings suggest that neural activity related

---

to motor preparation, starts its buildup well before the evidence accumulation process completes, effectively lagging the primary process of evidence accumulation. Consequently, the level of stimulus difficulty could have a direct impact on motor preparation, which in turn could influence the eventual choice. This entanglement (and its potential implications) is not accounted for in the current generation of computational models that describe choice RT data with noisy evidence accumulators (Palmer et al., 2005; Ratcliff, 1978; Usher and McClelland, 2001).

Customarily, evidence accumulation models require a criterion to decide when enough evidence has been accumulated to warrant a choice. Although some alternatives for the commonly used fixed decision boundary have been proposed (e.g., collapsing bounds, see Hawkins et al., 2015), the relation between the decision making process on the one hand and the motor execution process on the other, is always assumed to be limited to a mere choice relay, once the decision has been formed. In this paper, we model the entanglement between the evidence accumulation and motor preparation processes with a process boundary. This process boundary is formalized as a second leaky accumulation process (presumably implemented in the motor system) that takes the accumulated evidence of the decision process as a continuous input and triggers the actual physical response when it reaches a certain threshold of its own. In other words, the evidence accumulator relinquishes control of the eventual choice by passing the integrated evidence along to the motor system. In that sense, it is the dynamics observed in the latter that dictate the choice.

We first implement this process boundary, which we will call the Leaky Integrating Threshold (LIT), for the widely used one-dimensional constant drift diffusion model (DDM). We investigate this special case analytically and show what the implications of such a process boundary are, with special attention for the speed/accuracy trade-off (SAT) manipulation. Conceptually, the LIT framework enables a more intricate interplay between decision and motor processes and by extension could offer a more flexible account of how perceptual decisions, in general, are being implemented. We chose to put a special emphasis on the SAT primarily as a means of validating the model. More specifically, the SAT offers an intuitive opportunity for a robust manipulation of the "leak" parameter in the model and by extension it enables a direct comparison of the LIT with other incarnations of the more conventional DDM.

We use two different approaches to examining the data: the first uses a recently developed simulation based method of statistical inference called prepaid (Mestdagh et al., 2019) to directly fit the LIT to data; the second uses an approximation of the LIT that can be tackled with simpler numerical tools. We use both methods to analyse a diverse collection of datasets and compare the LIT to existing models. In both approaches, we first test if a leak parameter adds to a simple DDM description of choice RT data, in both the speed and the accuracy condition. We then test if the leak parameter constitutes a good alternative for boundary separation to model SAT. The second approach allows us to compare the leak parameter with the inter-trial variability parameters as defined in the popular Ratcliff extension (RDM). Depending on what is possible in each approach, we also investigate the selective influence of typical experimental manipulations on model parameters. This is a critical step since task manipulations, such as difficulty and SAT, should each only impact a designated parameter in order for the model to be easily interpretable.

Finally, we briefly investigate the impact of the LIT on more complex 2D models of choice

RT, more specifically on the Leaky Competing Accumulator (LCA) (Usher and McClelland, 2001) and the Ising Decision Maker (IDM) (Verdonck and Tuerlinckx, 2014).

The modeling proposed in this paper builds on the vast tradition of diffusion models for choice RT, as they are used in the fields of psychology and neuroscience. Although we drew inspiration from from both domains for the formulation of the model, in terms of experimental validation, this paper focuses purely on behavioral choice RT data. A complete assessment of the model will require independent neurobiological validation, ideally at different levels of abstraction.

1   The Leaky Integrating Threshold

The dynamics of the deterministic Leaky Integrating Threshold are formalized as

$$dy(t) = (\beta x(t) - \lambda y(t))dt, \tag{1}$$

with $x(t)$ the evidence buildup over time, $y(t)$ the resulting LIT buildup over time, λ the leak that determines their lag, and β a scaling parameter for the input. When this LIT process hits some boundary, the physical response ensues and a choice is conveyed. The LIT boundaries are set up similarly to the stopping criterion of the evidence accumulation process $x(t)$ under consideration. In the case of typical one-dimensional evidence accumulators like the constant drift diffusion model, both upper and lower boundaries are used, each representing one of the two possible choices. This means a decision is made once the process exceeds the upper or drops below the lower decision boundary. In the case of two-dimensional evidence accumulators like the LCA or the IDM, each dimension of evidence has their own LIT with a single upper boundary. In this case, whichever of the two dimension boundaries is crossed first, determines the choice. In this paper, we deliberately choose not to add Wiener noise to the LIT equation, to keep the mathematics of this additional level as simple as possible. In contrast, we fully respect the stochastic nature of the evidence accumulation process $x(t)$, most importantly when calculating its impact on the LIT signal $y(t)$. The solution for $y(t)$ in Equation 1 is

$$y(t) = \beta \int_{-\infty}^{t} dt' x(t') e^{\lambda(t'-t)}. \tag{2}$$

Indeed, substituting this solution in Equation 1, gives

$$
\begin{aligned}
\frac{dy(t)}{dt} &= \beta \frac{d}{dt} \left(e^{-\lambda t} \int_{-\infty}^{t} dt' x(t') e^{\lambda t'} \right) \\
&= \beta \left(-\lambda e^{-\lambda t} \int_{-\infty}^{t} dt' x(t') e^{\lambda t'} + e^{-\lambda t} x(t) e^{\lambda t} \right) \\
&= -\lambda y(t) + \beta x(t).
\end{aligned}
$$

As can be seen from Equation 2, the effect of the Leaky Integrating Threshold is that of a smoothing filter. It takes the average of all past states of x(t), but with exponentially decaying weights. Because this averaging flattens out random noise, the signal-to-noise ratio of the underlying process is increased.

Rescaling y and boundary a with the same factor $\frac{\lambda}{\beta}$,

$$y' = \frac{\lambda}{\beta} y$$

$$a' = \frac{\lambda}{\beta} a,$$

does not change first passage times nor choice RT outcomes. The equation can now be written as

$$\frac{\beta}{\lambda} \frac{dy'(t)}{dt} = -\beta y'(t) + \beta x(t)$$

$$\frac{dy'(t)}{dt} = -\lambda y'(t) + \lambda x(t),$$

with solution

$$y'(t) = \lambda \int_{-\infty}^{t} dt' x(t') e^{\lambda(t'-t)}. \tag{3}$$

This rescaling removes a redundant parameter by equating $\lambda = \beta$.

1.1   Applied to the constant drift diffusion model

In the following paragraphs, we apply the LIT to the constant drift diffusion model. We start out with a simpler problem – describing the situation without absorbing boundaries. The constant drift diffusion model is described by the following stochastic differential equation:

$$dx(t) = v dt + \sigma dW(t),$$

with v the drift speed, σ the diffusion rate and $W(t)$ a Wiener process. The solution to this equation is

$$x(t) = x_0 + vt + \sigma W_t, \tag{4}$$

with x_0 the starting position. Now, let's insert the diffusion process of Equation 4 into the LIT solution in Equation 2. Assuming that drift speed v and noise σ are zero before $t = 0$,

$$
\begin{aligned}
y(t) &= \lambda \int_{-\infty}^{t} dt' (x_0 + vt' + \sigma W_{t'}) e^{\lambda(t'-t)} \\
&= \lambda x_0 \underbrace{\int_{-\infty}^{t} dt' e^{\lambda(t'-t)}}_{A} + \lambda v \underbrace{\int_{0}^{t} dt' t' e^{\lambda(t'-t)}}_{B} \\
&\quad + \lambda \sigma \underbrace{\int_{0}^{t} dt' e^{\lambda(t'-t)} W_{t'}}_{C}.
\end{aligned}
\tag{5}
$$

Inserting the results from integrals A, B and C (for more details, see Appendix I), we can say $y(t)$ constitutes a stochastic process that is normally distributed at every time point t, with the mean evolving as

$$\langle y(t) \rangle = x_0 + v(t - \frac{1}{\lambda} + \frac{e^{-\lambda t}}{\lambda}),$$

and the variance as

$$\left\langle (y(t) - \langle y(t) \rangle)^2 \right\rangle = \sigma^2 \left(t - \frac{3}{2\lambda} + \frac{4e^{-\lambda t} - e^{-2\lambda t}}{2\lambda} \right).$$

For sufficiently large t, the mean and variance are respectively given by

$$\left\langle (y(t) - \langle y(t) \rangle)^2 \right\rangle = \sigma^2 \left(t - \frac{3}{2\lambda} \right),$$

$$\langle y(t) \rangle = x_0 + v(t - \frac{1}{\lambda})$$

$$= x_0 + v \left(t - (\frac{3}{2\lambda} - \frac{1}{2\lambda}) \right)$$

$$= (x_0 + \frac{v}{2\lambda}) + v(t - \frac{3}{2\lambda}). \tag{6}$$

This is the same distributional evolution as that of a simple constant drift diffusion process with the same drift speed v and noise σ, a stimulus dependent bias $y_0 = x_0 + \frac{v}{2\lambda}$ towards the correct choice and shifted $\frac{3}{2\lambda}$ time units backward in time. The origin of the stimulus dependent bias can be intuitively understood as follows. As can be seen from Equation 6, the LIT signal $y(t)$ lags the original evidence accumulation signal $x(t) = x_0 + vt + \sigma W_t$. This lag, however, is not the same for the expected mean and variance of those processes. As is apparent from Equation 6, the LIT signal mean is delayed by $\frac{1}{\lambda}$, while its variance is delayed by the slightly larger amount $\frac{3}{2\lambda}$. Because the variance only starts growing later, the first part of the evolution can be seen as deterministic, and its effect can be reduced to a bias proportional to the drift speed. At least, the LIT process at larger t behaves as if this would have happened: Equation 6 does not offer much insight into what really occurs at small t in the LIT, as the approximation is only valid for large t. This is something that is best assessed through simulations. Although mathematically presented as a bias of a DDM, it is important to stress that this stimulus dependent bias has no valid interpretation in that context. In a traditional interpretation of any diffusion model of choice RT, a bias can never be stimulus dependent, because this would mean that the initial value of the accumulated evidence contains information about the stimulus, even before it had the chance to convey it. In the same vein, a bias can never point towards the correct answer: it should always point to one of the two experimental options, regardless of the stimulus.

Now, let's consider the case with absorbing boundaries. Because the resulting signal $y(t)$ is one-dimensional like $x(t)$, two boundaries are required. Throughout this paper, boundaries on one-dimension models — on the level of the LIT or otherwise — always represent choice options, not correct or error responses (it is therefore perfectly possible to encounter negative drift speeds). As we defined the LIT to evolve to zero without any input, the two boundaries are at $\frac{a}{2}$ and $-\frac{a}{2}$, with a the boundary separation. The noise that ends up driving $y(t)$ in Equation 5 has become auto-correlated, or so-called colored noise, and the first passage times cannot be seen as those of a simple constant diffusion model with Wiener
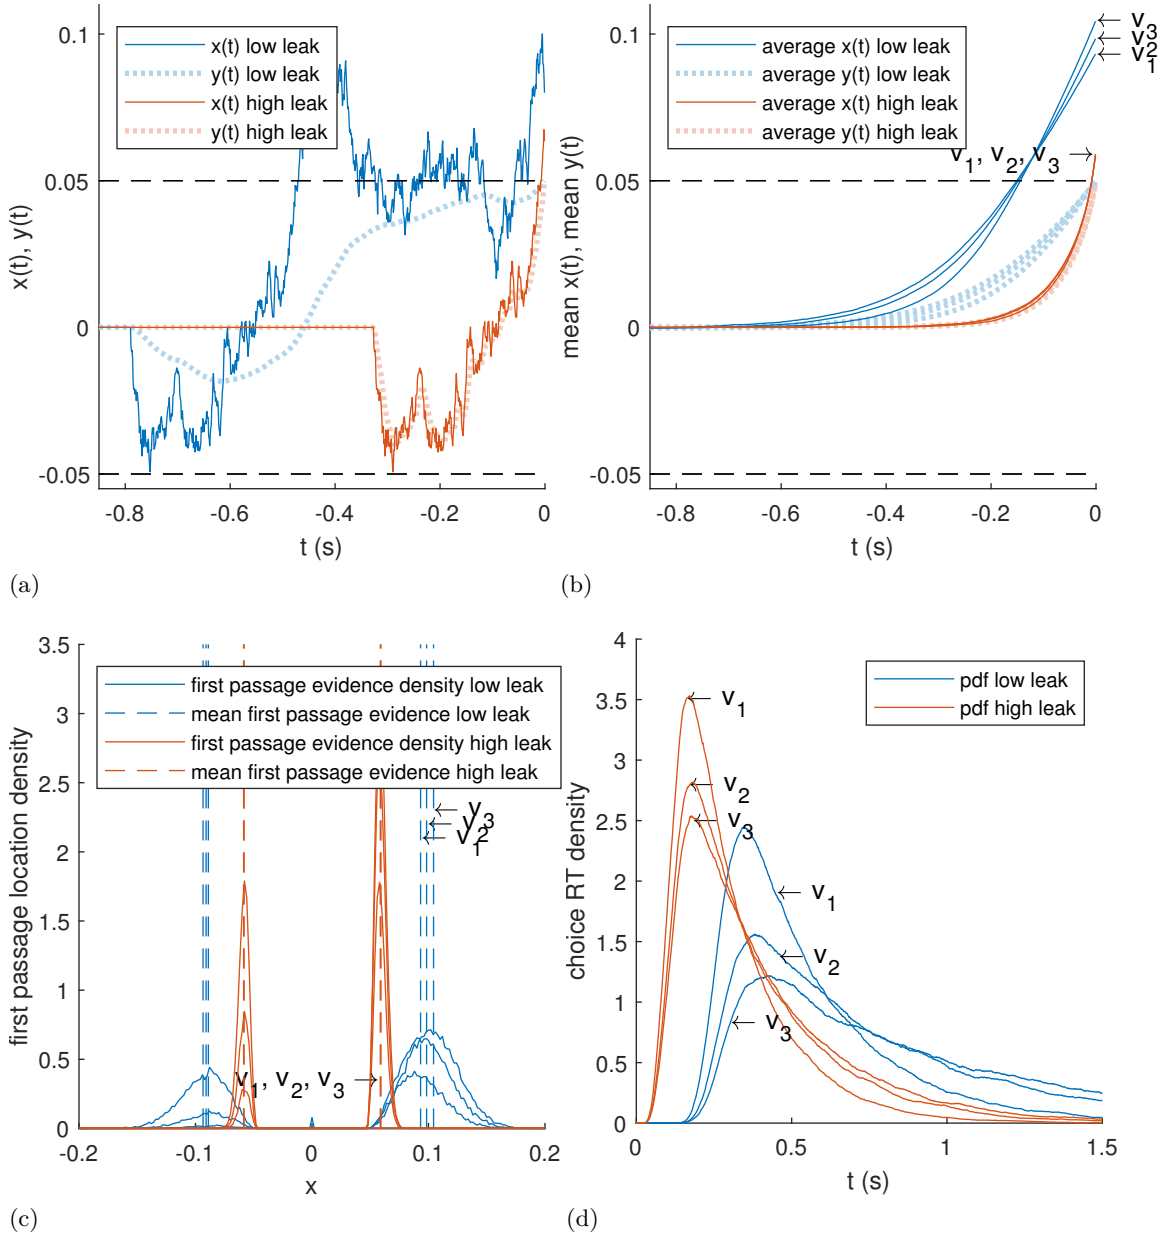
Figure 1. Evidence integration and lagged LIT buildup for a constant drift diffusion model with a Leaky Integrating Threshold. Plot (a) shows, response locked, the noisy evidence accumulation signal $x(t)$ (solid) and accompanying lagged and smoothed LIT signal $y(t)$ (dotted) for a trial with a high leak ($\lambda = 100$, red) and a trial with a low leak ($\lambda = 5$, blue). Plot (b) shows the average evidence (solid) and LIT signals (dotted) for the LIT hitting the upper boundary, again response locked, for the same two values of λ, but both for three different levels of stimulus difficulty ($v_1 = 0, v_2 = 0.1, v_3 = 0.2$). In plot (c), the distributions of the accumulated evidence $x(t_r)$ taken at the moment the first boundary on y is reached, i.e. $|y(t_r)| = \frac{a}{2}$, are shown. All previous combinations of leak and stimulus difficulty are considered. The means of these distributions are indicated with dashed vertical lines. Finally, plot (d) shows the respective choice RT probability density functions. For an overview of the parameters that are kept constant, see Table 2 in Appendix II.

noise. In Figure 1 we show both processes $x(t)$ and $y(t)$, for two values of λ. Figure 1a shows how $y(t)$ is a smoothed and lagged version of $x(t)$, resulting in an increased signal-to-noise ratio. The lower the λ ($\lambda = 5$ is in blue, $\lambda = 100$ in red), the more pronounced the effect. Plot (b) shows the resulting mean signals. For $\lambda = 100$, $y(t)$ tracks $x(t)$ closely (red lines). Plot c) shows the accompanying distributions of the accumulated evidence $x(t_r)$ taken at the moment t_r the first boundary on y is reached, i.e. $|y(t_r)| = \frac{a}{2}$. They take the form of two clearly defined boundaries on the level of x as well, much like would be the case for a normal diffusion model, without a LIT. For $\lambda = 5$ however, the smoothing and lag are substantial. The resulting $x(t_r)$ distributions (blue) have more variance, because there is a wider range of values the accumulated evidence can take at the moment of boundary crossing playing out on the level of y. Moreover, the $x(t_r)$ distributions are pushed outwards because a heavily smoothed process $y(t)$ is typically closer to the starting point than the underlying evidence accumulation process $x(t)$, so it takes the $y(t)$ signal longer to reach one of the boundaries. Although the boundaries at $\frac{a}{2}$ and $-\frac{a}{2}$ are imposed on the y level, $x(t_r)$ can be seen as a resulting effective decision boundary on the x level, leading to both slower and more accurate choices when they are further away from the starting point. In that sense, the LIT leak parameter λ is an interesting candidate for speed/accuracy trade-off regulation. The increase in accuracy obtained with lower values of λ, seems to coincide with some additional effective non-decision time, as can be seen in plot (d) from the shift of the onset of the low leak (blue, $\lambda = 5$) choice RT distributions. This is in line with what could be expected from the time shift appearing in Equation 6. Another feature of the LIT that deserves to be highlighted becomes apparent in plot (c) when looking at the means of the effective decision boundary distributions $x(t_r)$ (indicated with vertical dashed lines). In the low leak (blue, $\lambda = 5$) case, the effective decision bound is not the same for all three stimuli under consideration: the easier the stimulus, the further from the starting point the effective decision boundary distribution corresponding to the correct choice. Although there is still a common motor preparation bound for all stimuli at the level of the LIT signal $y(t)$, this result challenges the idea of a common decision bound for all stimuli at the level of the original evidence accumulation signal $x(t)$.

In what follows, we will perform two series of analyses showing evidence for the LIT in data. This is all but straightforward as models without an analytical expression for their likelihood are notoriously hard to estimate. For the first series of analyses, we will use a novel likelihood free estimation method, the prepaid method (Mestdagh et al., 2019). Because this method is still new, we will avoid adding extra layers of methodological complexity and limit our analysis to what can be based on direct prepaid estimates of the DDM and its LIT extension. A second series of analyses will be based on an analytical approximation of the LIT, which is easier to estimate. We will revisit the first series of analyses, but also include a comparison to the Ratcliff diffusion model (RDM), which includes extra inter-trial variability parameters, and check a number of predictions that can be made in this context.

To avoid confusion, we will label traditional diffusion model parameters with a small ring on top: boundary separation $\mathring{a}$, starting position $\mathring{x}_0$, drift speeds $\mathring{v}_i$, and for the Ratcliff diffusion model (RDM), additional inter-trial variability parameters $\mathring{\eta}$ (variability on drift speed) and $s\mathring{x}_0$ (variability on starting position). To account for parameter redundancy, we again fix $\mathring{\sigma} = 0.1$.

2   Direct parameter estimation with prepaid

2.1   Methods

Having established its main features, we will now show how the LIT extension of a constant drift diffusion model can be fitted to data. Because the model does not have an analytic expression, statistical inference is not trivial. Calculating probability density functions through numerical integration as in Voss and Voss (2008) becomes extremely slow as a 2D discretization is required, compared to the 1D discretization in the case of a standard diffusion model. Simulating the model with an Euler-Maruyama approximation is still straight forward, but simulation based techniques for inferring parameters (Beaumont et al., 2002; Heard et al., 2015; Wood, 2010) are known to have convergence issues and are tediously slow, especially if one wants to do many estimations to verify the method and model's recovery properties. We instead opt for a novel method of inference, which has recently been coined as the prepaid method (Mestdagh et al., 2019). In this approach, data is simulated for a comprehensive grid of parameter sets, covering the entire parameter space of interest. The resulting massive look-up table is called a prepaid database. For the particular case of choice RT models, a special grid is used that exploits a number of symmetries and scaling properties to keep the final grid dimensionality low. In what basically boils down to a parameter grid search, the prepaid grid is then used to find minimal values of an objective function operating on its simulated probability density functions. We will not go into the details of this technique (see Mestdagh et al. (2019) for an in-depth treatment), but we will show its accuracy using a comprehensive simulation study. Because many datasets, simulated as well as experimental, are to be evaluated in this grid, the one shot cost of creating the prepaid database of probability density functions, quickly becomes justified. We use the objective function of the D*M approach (Verdonck and Tuerlinckx, 2015), in order not to make shape assumptions on the non-decision distribution. The D*M objective function capitalizes on the assumption that the non-decision time, whatever shape it might have, is shared at least across some condition-response pairs. For instance, the error responses of stimulus A and the correct responses of stimulus B could be considered to have the same non-decision time distribution, without any extra constraints on the shape of that distribution. We can now simply convolve (or stochastically add) the observed error response distribution of stimulus A (also containing some non-decision time) with a model distribution for the correct responses of stimulus B (not containing any non-decision time), and vice versa (switching which is model and which is observed). This way, we end up with two convolutions that each contain three parts: the decision time part of the error responses of A, the decision time part of the correct responses of B and a non-decision part. For the first convolution the non-decision part comes from the observed error response distribution of stimulus A, for the second convolution, the non-decision part comes from the correct response distribution of stimulus B. As a whole, the two convolutions should be identical if the model matches the decision part of the observed data. By minimizing the difference between these two convolutions (or a sum of these, given there are a lot of these condition-response pairs to compare) we can now try and find the optimal parameters for the model. In Verdonck and Tuerlinckx (2015) it is shown that this approach significantly reduces estimation biases that are caused by the misspecification of the shape of the non-decision time distribution (e.g., assuming a uniformly distributed non-decision time). For a more

rigorous explanation of the D*M approach, we refer the reader to this paper.

To test the method, we generate data from a broad range of parameter sets and estimate them as described above. We assume four stimulus difficulty conditions (differing only in drift speed v_i with $i = 1 \ldots 4$) covering a typical range of accuracy (70 and 90 percent, for each direction of evidence). As is the case for most diffusion models of choice RT, there is one redundant parameter in the model, so we fix $\sigma = 0.1$. In Figure 2, the relation between actual and estimated parameters is shown in a set of scatter plots, for different numbers of trials per condition ($N = 10000$, $N = 300$), as well as the coverage plots based on non-parameterically bootstrapped confidence intervals. A coverage plot can be used to validate the way confidence intervals of parameter estimates are calculated (in this case the method is non-parametric bootstrap). The plot shows what fraction of the original parameter values for which data was generated, fall outside of the confidence intervals estimated for that data. For $N = 10000$ trials per stimulus condition (plot (a)), model parameters (boundary separation a, starting point x_0, inverse leak λ^{-1} and drift speeds v_i) are recovered almost perfectly, without any apparent biases. The exact distribution of test parameter sets is related to the underlying parametrization of the prepaid grid, which we will not expand on here (more details are provided in Appendix III); looking at the recovery plots however, one can assess the effective range. The bottom right plot is a parameter coverage plot, with each line representing the coverage for a different parameter. The horizontal axis indicates p in the confidence interval $[\frac{p}{2}, 1 - \frac{p}{2}]$, varying from 0.05 to 0.95 with steps of 0.05. The vertical axis shows the fraction q of the original parameter values falling outside of their estimated confidence intervals. Perfect coverage manifests as a straight diagonal line for each parameter. The obtained coverages are very acceptable in that sense. For $N = 300$ trials per stimulus condition (plot (b)), parameter recovery is of course a bit less accurate, but looking at the coverage plot, confidence intervals on the estimates are in line with what is expected. From these results, it's safe to assume the estimation algorithm behaves as it is supposed to.

The recovery results for the DDM (traditional diffusion model without variability parameters) and the RDM (traditional diffusion model with variability parameters) can be consulted in Figures 15 and 16 in Appendix III.

2.2   Results

2.2.1   The occurrence of leak.   With the fitting procedure in place, we now investigate to which extent the leak parameter as defined in the LIT extended constant drift diffusion model (which we will, as a whole, also refer to as LIT), adds to a simple DDM description of choice RT data. Because the DDM is nested in the LIT (the DDM is a LIT with inverse leak equal to zero), we could simply check if the estimated inverse leak is significantly higher than zero. However, because the LIT does not allow inverse leak estimates below zero, this would not constitute a good test. To account for the fact that estimates of inverse leak are automatically higher than zero, we take the following approach. For each participant, we fit the LIT on both the original data, as well as data simulated with parameters obtained through a traditional DDM estimate of the original data (a parametric DDM bootstrap). If the LIT indeed adds something unique to the DDM description (i.e., something that could not have been simulated by the best fitting DDM), the inverse leak
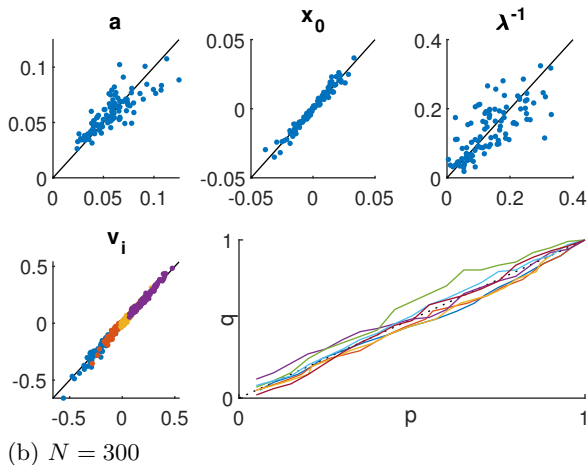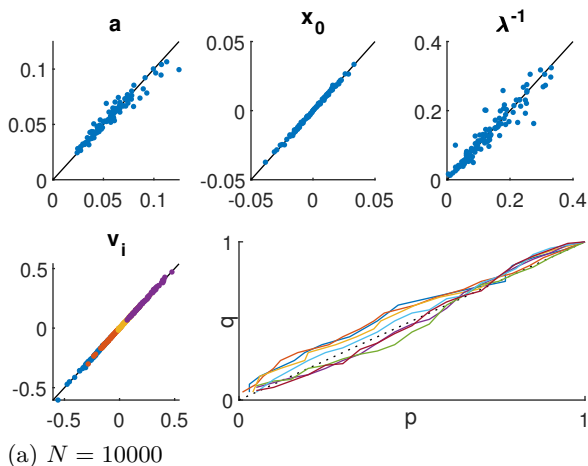
Figure 2. Scatter plots of original versus recovered parameters for the LIT extended DDM. Parameters are respectively boundary separation a, starting position x_0, inverse leak parameter λ^{-1}, and four drift speeds v_i. The last graph represents the coverage of the bootstrapped confidence intervals. Plot (a) shows the recovery results for $N = 10000$ trials per stimulus condition, plot (b) for $N = 300$.

estimates based on the original data should be higher than the ones based on the DDM parametric bootstrapped data.

We analyze a diverse collection of choice RT data sets. The first dataset is based on a face/car discrimination experiment containing 32 participants (Philiastides et al., 2014a). It comprises 8 stimulus conditions (4 difficulty levels for each of the face and car stimuli) and 60 trials per condition. Because that number is rather low, we collapse the 4 stimulus difficulty levels into 2 levels (difficult and easy) and do so for each of the face and car stimuli. The second dataset pertains to a lexical decision making experiment, where words are discriminated from non-words (Wagenmakers et al., 2008). The dataset covers 17 participants, 6 stimulus conditions (3 degrees of word and 3 degrees of non-word), and 160 trials per condition. The third dataset is based on a random-dot direction discrimination paradigm, where a cloud of moving dots is judged to be going to the left or the right (Dutilh

Figure 3. Evidence for the inverse leak parameter. In this plot, within-person differences are shown between the estimate of inverse leak based on the original data on the one hand and on DDM bootstrapped data on the other. Values higher than zero indicate an inverse leak estimate that is higher than what is to be expected according to a simple DDM. Accuracy (blue) and speed (red) conditions are considered independently.

et al., 2018). It consists of 20 participants, each performing an extensive choice RT task that can be split up in three different datasets according to three bias conditions (without bias, with bias to the left, with bias to the right). There are 4 stimulus conditions (2 difficulty levels, for each of the left and right directions), with trial numbers per condition varying between 50 and 230. For the latter two datasets, there is separate choice RT data for a speed and an accuracy instruction.

In first instance, we analyze the speed and accuracy conditions separately (all parameters are allowed to differ), with only stimulus difficulty level to be accounted for in terms of experimental conditions. In Figure 3 we show the participant distribution of the difference between the inverse leak estimates based on the original data and those based on a DDM parametric bootstrap, for both the accuracy and the speed condition.

We see that for the default accuracy instruction (blue), on the level of the participant population, the inverse leak estimates are significantly greater for the original data than the DDM bootstrapped data (two-sided paired t-test: $p = 0.02, p = 0.02, p < 10^{-5}$ for the random-dot data, the lexical data and the face/car discrimination data respectively). For the speed instruction however, the inverse leaks estimated on the original data do not significantly exceed the DDM bootstrap estimates ($p = 0.7$ for the lexical data; for the random-dot data, they are even slightly below with $p < 10^{-5}$). It should be stressed that the distributions plotted in Figure 3 are based on individual participant estimates and do not convey the uncertainty of their statistical means, which are the quantities that are tested.

The fact that the leak only contributes additional explanatory power in the accuracy condition is completely inline with its supposed effect on SAT. For low values of inverse leak, we don't expect the leak estimate to significantly exceed the benchmark estimate based on DDM bootstrapped data.

2.2.2   SAT manipulation: leak versus boundary separation.   To address the question of selective parameter influence of the SAT manipulation more carefully, we estimate the speed and accuracy conditions together, in a single model. We fix all parameters across speed and accuracy conditions, except for a and λ^{-1}, respectively (details on how

Figure 4. Boundary separation a versus inverse leak λ^{-1} as a mechanism for SAT. In this plot, within-person differences $fit_a - fit_{\lambda^{-1}}$ of the LIT fit for boundary separation on the one hand and inverse leak on the other is shown. Participants of both speed/accuracy experiments are included. Positive values indicate a better fit for the inverse leak parameter as a mechanism for SAT.
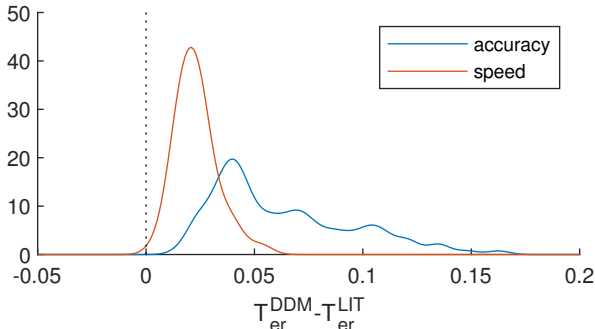

Figure 5. Differences in mean non-decision time estimates between LIT and DDM. Accuracy condition is in blue, speed condition in red.

this is done with the prepaid method can be found in Appendix IV). Because both models under consideration have the same amount of parameters, we can simply compare fit to determine the better model. The difference in fit is shown in Figure 4. Both data sets with a speed accuracy manipulation, the larger random-dot data set (N=60), and to a lesser extent the lexicon data set (N=17), support the conclusion that inverse leak outperforms boundary separation as a mechanism for SAT ($p < 10^{-5}$ and $p = 0.12$, respectively).

2.2.3 Non-decision time: LIT versus DDM. Equation 6 suggests that the LIT can account for some of the stimulus independent non-decision time by itself, which should lead to smaller non-decision time estimates, especially in the accuracy condition where the inverse leak is higher. We estimate mean non-decision time by subtracting the mean of the model distribution (obtained through the initial D*M procedure) from the mean of the empirical distribution. In Figure 5 we plot the DDM-LIT differences in mean non-decision time T_{er} for all experiments combined, using the initial, non-restricted estimates. We see that the mean non-decision time estimated under the LIT assumption is indeed always smaller than its DDM counterpart (for all participants, across all experiments and conditions), and that this effect is most pronounced in the accuracy condition.

Another way to differentiate LIT from DDM is to look at the difference between T_{er}

Figure 6. Differences in mean non-decision time estimates between the accuracy and the speed condition, for both the DDM (dashed line) and the LIT (solid line).

estimates under a SAT manipulation. If SAT is indeed driven by the inverse leak of a LIT and not simply by the boundary separation in a DDM context, there should be a noticeable difference between speed and accuracy estimates of non-decision time for a DDM analysis, that would largely disappear in a proper LIT analysis. Figure 6 shows the distributions of these differences for both a DDM and a LIT analysis. As predicted, there are significant differences between the estimates of T_{er} in the speed and accuracy condition ($p = 0.02, p < 10^{-5}$ for lexicon and random-dot data respectively) for the DDM analysis, but they largely disappear in the LIT analysis ($p = 0.7, p = 0.1$).

3 Approximating the LIT as a functionally reparameterized drift diffusion model

To get a better understanding of the reasons why the LIT performs better than a simple threshold, we use an analytical approximation of the LIT that presents it as a functionally reparameterized standard diffusion model. Based on the results of (Hagan et al., 1989), the choice RT distributions of the LIT can be reduced to those of a DDM with a more complicated influence of stimulus and SAT conditions on model parameters.

3.1 Methods

Equation 6 suggests that, assessed through the lens of a traditional DDM, the LIT predicts a starting position that is dependent on the stimulus. This would result in a violation of a common assumption regarding the selective influence of the stimulus condition on that level of analysis, i.e., that only drift speed is affected by the stimulus. We can incorporate this relationship into a standard diffusion model by defining the effective starting position for stimulus i as

$$\dot{x}_{0,i} = \dot{x}_0 + \frac{\dot{v}_i}{2\dot{\lambda}},$$

with drift speed $\dot{v}_i$, shared starting position $\dot{x}_0$ and pseudo leak $\dot{\lambda}$, parameters that have to be estimated in the context of a standard diffusion model. We have dotted the symbols to avoid confusion with earlier notations. A problem with this expression is that $\dot{x}_{0,i}$ can exceed boundaries $\frac{\dot{a}}{2}$ or $-\frac{\dot{a}}{2}$, which is of course not allowed. We solve this problem by using a slightly adjusted expression that prevents this:

parameter	grid values
a	$0.04, 0.06, 0.08, \dots, 0.2$
x_0	$[-0.1, 0, 0.1] \times a$
λ^{-1}	$0.001, 0.05, 0.1, 0.15, \dots, 0.5$
	fixed values
v_i, with $i = 1, \dots, 7$	$0.15, 0.1, 0.05, 0, -0.05, -0.1, -0.15$
T_{er}	0.3

Table 1

Parameter grid used for simulating a broad selection of LIT choice RT data. All combinations of the values listed for a, x_0 and λ^{-1} are used. For each combination, the same seven values of v_i and $T_{er} = 0.3$ are used.
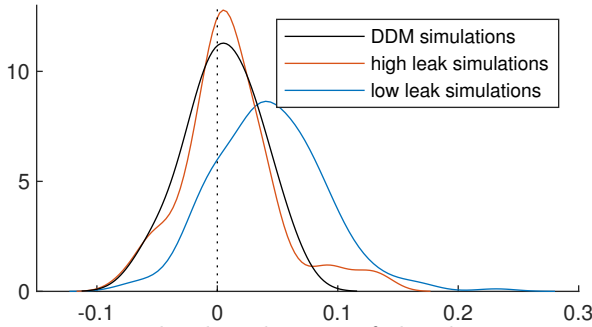

Figure 7. The distribution of the slopes $\dot{s}$ estimated for DDM simulated data (black), high leak LIT simulated data (red) and low leak LIT simulated data (blue).

$$\dot{x}_{0,i} = \frac{\dot{a}}{\pi} \arctan(\tan(\frac{\pi \dot{x}_0}{\dot{a}}) + \dot{s}\dot{v}_i),$$

with $\dot{s}$ a simplified slope parameter. Given this expression, estimates of slope $\dot{s}$ should become greater than zero for data generated with low values of the original leak parameter λ. To check if this holds in practice, we simulate a broad selection of choice RT data from all possible combinations of the a, x_0, and λ^{-1} values given in Table 1, consistently assuming seven stimuli with drift speeds v_i ($i = 1, \dots, 7$) given in that same table. For each set of parameters, 300 trials were simulated for each of the seven stimuli. Using a DDM with a starting point/drift speed dependency as described above, we estimate the parameters of all these simulated data. The probability density functions of the underlying DDM are calculated as in Voss and Voss (2007), the objective function again is D*M, and finally, as a global minimization heuristic we opt for differential evolution (Storn and Price, 1997). We repeat this exercise for data generated with a normal DDM. The same parameters are used as in Table 1 for the simulations, ignoring the LIT specific λ parameter.

In Figure 7, the estimated slopes $\dot{s}$ are shown for LIT and DDM generated data. As expected, data simulated with the DDM shows no clear sign for $\dot{s}$, nor does data simulated with a high leak LIT ($\lambda^{-1} \leq 0.1$), which is close to a simple DDM). LIT data simulated with a low leak ($\lambda^{-1} > 0.1$), on the other hand, does show a clear positive $\dot{s}$.

Figure 8. Density of slopes $\dot{s}$ - a proxy for inverse leak - for all three experiments combined. The colors red and blue are used to distinguish between speed and accuracy conditions.

3.2   Results

3.2.1   Reverifying the occurence of leak.   Figure 8 shows the distribution of slopes $\dot{s}$ found for all datasets combined, seperated for speed (red) and accuracy conditions (blue). Because negative slopes are possible in this model, and DDM does not create a bias on the leak parameter (see Figure 7), we can perform a simple test of the mean with regards to zero, without the need for bootstrapping as we did in the case of direct estimation.

In line with the results shown in Figure 3, estimated slopes $\dot{s}$ - which can be seen as a proxy for inverse leak - are significantly larger than zero for the accuracy condition ($p = 0.004$, $p < 10^{-5}$, $p < 10^{-5}$ for the lexicon, random dot and face/car data, respectively), and not for the speed conditions. On the contrary, a slightly negative result was found for the speed condition of the random-dot experiment ($p = 0.04$).

This suggests that the assumption that a change in stimulus difficulty only influences the drift speed is indeed violated in the accuracy condition in the context of the DDM (slopes introduce bias on the starting position that depends on the drift speed and thus differs from stimulus to stimulus) and that the LIT can account for at least some of it.

3.2.2   Leak outperforms inter-trial variability of starting position and drift speed.   In Figure 9, we show the differences in fit between three single parameter extensions of the DDM. We find that inverse leak λ^{-1} outperforms both inter-trial variability of starting position and drift speed in the accuracy condition ($p < 10^{-5}$ and $p < 10^{-5}$ respectively, for all datasets combined). In the speed condition, inverse leak also outperforms variability of drift speed ($p < 10^{-5}$), and performs comparable to variability of starting position ($p = 0.4$). Comparing between the two variability parameters of DDM, it can be seen that variability of drift speed is more important for the accuracy condition ($p < 10^{-5}$), and variability of starting position is more important in the speed condition ($p < 10^{-5}$). Looking at the smaller lexicon dataset independently (N=17), results were not always significant, but always with the correct sign.

3.2.3   Selective parameter influence of SAT manipulation: leak versus boundary separation.   As it stands, boundary separation is typically considered to be the parameter that operationalizes SAT regulation. This reflects another common selective influence assumption for the DDM, namely that drift speeds do not change under SAT manipulations, nor does non-decision time. However, traditional diffusion modeling results from the literature (Lerche and Voss, 2018; Rae et al., 2014; Zhang and Rowe, 2014) suggest

Figure 9. Differences in fit between three single parameter extension of the DDM: inverse leak λ^{-1}, variability of starting position zr and variability of drift speed η. Each box plot shows the distribution of the participants differences in fit between an extension that uses the parameter labeled on the left versus the right. The more the distribution leans to the side of a parameter, the more that parameter outperforms the other. The colors red and blue are used to distinguish between speed and accuracy conditions.

that the parameter influence of a SAT manipulation is not limited to a change in boundary separation. There seems to be an additional change in non-decision time (smaller for the speed condition) and drift speed (also smaller for the speed condition). In what follows we investigate if a leak driven SAT, predicts this typically observed pattern of parameter selectivity for the SAT manipulation.

In other words, we want to investigate what can be predicted about the parameter influence of a SAT manipulation through the lens of a traditional diffusion analysis, using either boundary separation, or leak as the operative SAT mechanism. We explore three distinct cases: data generated with a constant drift diffusion model and SAT driven by simple boundary separation $\mathring{a}$, data generated with a LIT extended version of said model but SAT still driven by — now motor level — boundary separation a, and finally, data generated with the LIT extended version but SAT driven by leak parameter λ. We simulate data for these three combinations using parameters in the range of what we found from analyzing the datasets discussed above. More specifically, for the DDM we use basic parameters $\mathring{a} = 0.1, \mathring{x}_0 = 0, \mathring{T}_{er} = 0.3$ and $\mathring{v}_i = [0.2, 0.1, -0.1, -0.2]$ with $i = 1, \ldots, 4$ representing four different stimulus conditions. As SAT driving parameter we use boundary separation $\mathring{a}_c = [0.1, 0.2]$, with $c = 1, 2$ representing speed and accuracy conditions. For the LIT extended version we use basic parameters $a = 0.1, x_0 = 0, T_{er} = 0.3$ and $v_i = [0.2, 0.1, -0.1, -0.2]$. As SAT driving parameter we use either boundary separation $a_c = [0.1, 0.2]$, or leak parameter $\lambda_c = [500, 5]$.

To approach realistic effect-sizes, we simulate 300 trials for each of the 4 stimuli, both in a speed and an accuracy condition. We then perform a traditional diffusion analysis according to the BIC-optimal parameter selectivity found in Rae et al. (2014). More specifically, this means that drift speed is the only parameter that can differ across stimulus conditions, and that the SAT manipulation is allowed to be driven by changes in boundary separation, drift speed and non-decision time. As this is a traditional diffusion analysis the estimated parameters are labeled with a small ring on top. For each of the three combinations described above, we simulate 100 datasets and repeat the estimation procedure, to get an idea

about theoretical confidence intervals. To remain consistent with standard methodology (we want to be able to compare with previous results from literature), we opt for a maximum likelihood estimation procedure with a uniformly distributed non-decision time with mean $\mathring{T}_{er}$ and width $s\mathring{T}_{er}$, as implemented in Voss and Voss (2007). Inter-trial variability for starting position and drift speed are allowed as condition common parameters for the same reason.

The results are shown in Figure 10. If boundary separation is the parameter driving SAT (first two rows), as expected, no other parameters are influenced by the manipulation. If leak parameter λ is driving SAT, however, mean non-decision time $\mathring{T}_{er}$ and to a lesser extent drift speed $\mathring{v}_i$, also seem to be affected, more specifically with the speed instruction leading to smaller non-decision times and smaller drift speeds. This prediction is in agreement with the analyses of Rae et al. (2014), Zhang and Rowe (2014), and Lerche and Voss (2018), that find similar changes in non-decision time and/or drift speed above and beyond the change in boundary separation.

We run the same diffusion analysis on the two datasets in this paper that have a SAT manipulation. The results are shown in Figure 11 as density plots of the within-person differences between speed and accuracy conditions for boundary separation $\Delta\mathring{a} = \mathring{a}^{acc} - \mathring{a}^{speed}$, mean non-response time $\Delta\mathring{T}_{er} = \mathring{T}_{er}^{acc} - \mathring{T}_{er}^{speed}$ and $\Delta\mathring{v} = \frac{1}{4}\sum_{i=1}^{4}(|\mathring{v}_i^{acc}| - |\mathring{v}_i^{speed}|)$, respectively. For all these differences, we find significant results in both the random dot data ($p < 10^{-5}$, $p < 10^{-5}$, $p < 10^{-5}$, respectively), and the lexicon data ($p < 10^{-5}$, $p = 0.03$, $p = 0.05$, respectively). These results again echo what is found in the literature (Lerche and Voss, 2018; Rae et al., 2014; Zhang and Rowe, 2014), i.e., apart from an increase in boundary separation, a traditional diffusion analysis reveals additional increases in non-decision time and drift speed, for the accuracy condition. From the three theoretical cases studied above, this seems to be compatible only with a leak driven SAT.

Finally, we check if the LIT approximating diffusion model introduced in the beginning of this section is indeed able to remove the influence of the speed-accuracy manipulation on drift speeds (we do not predict this for boundary separation or non-decision time, which are still parameterized in a way that should change under leak induced SAT). In Figure 12 we plot the difference between the absolute values of the drift speeds in the speed and accuracy condition, averaged across stimuli. In line with the last plot of Figure 11, we see that the absolute drift speeds are larger in the accuracy condition when estimating a normal DDM ($p < 10^{-5}$ for the random dot data; $p = 0.004$ for the lexicon data). This difference is no longer significant when estimating the LIT approximating diffusion model ($p = 0.1$ and $p = 0.7$, respectively).

Figure 10. Traditional maximum likelihood estimates of three sets of simulated data, each generated with a different model/SAT mechanism (see row labels). In the recovery, SAT effects are allowed on boundary separation $\mathring{a}$, drift speeds $\mathring{v}_i$ and non-decision time $\mathring{T}_{er}$, while stimulus effects are only allowed on $\mathring{v}_i$. Speed condition bootstrap distributions are in red, accuracy condition bootstrap distributions are in blue.


Figure 11. Influence of a SAT manipulation on traditional diffusion model parameters in real data (random-dot and lexicon data combined). The same design and estimation procedure was used as for the theoretical predictions. Each density plot shows the distribution of the within-person shifts for the respective parameter going from speed to accuracy conditions.

Figure 12. Influence of SAT on drift speed, for both LIT ($\Delta\dot{v}$, solid line) and DDM ($\Delta\mathring{v}$, dotted line). Each line represents the participant distribution of the difference between the absolute values of drift speed for the speed and accuracy condition respectively, averaged across stimuli.

4   Applied to the Leaky Competing Accumulator and Ising Decision Maker

Finally, we briefly investigate the effect of the LIT on more complicated models of choice RT such as the Leaky Competing Accumulator (LCA) and the Ising Decision Maker (IDM). Both models are two-dimensional evidence accumulators, and are given a separate, single boundary LIT for each dimension. The two LITs then race each other for the eventual choice. Applying a LIT in this way, results in much of the same effects we found when applying the LIT to the DDM. Figures 13 and 14 are structured in the same way as Figure 1, but now deal with the LIT extensions of the LCA and the IDM respectively. Analogous to the LIT extended DDM, the two LIT signals $(y_1(t), y_2(t))$ are lagging the actual evidence accumulation signals $(x_1(t), x_2(t))$, and have an increased signal-to-noise ratio because of the LIT's smoothing effect. Again, this leads to effective decision boundary distributions and a shift in effective non-decision time for low leaks. Each trial's effective decision boundary however, now has double coordinates and can be represented as a point $(x_1(t_r), x_2(t_r))$ (with t_r the moment the first LIT reaches its boundary). The effective decision boundary distributions this amounts to, are, for both figures, shown in plot c) as two-dimensional distributions of points. Interestingly, for lower values of λ, the LIT creates an effective boundary distribution on the evidence space that starts conforming to the associated minimum of the potential surface underlying the decision dynamics — the solid black lines in the background of plot c) are isoclines of this surface (for more details on potential surfaces, see Verdonck and Tuerlinckx (2014)). This seems to be the case for both the IDM and the LCA. Finally, plot d) shows that the LIT can account for a SAT dependent non-decision time, also for these more complicated models. Fitting these advanced models to data is beyond the scope of this paper.

Figure 13. Evidence integration and lagged LIT buildup for a LCA with Leaky Integrating Thresholds. Plot (a) shows, response locked, the noisy evidence accumulation signal (solid) and accompanying lagged and smoothed LIT signal (dotted) for a trial with a high leak ($\lambda = 200$, red) and a trial with a low leak ($\lambda = 5$, blue). Plot (b) shows the average evidence (solid) and LIT signals (dotted) where the first LIT wins the race, response locked, for the same two values of λ, but both for three different levels of stimulus difficulty. In plot (c), the distribution of evidence values (x_1, x_2) at which the first passage on the level of y_1 or y_2 is reached, is shown. Finally, plot (d) shows the respective choice RT probability density functions. For an overview of the parameters that are kept constant, see Table 3 in Appendix II.

Figure 14. Evidence integration and lagged LIT buildup for an IDM with Leaky Integrating Thresholds. Plot (a) shows, response locked, the noisy evidence accumulation signal (solid) and accompanying lagged and smoothed LIT signal (dotted) for a trial with a high leak ($\lambda = 200$, red) and a trial with a low leak ($\lambda = 5$, blue). Plot (b) shows the average evidence (solid) and LIT signals (dotted) where the first LIT wins the race, response locked, for the same two values of λ, but both for three different levels of stimulus difficulty. In plot (c), the distribution of evidence values (x_1, x_2) at which the first passage on the level of y_1 or y_2 is reached, is shown. Finally, plot (d) shows the respective choice RT probability density functions. For an overview of the parameters that are kept constant, see Table 4 in Appendix II.

5 Discussion

In this paper, we propose a Leaky Integrating Threshold process boundary as an alternative for the traditional simple threshold for evidence accumulation models of choice RT. This relaxes the assumption that the executive process leading to the physical response is – apart from receiving the choice once it has been decided – independent of the evidence accumulation process. The LIT is defined as a simple leaky integrator of the evidence variable(s) of the decision model. The integrated signal is associated with motor preparation and can be seen as a lagged and smoothed version of the original evidence accumulation signal. It is the latter signal that is compared to a final threshold. In other words, the primary process of evidence accumulator hands control over to the motor accumulation process, which becomes the main determinant of the eventual choice.

Importantly, the smoothing of the evidence accumulation process happening at the LIT stage makes the effective noise of the final accumulator auto correlated, which is a fundamental difference from most choice RT models, including the Leaky Competing Accumulator and the Urgency Gating Model with evidence filtering (Cisek et al., 2009). Although both of these models contain some form of leakage, it is not in the context of double integration of system-noise, which is the fundamental mathematical feature of the LIT setup. We view this as the main novelty of the work, which we believe can have wider implications on how (perceptual) decisions are being modelled.

Applied to the simple constant drift diffusion model, the LIT is shown to be superior, both with and without inter-trial variability. We replicated these results across all three datasets. Additionally, as the LIT increases the effective signal-to-noise ratio of the evidence accumulation signal at the expense of response time, the LIT leak parameter can be seen as an alternative to boundary separation for operationalizing SAT regulation. We tested this alternative formally and showed that the leak parameter is better at explaining the differences between speed and accuracy conditions than the boundary separation.

We obtained further validation of our model through the lens of traditional diffusion model analyses. First, we show how the LIT can be approximated by a diffusion model that is functionally reparametrized. More specifically, a change in drift speed on the level of the evidence accumulation process within the DDM-LIT architecture, when looking at the corresponding parameters of a traditional diffusion model approximating that architecture, results in a change of not only the drift speed, but also the starting position and the non-decision time. This predicts a violation of the selective influence of stimulus difficulty on drift speed in a DDM analysis where starting position is also allowed to vary for different levels of stimulus difficulty. In strong support of the LIT, we discovered this violation in all three datasets, at least for the accuracy condition, where the effect of the leak parameter should be strongest.

Moreover, generating choice RT data with the leak as the underlying driver for SAT and then performing a traditional diffusion analysis (without allowing starting position to vary in tandem with drift speed), resulted in a specific pattern of SAT sensitive parameters: in the speed condition, we observed decreases in non-decision time and drift speed, along with decreases in boundary separation. This pattern is in line with what is found previously for a similar analysis of the data and with what has been repeatedly found for SAT manipulations in the literature (Lerche and Voss, 2018; Rae et al., 2014; Zhang and Rowe, 2014). We find

the same pattern of SAT sensitivity in our own analyses. Crucially, in a LIT compatible diffusion analysis (that allows starting position to covary with drift speed for changing stimulus difficulty levels), drift speed should no longer change under SAT manipulation. This was fully corroborated in both SAT datasets.

Finally, we showed that for more advanced two-dimensional models of decision making like the LCA or the IDM, the LIT results in effective decision boundaries that take the shape of the underlying potential function, constituting a more natural decision convergence criterion.

5.1   Connection with earlier models of choice RT

Formally, the key property of the LIT is that the leaky integration is done on already stochastically integrating evidence and that the system-noise driving the primary evidence accumulator is, therefore, doubly integrated. This results in auto-correlated noise on the level of the final motor accumulator, which is the main feature that sets the LIT apart from traditional diffusion models. The interplay between this type of noise and an absorbing boundary is far from trivial (Hagan et al., 1989), and gives the model a unique signature, with important implications for commonly accepted assumptions about what influence experimental manipulations should have on traditional diffusion model parameters.

The cascade model as described by McClelland (1979) adopts a comparable approach with regards to the deterministic part of the accumulation processes, but is quite different in the way it introduces variability in choice responses. In contrast to even a standard diffusion model, variability in the cascade model is not obtained with the integration of system noise throughout the trial, but added in through the variability of parameters. In a sense, the LIT marries the idea of cascading explored by McClelland with the stochastic integration typical for normal diffusion models. This fundamentally results in a double integration of system-noise, which brings the LIT into a different class of diffusion models (with auto-correlated noise).

Moreover, there are a number of functional/interpretational differences between the LIT and cascade models. Firstly, in most cascade models, the response execution (final stage in the cascade) is assumed to be a discrete event. In the LIT, the motor system is continuously active and it is the properties/dynamics of the motor accumulator itself that determine the decision – not the stage prior. In other words, the motor preparation deserves a dedicated level in the cascade. This aspect of the LIT design is motivated by recent neural observations in the literature (as we discussed in the Introduction) and in this sense the LIT can be thought of as having a more biologically plausible architecture.

Secondly, in cascade models, the layers are forever connected in the same way, with no motivation or pathway towards the implementation of external regulation. In principle, in the LIT (although it was not required for the applications analyzed in this paper), the motor accumulator should terminate and return to baseline after motor execution (as one would expect electrophysiologically), while the primary accumulator can continue to linger on past the response. In scenarios in which a secondary decision is required (or triggered), the motor accumulator can (will) start up again and proceed to receive information from the primary accumulator. This would open up possibilities for modeling more advanced features of decision making, like change-of-mind decisions, post-decision confidence reporting, etc.

We discuss, qualitatively, how such secondary decisions might be implemented within the LIT framework below.

Furthermore, we compare the idea of using the LIT for triggering the actual decision with the traditional view of using a simple threshold, be it fixed or collapsing over time. One important difference between these three accounts is the value of accumulated evidence at the moment the actual decision is made. For a fixed boundary this is obviously the same regardless of stimulus difficulty. For the collapsing bounds criterion, stronger stimuli hit the bounds for higher values of accumulated evidence, seeing the collapse has progressed less for earlier responses. This leads to stimulus-dependent values of accumulated evidence at the moment of decision, with higher values for stronger stimuli. A similar pattern emerges from the LIT, albeit through an entirely different mechanism, and is also born out in EEG data (Philiastides et al., 2014b).

Similar in all three SAT mechanisms is that they do not fundamentally interfere with the actual evidence accumulation process. The main difference is the way they use the noisy accumulated evidence to come to a decision. For both fixed and collapsing bounds, SAT is manipulated by changing the evidence criterion (require more accumulated evidence or make do with less). In the case of collapsing bounds the same philosophy applies, but a continuously encroaching deadline is built in. SAT can be controlled by either changing the initial value of the threshold or the rate at which it collapses. In both cases however, the bounds only monitor the noisy accumulated information signal at that given point in time, and do not "remember" what they already saw before. In the LIT, instead of changing the end criterion, time integration is used to make the signal less noisy. In a setting with higher accuracy, the boundary is no longer simply monitoring the noisy evidence accumulation signal for an immediate crossing, but uses a time smoothing window to increase the signal–to–noise ratio and avoid an unwanted accidental crossing. In other words, the boundary has a leaky memory, and uses this property to consider past evidence in the accumulated signal, rather than just its current state. Mathematically, when the leak parameter approaches infinity, the system becomes "memoryless" and the LIT reduces to a simple DDM.

Finally, there are other models that share some aspects with the LIT explored in this paper. For instance, the Leaky Competing Accumulator also has a leak parameter and one of the Urgency Gating Models proposed in Cisek et al. (2009) performs some kind of filtering on the level of the input. However, none of these models share the LIT's doubly integrated system-noise which has a specific effect on its first passage times (Hagan et al., 1989). In our analyses we limit ourselves to comparisons with the benchmark DDM and RDM, of which the LIT is a natural extension. It is possible, however, that another model could resolve some of the same issues in a different way, or that vice versa the LIT may provide a viable alternative explanation for the paradigms and data used to validate these other models. For example, the gating in Cisek et al. (2009) is used as a way to prevent early random fluctuations of the input to trigger premature responses. The LIT on the other hand smooths out fluctuations on the level of the accumulated evidence rather than acting immediately on the level of the input signal, also resulting in robustness to early random fluctuations of the input, if they were to be considered. Ultimately, it would be an interesting exercise to check how different models compare to each other under all commonly used experimental paradigms. This, however, would also include combinations of models that are not mutually exclusive – the LIT can be combined with almost every existing model

of evidence accumulation, creating a host of new models. The comparison could be done by straight forward model selection, or perhaps just as importantly, in terms of parameter influence of experimental manipulations. In sum, formally comparing all existing models of choice RT on a variety of benchmark choice RT data is work long overdue, but it would be a project in its own right and is beyond the scope of this paper.

5.2   Future utility of the LIT

To our knowledge, this is the first formal, fully estimatable model in which a separate (secondary) motor accumulation process becomes part of the causal chain of events, by receiving stochastically accumulated decision evidence and ultimately taking control over the primary accumulation process to drive the final decision. We view this as the main novelty of the work, which we believe can have wider implications in how (perceptual) decisions are being modeled beyond the specific application of the SAT. For example, the LIT could be relevant in scenarios involving integration of non-stationary evidence with varying degrees of temporal uncertainty (Ossmy et al., 2013) by controlling the integration time constant and adjusting the leakage from the primary to the motor accumulator, such that slow leakage could be used to limit the integration of pre–signal noise, similar to what a leak on the decision level would accomplish. One advantage of the LIT over previous work is that commitment to a choice is driven by the motor (secondary) accumulator, which in turn affords additional flexibility in the primary accumulator that could continue to accumulate past the decision and potentially inform secondary (overt) decisions involving additional post-decisional deliberation (e.g."change-of-mind" decisions, post-decision metacognitive appraisal, etc) (Pleskac and Busemeyer, 2010; Resulaj et al., 2009; van den Berg et al., 2016).

For example, consider the case of change-of-mind decisions or double responses (Evans et al., 2020; Resulaj et al., 2009). Because the LIT is in effect a separate motor accumulator, resetting it after a motor action has been triggered does not reset the primary evidence accumulator feeding into it. In the time it takes to complete the motor action itself (time for the actual movement to take place, button press detection) the primary accumulator could have accumulated additional evidence. Immediately after (or even during) this resetting, the motor accumulator could resume accumulating, swiftly picking up on the new evidence accumulated by the primary accumulator in the interim. In turn, this second wave of motor accumulation can enable the selection of a different response in cases where the new evidence now points towards the alternative choice.

The traditional modeling approach for secondary responses in simple threshold systems, also consists of continuing the primary evidence accumulation process after the first decision has been made and monitoring the alternative boundary (Evans et al., 2020), sometimes accompanied by an additional change to the threshold value for this boundary (Resulaj et al., 2009). Although this approach can account for mostly corrective second responses and fewer double responses when decision accuracy is emphasized (two important qualitative features that are observed in this context), it still cannot account for the experimentally observed right skew in the double response times (the time between the first and second response) (Evans et al., 2020). The architecture of the LIT, however, allows the model to properly account for all three observed features. In contrast to the unperturbed primary evidence accumulator, the motor accumulator that triggers the actual response is reset, and has to

evolve to threshold-level again. This results in a typical first passage time distribution, which is right skewed. By resetting the motor accumulator but not the evidence accumulator, the second response can benefit from the continued evidence processing after the first decision, which is required for the second responses to be mostly corrective.

Finally, our general modelling approach may also prove useful for linking other functionally distinct but overlapping processes in the causal chain of events from stimulus to choice like early encoding, memory retrieval (Cox and Shiffrin, 2017), etc.

5.3 Potential neurobiological implementation of the LIT

The observed violations of commonly held assumptions about selective parameter influence of typical experimental manipulations in the DDM and the way the LIT can resolve them are intriguing since they highlight crucial points of departure from conventional modelling and electrophysiological work. The literature has focused extensively on the notion of a single (common) boundary regardless of stimulus condition/difficulty, with evidence arriving from both animal and human work (Kim and Shadlen, 1999; O'Connell et al., 2012; Ratcliff and McKoon, 2008; Shadlen and Newsome, 2001). More recent evidence, however, paints a more complex picture with additional brain responses relating to boundary changes that scale with the degree of decision difficulty as predicted by the LIT, both at the level of individual neurons as well as at the macroscopic level of scalp responses (Bennur and Gold, 2011; Ding and Gold, 2010; Gherman and Philiastides, 2015; Philiastides et al., 2014a; Scott et al., 2017). Similarly, neurobiological work for SAT has focused nearly exclusively on boundary adjustments (Bogacz et al., 2010; Forstmann et al., 2008b; Ivanoff et al., 2008), however, more recent work across species has also provided support for changes in the rate of evidence accumulation and non-decision times (Hanks et al., 2014; Heitz and Schall, 2012; Wenzlaff et al., 2011), in line with predictions stemming from the LIT.

The LIT formulation also has some interesting neurobiological implications as it highlights the need to differentiate between two inter-related but largely separate processes that are likely to take place in different brain structures. While the former can be independent of sensory and response modality, as has been reported in regions of the dorsolateral prefrontal cortex (DLPFC) and lateral intraparietal sulcus/cortex (LIP/IPS) (Filimon et al., 2013; Heekeren et al., 2006; Philiastides et al., 2011; Ploran et al., 2011), the latter would be tightly coupled with structures controlling the specific motor effectors involved in implementing the decision (Donner et al., 2009; Filimon et al., 2013; Tosoni et al., 2008).

We speculate that at the neural level the interplay between evidence and motor accumulation will be implemented via an interaction between prefrontal and parietal regions known to be involved in evidence accumulation (e.g. left DLPFC, LIP/IPS) and (pre)motor structures previously linked to perceptual decision making (e.g. supplementary and pre-supplementary motor areas; SMA/pre-SMA). The latter structures in particular, have traditionally been thought to be directly involved in implementing the SAT by controlling decision boundaries in the brain (Bogacz et al., 2010; Forstmann et al., 2008a; Ivanoff et al., 2008). Using simultaneous EEG-fMRI, however, we have recently shown that the SMA/pre-SM has an even more intricate role in decision making by capturing the full moment-by-moment (temporal) dynamics of the process of evidence accumulation and covarying systematically with traditional evidence accumulation regions in prefrontal and parietal cortex (Pisauro et al., 2017). These recent findings paint a picture in which

(pre)motor structures might take a more elaborate role in decision making, extending beyond a mere involvement in boundary adjustments, consistent with the general framework of the LIT.

Ultimately, the LIT (as any other computational account of decision making) will require independent neurobiological validation. Overall, our work provides a new, viable — and likely more biologically plausible — alternative to the traditional simple threshold evidence accumulation models of choice RT. As such it offers a novel benchmark against which to compare neural data to offer further neurobiological validation for the proposed evidence-vs-motor accumulation processes.

References

Beaumont, M. A., Zhang, W., and Balding, D. J. (2002). Approximate Bayesian Computation in Population Genetics. *Genetics*, 162(4):2025–2035.

Bennur, S. and Gold, J. I. (2011). Distinct representations of a perceptual decision and the associated oculomotor plan in the monkey lateral intraparietal area. *Journal of Neuroscience*, 31(3):913–921.

Bogacz, R., Wagenmakers, E.-J., Forstmann, B. U., and Nieuwenhuis, S. (2010). The neural basis of the speed–accuracy tradeoff. *Trends in neurosciences*, 33(1):10–16.

Cisek, P., Puskas, G. A., and El-Murr, S. (2009). Decisions in changing conditions: the urgency-gating model. *Journal of Neuroscience*, 29(37):11560–11571.

Cox, G. E. and Shiffrin, R. M. (2017). A dynamic approach to recognition memory. *Psychological Review*, 124(6):795.

Ding, L. and Gold, J. I. (2010). Caudate encodes multiple computations for perceptual decisions. *Journal of Neuroscience*, 30(47):15747–15759.

Ding, L. and Gold, J. I. (2012). Neural correlates of perceptual decision making before, during, and after decision commitment in monkey frontal eye field. *Cereb Cortex*, 22(5):1052–67.

Donner, T. H., Siegel, M., Fries, P., and Engel, A. K. (2009). Buildup of Choice-Predictive Activity in Human Motor Cortex during Perceptual Decision Making. *Current Biology*, 19(18):1581–1585.

Donner, T. H., Siegel, M., Oostenveld, R., Fries, P., Bauer, M., and Engel, A. K. (2007). Population activity in the human dorsal pathway predicts the accuracy of visual motion detection. *J Neurophysiol*, 98(1):345–59.

Dutilh, G., Annis, J., Brown, S. D., Cassey, P., Evans, N. J., Grasman, R. P. P. P., Hawkins, G. E., Heathcote, A., Holmes, W. R., Krypotos, A.-M., Kupitz, C. N., Leite, F. P., Lerche, V., Lin, Y.-S., Logan, G. D., Palmeri, T. J., Starns, J. J., Trueblood, J. S., van Maanen, L., van Ravenzwaaij, D., Vandekerckhove, J., Visser, I., Voss, A., White, C. N., Wiecki, T. V., Rieskamp, J., and Donkin, C. (2018). The Quality of Response Time Data Inference: A Blinded, Collaborative Assessment of the Validity of Cognitive Models. *Psychonomic Bulletin & Review*.

Evans, N. J., Dutilh, G., Wagenmakers, E.-J., and van der Maas, H. L. (2020). Double responding: A new constraint for models of speeded decision making. *Cognitive Psychology*, 121:101292.

Filimon, F., Philiastides, M. G., Nelson, J. D., Kloosterman, N. A., and Heekeren, H. R.

(2013). How embodied is perceptual decision making? evidence for separate processing of perceptual and motor decisions. *J Neurosci*, 33(5):2121–36.

Forstmann, B. U., Dutilh, G., Brown, S., Neumann, J., Cramon, D. Y. v., Ridderinkhof, K. R., and Wagenmakers, E.-J. (2008a). Striatum and pre-SMA facilitate decision-making under time pressure. *Proceedings of the National Academy of Sciences*, 105(45):17538–17542.

Forstmann, B. U., Dutilh, G., Brown, S., Neumann, J., von Cramon, D. Y., Ridderinkhof, K. R., and Wagenmakers, E.-J. (2008b). Striatum and pre-SMA facilitate decision-making under time pressure. *Proceedings of the National Academy of Sciences of the United States of America*, 105(45):17538–17542.

Gherman, S. and Philiastides, M. G. (2015). Neural representations of confidence emerge from the process of decision formation during perceptual choices. *NeuroImage*, 106(C):1–10.

Gold, J. I. and Heekeren, H. R. (2013). *Neural mechanisms for perceptual decision making*, book section 19, pages 355–372. Academic Press, London, 2nd edition.

Gold, J. I. and Shadlen, M. N. (2007). The neural basis of decision making. *Annu Rev Neurosci*, 30:535–74.

Hagan, P. S., Doering, C. R., and Levermore, C. D. (1989). The Distribution of Exit Times for Weakly Colored Noise. *Journal of Statistical Physics*, 54(5-6):1321–1352.

Haggard, P. (2019). The neurocognitive bases of human volition. *Annual Review of Psychology*, 70(1):9–28.

Hanks, T., Kiani, R., and Shadlen, M. N. (2014). A neural mechanism of speed-accuracy tradeoff in macaque area lip. *eLife*, 3.

Hawkins, G. E., Forstmann, B. U., Wagenmakers, E.-J., Ratcliff, R., and Brown, S. D. (2015). Revisiting the Evidence for Collapsing Boundaries and Urgency Signals in Perceptual Decision-Making. *Journal of Neuroscience*, 35(6):2476–2484.

Heard, D., Dent, G., Schifeling, T., and Banks, D. (2015). Agent-based models and microsimulation. *Annual Review of Statistics and Its Application*, 2:259–272.

Heekeren, H. R., Marrett, S., Bandettini, P. A., and Ungerleider, L. G. (2004). A general mechanism for perceptual decision-making in the human brain. *Nature*, 431(7010):859–62.

Heekeren, H. R., Marrett, S., Ruff, D. A., Bandettini, P. A., and Ungerleider, L. G. (2006). Involvement of human left dorsolateral prefrontal cortex in perceptual decision making is independent of response modality. *Proceedings of the National Academy of Sciences of the United States of America*, 103(26):10023–10028.

Heekeren, H. R., Marrett, S., and Ungerleider, L. G. (2008). The neural systems that mediate human perceptual decision making. *Nat Rev Neurosci*, 9(6):467–79.

Heitz, R. P. and Schall, J. D. (2012). Neural mechanisms of speed-accuracy tradeoff. *Neuron*, 76(3):616–628.

Ivanoff, J., Branning, P., and Marois, R. (2008). fmri evidence for a dual process account of the speed-accuracy tradeoff in decision-making. *PloS one*, 3(7):e2635.

Kayser, A. S., Buchsbaum, B. R., Erickson, D. T., and D'Esposito, M. (2010). The functional anatomy of a perceptual decision in the human brain. *J Neurophysiol*, 103(3):1179–94.

Kim, J. N. and Shadlen, M. N. (1999). Neural correlates of a decision in the dorsolateral

prefrontal cortex of the macaque. *Nat Neurosci*, 2(2):176–85.

Klein-Flugge, M. C. and Bestmann, S. (2012). Time-dependent changes in human corticospinal excitability reveal value-based competition for action during decision processing. *J Neurosci*, 32:8373–8382.

Lerche, V. and Voss, A. (2018). Speed-accuracy manipulations and diffusion modeling: Lack of discriminant validity of the manipulation or of the parameter estimates? *Behavior Research Methods*, 50(6):2568–2585.

McBride, J., Sumner, P., and Husain, M. (2018). Masked primes evoke partial responses. *Q J Exp Psychol*, 71:1431–1439.

McClelland, J. L. (1979). On the time relations of mental processes: an examination of systems of processes in cascade. *Psychological review*, 86(4):287.

Meister, M. L., Hennig, J. A., and Huk, A. C. (2013). Signal multiplexing and single-neuron computations in lateral intraparietal area during decision-making. *J Neurosci*, 33(6):2254–67.

Mestdagh, M., Verdonck, S., Meers, K., Loossens, T., and Tuerlinckx, F. (2019). Prepaid parameter estimation without likelihoods. *PLoS computational biology*, 15(9):e1007181.

Noppeney, U., Ostwald, D., and Werner, S. (2010). Perceptual decisions formed by accumulation of audiovisual evidence in prefrontal cortex. *J Neurosci*, 30(21):7434–46.

O'Connell, R. G., Dockree, P. M., and Kelly, S. P. (2012). A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. *Nature Neuroscience*, 15(12):1729–1735.

Ossmy, O., Moran, R., Pfeffer, T., Tsetsos, K., Usher, M., and Donner, T. H. (2013). The timescale of perceptual evidence integration can be adapted to the environment. *Current Biology*, 23(11):981–986.

Palmer, J., Huk, A. C., and Shadlen, M. N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *J Vis*, 5(5):376–404.

Philiastides, M. G., Auksztulewicz, R., Heekeren, H. R., and Blankenburg, F. (2011). Causal role of dorsolateral prefrontal cortex in human perceptual decision making. *Current Biology*, pages 1–4.

Philiastides, M. G., Heekeren, H. R., and Sajda, P. (2014a). Human scalp potentials reflect a mixture of decision-related signals during perceptual choices. *J Neurosci*, 34(50):16877–89.

Philiastides, M. G., Heekeren, H. R., and Sajda, P. (2014b). Human scalp potentials reflect a mixture of decision-related signals during perceptual choices. *The Journal of Neuroscience*, 34(50):16877–16889.

Pisauro, M. A., Fouragnan, E., Retzler, C., and Philiastides, M. G. (2017). Neural correlates of evidence accumulation during value-based decisions revealed via simultaneous eeg-fmri. *Nature Comms*, 8:15808.

Pleskac, T. J. and Busemeyer, J. R. (2010). Two-stage dynamic signal detection: A theory of choice, decision time, and confidence. *Psychological Review*, 117(3):864–901.

Ploran, E. J., Tremel, J. J., Nelson, S. M., and Wheeler, M. E. (2011). High quality but limited quantity perceptual evidence produces neural accumulation in frontal and parietal cortex. *Cereb Cortex*, 21:2650–2662.

Polania, R., Krajbich, I., Grueschow, M., and Ruff, C. C. (2014). Neural oscillations and synchronization differentially support evidence accumulation in perceptual and value-

based decision making. *Neuron*, 82(3):709–20.

Rae, B., Heathcote, A., Donkin, C., Averell, L., and Brown, S. (2014). The hare and the tortoise: Emphasizing speed can change the evidence used to make decisions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 40(5):1226–1243.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychol Rev*, 85(2):59–108.

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

Ratcliff, R., Philiastides, M. G., and Sajda, P. (2009). Quality of evidence for perceptual decision making is indexed by trial-to-trial variability of the eeg. *Proc Natl Acad Sci U S A*, 106(16):6539–44.

Resulaj, A., Kiani, R., Wolpert, D. M., and Shadlen, M. N. (2009). Changes of mind in decision-making. *Nature*, 461(7261):263–266.

Rochet, N., Spieser, L., Casini, L., Hasbroucq, T., and Burle, B. (2014). Detecting and correcting partial errors: Evidence for efficient control without conscious access. *Cogn Affect Behav Neurosci*, 14:970–982.

Sajda, P., Philiastides, M. G., and Parra, L. C. (2009). Single-trial analysis of neuroimaging data: inferring neural networks underlying perceptual decision-making in the human brain. *IEEE Rev Biomed Eng*, 2:97–109.

Schurger, A., Sitt, J. D., and Dehaene, S. (2012). An accumulator model for spontaneous neural activity prior to self-initiated movement. *Proceedings of the National Academy of Sciences*, 109(42):E2904–E2913.

Scott, B. B., Constantinople, C. M., Akrami, A., Hanks, T. D., Brody, C. D., and Tank, D. W. (2017). Fronto-parietal cortical circuits encode accumulated evidence with a diversity of timescales. *Neuron*, pages 1–20.

Servant, M., White, C., Montagnini, A., and Burle, B. (2015). Using covert response activation to test latent assumptions of formal decision-making models in humans. *J Neurosci*, 35:10371–10385.

Shadlen, M. N. and Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area lip) of the rhesus monkey. *J Neurophysiol*, 86(4):1916–36.

Spieser, L., Servant, B., Hasbroucq, T., and Burle, B. (2017). Beyond decision! motor contribution to speed-accuracy trade-off in decision-making. *Psychon Bull Rev*, 24:950–956.

Storn, R. and Price, K. (1997). Differential evolution–a simple and efficient heuristic for global optimization over continuous spaces. *Journal of global optimization*, 11(4):341–359.

Tosoni, A., Galati, G., Romani, G. L., and Corbetta, G. W. (2008). Sensory-motor mechanisms in human parietal cortex underlie arbitrary visual decisions. *Nature Neuroscience*, 11(12):1446–1453.

Usher, M. and McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, 108(3):550–592.

van den Berg, R., Anandalingam, K., Zylberberg, A., Kiani, R., Shadlen, M. N., and Wolpert, D. M. (2016). A common mechanism underlies changes of mind about decisions and confidence. *eLife*, 5(12192).

Verdonck, S. and Tuerlinckx, F. (2014). The Ising Decision Maker: A binary stochastic network for choice response time. *Psychological Review*, 121(3):422–462.

Verdonck, S. and Tuerlinckx, F. (2015). Factoring out non-decision time in choice RT data: Theory and implications. *Psychological review.*

Voss, A. and Voss, J. (2007). Fast-dm: A free program for efficient diffusion model analysis. *Behavior Research Methods*, 39(4):767–775.

Voss, A. and Voss, J. (2008). A fast numerical algorithm for the estimation of diffusion model parameters. *Journal of Mathematical Psychology*, 52(1):1–9.

Wagenmakers, E.-J., Ratcliff, R., Gomez, P., and McKoon, G. (2008). A diffusion model account of criterion shifts in the lexical decision task. *Journal of Memory and Language*, 58(1):140–159.

Wenzlaff, H., Bauer, M., Maess, B., and Heekeren, H. R. (2011). Neural Characterization of the Speed–Accuracy Tradeoff in a Perceptual Decision-Making Task. *Journal of Neuroscience*, 31(4):1254–1266.

Wood, S. N. (2010). Statistical inference for noisy nonlinear ecological dynamic systems. *Nature*, 466(7310):1102–1104.

Wu, Z., Litwin-Kumar, A., Shamash, P., Taylor, A., Axel, R., and Shadlen, M. N. (2019). Context-dependent decision making in a premotor circuit. *bioRxiv*, page 757104.

Wyart, V., de Gardelle, V., Scholl, J., and Summerfield, C. (2012). Rhythmic fluctuations in evidence accumulation during decision making in the human brain. *Neuron*, 76(4):847–58.

Zhang, J. and Rowe, J. B. (2014). Dissociable mechanisms of speed-accuracy tradeoff during visual perceptual learning are revealed by a hierarchical drift-diffusion model. *Frontiers in Neuroscience*, 8.

Appendix I

Integral A

$$\int_{-\infty}^{t} dt' e^{\lambda(t'-t)} \ = \ e^{-\lambda t} \int_{-\infty}^{t} dt' e^{\lambda t'}$$

$$= \ e^{-\lambda t} \left[\frac{e^{\lambda t'}}{\lambda} \right]_{-\infty}^{t}$$

$$= \ \frac{e^{-\lambda t}}{\lambda} (e^{\lambda t} - 0)$$

$$= \ \frac{1}{\lambda}$$

Integral B

$$\int_0^t dt' t' e^{\lambda(t'-t)} = e^{-\lambda t} \int_0^t dt' t' e^{\lambda t'}$$

$$= e^{-\lambda t} \left[\frac{e^{\lambda t'}(\lambda t' - 1) + 1}{\lambda^2} \right]_0^t$$

$$= \frac{1}{\lambda^2}(\lambda t - 1) + \frac{1}{\lambda^2}e^{-\lambda t}$$

$$= \frac{1}{\lambda}(t - \frac{1}{\lambda} + \frac{e^{-\lambda t}}{\lambda})$$

Integral C

Stochastic process C in Equation 5 is a weighted time-integral of a Wiener process, which is normally distributed at every time point, with mean zero. We can calculate its variance at every time point.

In order to do this, we first calculate an intermediary result:

$$d(e^{\lambda(s-t)}W_s) = \lambda e^{\lambda(s-t)}W_s ds + e^{\lambda(s-t)}dW_s$$

$$\frac{1}{\lambda}d(e^{\lambda(s-t)}W_s) = e^{\lambda(s-t)}W_s ds + \frac{1}{\lambda}e^{\lambda(s-t)}dW_s$$

$$e^{\lambda(s-t)}W_s ds = \frac{1}{\lambda}d(e^{\lambda(s-t)}W_s) - \frac{1}{\lambda}e^{\lambda(s-t)}dW_s$$

Integrating this result gives us integral C in Equation 5:

$$\int_0^t ds e^{\lambda(s-t)}W_s = \left[\frac{1}{\lambda}e^{\lambda(s-t)}W_s \right]_0^t - \frac{1}{\lambda}\int_0^t dW_s e^{\lambda(s-t)}$$

$$= \frac{W_t}{\lambda} - \frac{1}{\lambda}e^{-\lambda t}W_0 - \frac{1}{\lambda}\int_0^t dW_s e^{\lambda(s-t)}$$

$$= \frac{W_t}{\lambda} - \frac{1}{\lambda}\int_0^t dW_s e^{\lambda(s-t)}$$

$$= \frac{1}{\lambda}\int_0^t dW_s \left\{ 1 - e^{\lambda(s-t)} \right\}$$

Using the Itô isometry,

$$E\left[\left(\int_0^t H_s dW_s \right)^2 \right] = E\left[\int_0^t H_s^2 ds \right],$$

we can calculate the variance of the process at time t:

$$\frac{1}{\lambda^2} \int_0^t ds \left\{1 - e^{\lambda(s-t)}\right\}^2 = \frac{1}{\lambda^2} \int_0^t ds(1 + e^{2\lambda(s-t)} - 2e^{\lambda(s-t)})$$

$$= \frac{1}{\lambda^2} \left([s]_0^t + \left[\frac{e^{2\lambda(s-t)}}{2\lambda}\right]_0^t - 2\left[\frac{e^{\lambda(s-t)}}{\lambda}\right]_0^t\right)$$

$$= \frac{1}{\lambda^2} \left(t + \frac{1 - e^{-2\lambda t}}{2\lambda} - \frac{2 - 2e^{-\lambda t}}{\lambda}\right)$$

$$= \frac{1}{\lambda^2} \left(t + \frac{1}{2\lambda} - \frac{e^{-2\lambda t}}{2\lambda} - \frac{2}{\lambda} + \frac{2e^{-\lambda t}}{\lambda}\right)$$

$$= \frac{1}{\lambda^2} \left(t + \frac{1 - 4}{2\lambda} + \frac{4e^{-\lambda t} - e^{-2\lambda t}}{2\lambda}\right)$$

$$= \frac{1}{\lambda^2} \left(t - \frac{3}{2\lambda} + \frac{4e^{-\lambda t} - e^{-2\lambda t}}{2\lambda}\right)$$

Appendix II

parameter	value
σ	0.1
a	0.1
x_0	0
v_i, with $i = 1, \ldots, 3$	$0, 0.1, 0.2$
T_{er}	0

Table 2

Parameters for the LIT extension of the constant drift diffusion model as used for the plots in Figure 1. Boundary separation $a = 0.1$ is used on the level of the LIT and corresponds to an upper boundary of 0.05 and a lower boundary of -0.05.

parameter	value
c	0.1
γ	1
κ	10
$v_{1,i}$, with $i = 1, \ldots, 3$	$0.15, 0.2, 0.3$
$v_{2,i}$, with $i = 1, \ldots, 3$	$0.15, 0.1, 0$
T_{er}	0

Table 3

Parameters for the LIT extension of the LCA as used for the plots in Figure 13 (notation according to Verdonck and Tuerlinckx, 2014). Single boundary thresholds of 0.08 are used on the level of the LITs.

parameter	value
β	$\frac{1}{28}$
D	0.05
N	8000
W^+	210000
W^-	33600
Θ	205800
$B_{1,i}$, with $i = 1, \ldots, 3$	$11000, 11200, 11400$
$B_{2,i}$, with $i = 1, \ldots, 3$	$11000, 10800, 10600$
T_{er}	0

Table 4

Parameters for the LIT extension of the IDM as used for the plots in Figure 14 (notation according to Verdonck and Tuerlinckx, 2014). We use a spontaneous relaxation time of 2 seconds. Single boundary thresholds of 0.5 are used on the level of the LITs.

Appendix III

The prepaid method as described in Mestdagh et al. (2019), application 3, internally uses a special parametrization of diffusion models where a free time-scaling parameter is added while the boundary separation is fixed to some constant. Exactly the same choice RT distributions can be reached as before, but the time-scaling parameter has the advantage that it can be estimated analytically (conditional on the other parameters), effectively reducing the amount of parameters that have to be accounted for in the prepaid database. The prepaid grids are cast as follows (for fixed boundary separation 0.1 and a free time-scaling parameter): starting position relative to boundary separation zr is uniformly distributed between 0.1 and 0.9 for all models (in this notation $zr = 0.5$ means no bias), inverse leak λ^{-1} is uniformly distributed between between 0 and 4 (only for LIT), inter-trial variability on relative starting point szr is uniformly distributed between 0 and 1 (only for RDM), and inter-trial variability on drift speed is uniformly distributed between 0 and 0.5 (only for RDM). For all three models, 100 drift speeds are incorporated per grid point, covering choice probabilities from 0.001 to 0.999. For the recovery study, parameters are drawn from the same distributions that are used for the creation of the prepaid grids, with a time-scale parameter uniformly drawn between 0.25 and 1.75 and a fixed non-decision time uniformly drawn between 0 and 0.3 seconds. The effective parameter ranges that this results in - in terms of the parametrizations used in this paper - can be assessed on the scatter plots in Figure 2, Figure 15, and Figure 16, for LIT, DDM and RDM respectively.
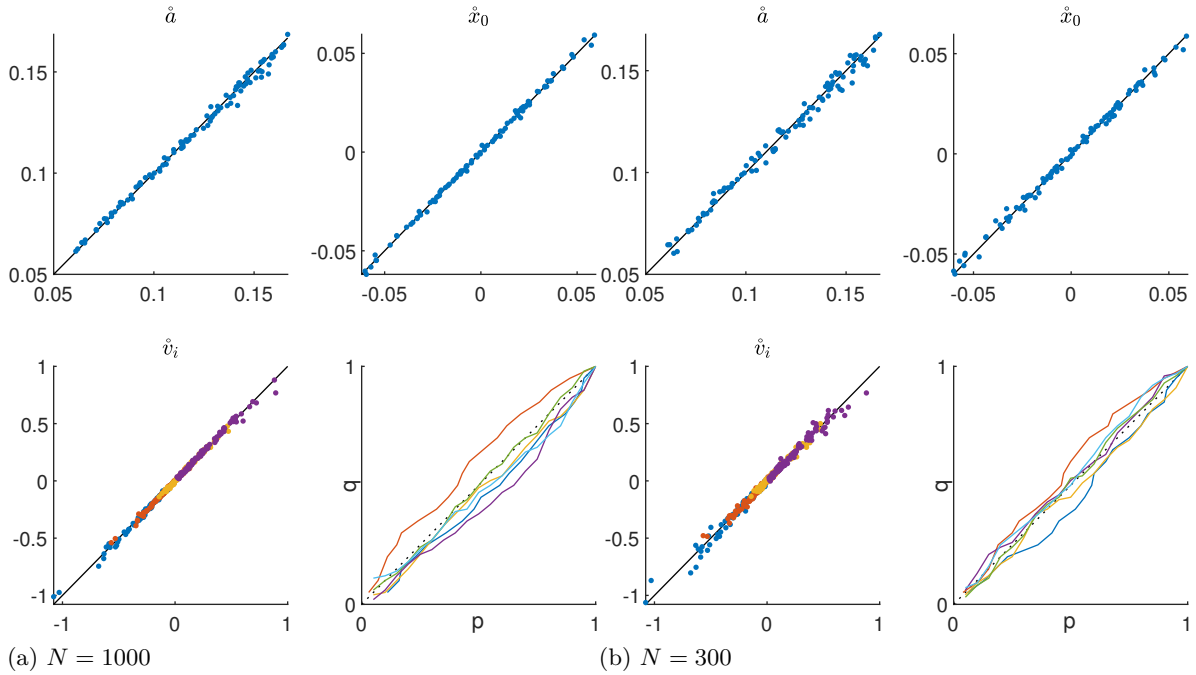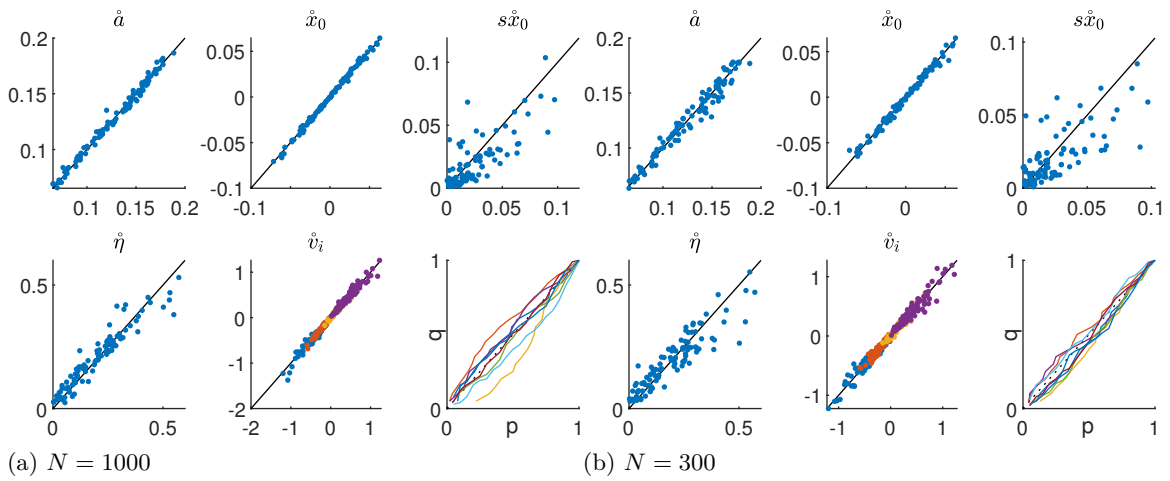
(a) $N = 1000$ (b) $N = 300$

Figure 15. Scatter plots of original versus recovered parameters for the DDM. Parameters are respectively boundary separation $\mathring{a}$, starting position $\mathring{x}_0$ and four drift speeds $\mathring{v}_i$. The last graph represents the coverage of the bootstrapped confidence intervals. The horizontal axis indicates the confidence interval $[p, 1-p]$, the vertical axis shows the fraction q of the original parameter values falling outside of their estimated confidence interval. Perfect coverage manifests as a straight diagonal line for each parameter. Plot (a) shows the recovery details for $N = 1000$ trials per stimulus condition, plot (b) for $N = 300$.

(a) $N = 1000$ (b) $N = 300$

Figure 16. Scatter plots of original versus recovered parameters for the RDM. Parameters are respectively boundary separation $\mathring{a}$, starting position $\mathring{x}_0$, variability of starting position $s\mathring{x}_0$, variability of drift speed $\mathring{\eta}$ and four drift speeds $\mathring{v}_i$. The last graph represents the coverage of the bootstrapped confidence intervals. The horizontal axis indicates the confidence interval $[p, 1 - p]$, the vertical axis shows the fraction q of the original parameter values falling outside of their estimated confidence intervals. Perfect coverage manifests as a straight diagonal line for each parameter. Plot (a) shows the recovery details for $N = 1000$ trials per stimulus condition, plot (b) for $N = 300$.

Appendix IV

The prepaid method as described in Mestdagh et al. (2019), application 3, can readily be used for generic diffusion models of choice RT where drift speed (or stimulus coherence) is the only experimental condition dependent parameter. To determine if either leak or boundary separation is the best single parameter explanation of the SAT manipulation, the method needs to be extended to allow other cross-conditional constraints and the following procedure is used: for every participant, the speed and accuracy data are estimated separately and two sets of non-parametric bootstrap estimates are obtained (both with 100 resamplings). We now have two overlapping clouds of parameter locations that we can use to create a new prepaid grid with either leak or boundary separation being free to change across SAT conditions. Practically, we take all possible pairs of bootstrap estimates - one estimate from each cloud - and each time we average the parameters that are shared across SAT conditions, while keeping the designated free SAT parameter separate for the two conditions. This results in a new parameter grid that has one extra parameter dimension (100x100 points, now covering two 2 SAT conditions). The choice RT distributions connected to this grid are then simulated at an accuracy of 10000 trials per stimulus-SAT condition. Based on this new prepaid grid the final estimate is done and resulting fits are compared. This procedure is done separately for every single participant and for both free parameter proposals (leak and boundary separation respectively).