Fitting drift-diffusion decision models to trial-by-trial data

Q. Feltgen[1], *J. Daunizeau*[1,2]

[1] Brain and Spine Institute, Paris, France

[2] ETH, Zurich, Switzerland

Address for correspondence:

Jean Daunizeau

Motivation, Brain and Behaviour Group

Brain and Spine Institute (ICM), INSERM U1127.

47, bvd de l'Hopital, 75013, Paris, France.

Tel: +33 1 57 27 43 26

Mail: jean.daunizeau@gmail.com

Web: https://sites.google.com/site/jeandaunizeauswebsite/

Abstract

Drift-diffusion models or DDMs are becoming a standard in the field of computational neuroscience. They extend models from signal detection theory by proposing a simple mechanistic explanation for the observed relationship between decision outcomes and reaction times (RT). In brief, they assume that decisions are triggered once the accumulated evidence in favor of a particular alternative option has reached a predefined threshold. Fitting a DDM to empirical data then allows one to interpret observed group or condition differences in terms of a change in the underlying model parameters. However, current approaches do not provide reliable parameter estimates when, e.g., evidence strength is varied over trials. In this note, we propose a fast and efficient approach that is based on fitting a self-consistency equation that the DDM fulfills. Using numerical simulations, we show that this approach enables one to extract relevant information from trial-by-trial variations of RT data that would typically be buried in the empirical distribution. Finally, we demonstrate the added-value of the approach, when applied to a recent value-based decision making experiment.

Over the past two decades, neurocognitive processes of decision making have been extensively studied under the framework of so-called *drift-diffusion models* or DDMs. These models tie together decision outcomes and reaction times (RT) by assuming that decisions are triggered once the accumulated evidence in favor of a particular alternative option has reached a predefined threshold (Ratcliff and McKoon, 2008; Ratcliff et al., 2016). They owe their popularity both to experimental successes in capturing observed data in a broad set of behavioral studies (De Martino et al., 2012; Gold and Shadlen, 2007; Hanks et al., 2014; Milosavljevic et al., 2010; Pedersen et al., 2017; Resulaj et al., 2009), and to theoretical work showing that DDMs can be thought of as optimal inference problem solvers (Drugowitsch et al., 2012; Tajima et al., 2016). Now current decision making experiments typically consider situations in which decision-relevant variables are manipulated on a trial-by-trial basis. For example, the reliability of perceptual evidence (e.g., the psychophysical contrast in a perceptual decision) is varied from one trial to the next. Under current applications of the DDM, this implies that some internal model variables (e.g., the drift rate, see below) effectively vary over trials. Nevertheless, current statistical techniques are limited in their ability to exploit predictable inter-trial variations of DDM variables to fit the model to RT data. The reason is at least twofold. First, DDMs suffer from inherent non-identifiability issues, which are difficult to anticipate from the basic form of the model (Ratcliff and Tuerlinckx, 2002). Second, the DDM is typically fitted to the empirical distribution of reaction time data (conditional on each alternative decision outcome). This is because the information that is deemed relevant for estimating DDM parameters lies in the moments of the RT distribution, whose empirical estimate requires many trial repetitions. However, we will see that the RT distribution is blind to most trial-by-trial variations of DDM parameters.

This note is concerned with the issue of obtaining reliable parameter estimates for drift-diffusion models from concurrent trial-by-trial choice and reaction time data. We propose a fast and efficient approach that is based on fitting a self-consistency equation, which is derived from a simple approximation to the DDM. In particular, we show that this approach

enables one to extract relevant information from trial-by-trial variations of RT data that would typically be buried in the empirical distribution.

In the first section of this document, we briefly recall the derivation of the DDM, and summarize the impact of DDM parameters onto the conditional RT distributions. In the second and third sections, we derive the DDM's self-consistency equation and summarizes the ensuing overcomplete approach to DDM-based data analysis. In the fourth section, we summarize the way in which trial-by-trial variations can be accounted for. In the fifth section, we perform parameter recovery analyses for both the standard DDM fitting procedure and the overcomplete approach. In the sixth section, we demonstrate the added-value of the overcomplete approach, when applied to a value-based decision making experiment. In the seventh and last section, we discuss our results in the context of the existing literature.

1. Model formulation and impact of DDM parameters

First, let us recall the typical form of a drift-diffusion model or DDM. In brief, the decision variable $x(t)$ is supposed to follow the following stochastic differential equation:

$$dx = v \times dt + \sigma \times \sqrt{dt} \times d\eta(t) \quad (1)$$

where v is the drift rate, $d\eta(t)$ is a standard Wiener process, and σ is the standard deviation of the stochastic (diffusion) perturbation term.

In discrete time, Equation 1 can be rewritten as follows:

$$x_{t+1} = x_t + v + \sigma\eta_t \quad (2)$$

where t indexes time (assuming unitary sampling rate for the sake of simplicity) and $\eta_t \sim N(0,1)$ is a standard normal random variable. By convention, the system's initial condition is denoted as x_0, which we refer to as the "initial bias".

By assumption, a decision is triggered whenever the stochastic process x_t hits either of two bounds ($x*$ or $-x*$). When a bound hit occurs defines the "hitting time" (HT), and which bound is hit defines the (binary) decision outcome o. Both depend upon the DDM parameters, which include: the drift rate v, the threshold's height $x*$, the noise's standard deviation σ and the initial bias x_0. Note that typical DDMs also include a so-called "non-decision" time parameter T_{ND}, which captures non-specific delays between hitting times and overt reaction times.

Fitting the DDM can be done if the joint distribution of reaction times and decision outcomes is available, because the model parameters control the sufficient statistics of this distribution. This is typically done by matching observed and predicted sufficient statistics (e.g., distribution quantiles or moments) of response times for both decision outcomes. Figures 1 to 4 below demonstrate the impact of model parameters on the decision outcome ratio $P\left(o\middle|v,x_0,x*,\sigma\right)$ and the first three moments of conditional hitting time distributions, namely: their mean $E\left[HT\middle|o,v,x_0,x*,\sigma\right]$, variance $V\left[HT\middle|o,v,x_0,x*,\sigma\right]$ and skewness $S_k\left[HT\middle|o,v,x_0,x*,\sigma\right]$ As we will see, each DDM parameter has a specific signature, in terms of its joint impact on these seven quantities. This does not imply however, that different parameter settings necessarily yield distinct moments. In fact, there are changes in the DDM parameters that leave the predicted moments unchanged. This will induce parameter recovery issues, which we will demonstrate later.

But first, let us summarize the impact of DDM parameters. To do this, we first set model parameters to the following "default" values: $v=1/2$, $x_0=1$, $x*=10$ and $\sigma=4$. Then, we vary each model parameter one by one, keeping the other ones at their default value.

Figure 1 below shows the impact of initial bias x_0.

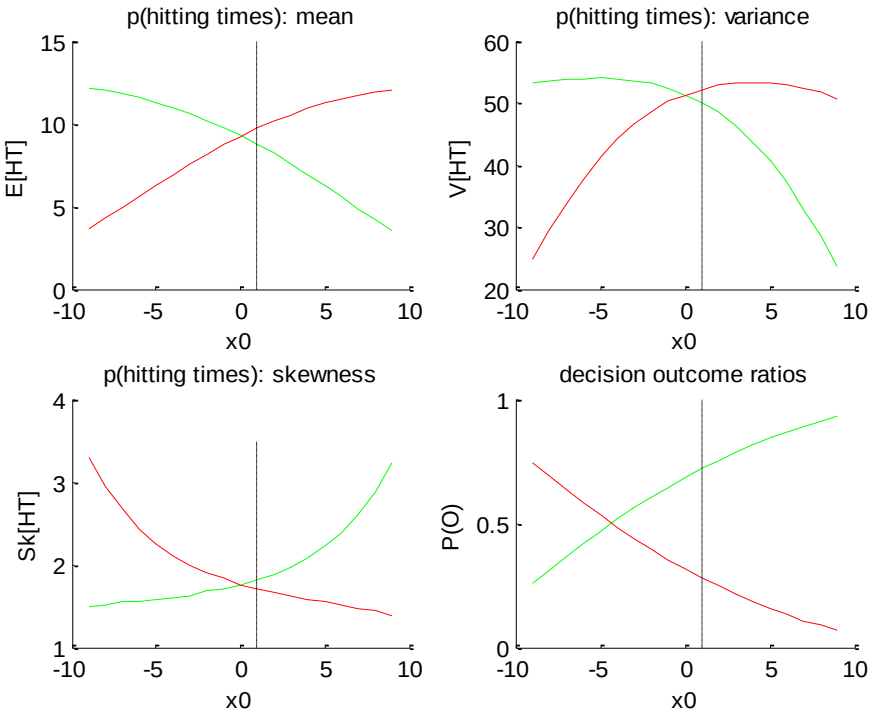

Figure 1: impact of initial bias x_0. In all panels, the colour code indicates the decision outcomes (green: 'up' decisions, red: 'down' decisions). The black dotted line indicates the default parameter value (for ease of comparison with other figures below). **Upper-left panel**: mean hitting times (y-axis) as a function of initial bias (x-axis). **Upper-right panel**: hitting times' variance (y-axis) as a function of initial bias (x-axis). **Lower-left panel**: hitting times' skewness (y-axis) as a function of initial bias (x-axis). **Lower-right panel**: outcome ratios (y-axis) as a function of initial bias (x-axis).

One can see that increasing the initial bias accelerates decision times for 'up' decisions, and decelerates decision times for 'down' decisions. This is because increasing x_0 mechanically increases the probability of an early upper bound hit, and decreases the probability of an early lower bound hit. Increasing x_0 also decreases (resp., increases) the variance for 'up' (resp., 'down') decisions, and increases (resp., decreases) the skewness for 'up' (resp., 'down') decisions. Finally, increasing the initial bias increases the ratio of 'up' decisions. These are corollary consequences of increasing (resp. decreasing) the probability of an early upper (resp., lower) bound hit. This is because when an increasing proportion of stochastic paths eventually hit a bound very early, this squeezes the distribution of hitting times just above null hitting times. Note that the outcome ratios are not equal to $1/2$ when $x_0 = 0$. This

is because the default drift rate v is positive, and therefore favours 'up' decisions. Most importantly, the initial bias is the only DDM parameter that has opposite effects on HT for 'up' and 'down' decision outcomes.

Figure 2 below shows the impact of drift rate v.

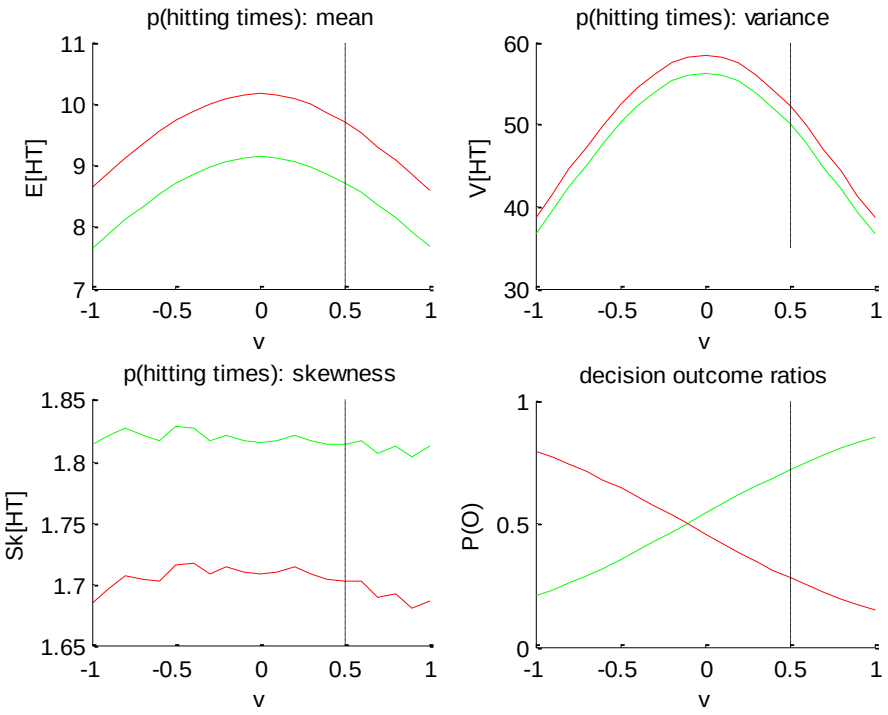

Figure 2: impact of drift rate v. Same format as Fig 1.

One can see that the mean and variance of decision times are maximal when the drift rate is null. This is because the probability of an early (upper or lower) bound hit decreases as $v \to 0$. Also, the drift rate has no impact on the HT skewness (visible variations on the lower-left panel of Figure 2 are due to numerical imprecisoins). In fact, it is the only parameter that does not change HT skewness. Note that, in contrast to the initial bias, the impact of the drift rate on HT is identical for both 'up' and 'down' decisions. Finally, and as expected, increasing the drift rate increases the ratio of 'up' decisions.

Figure 3 below shows the impact of the noise's standard deviation σ.

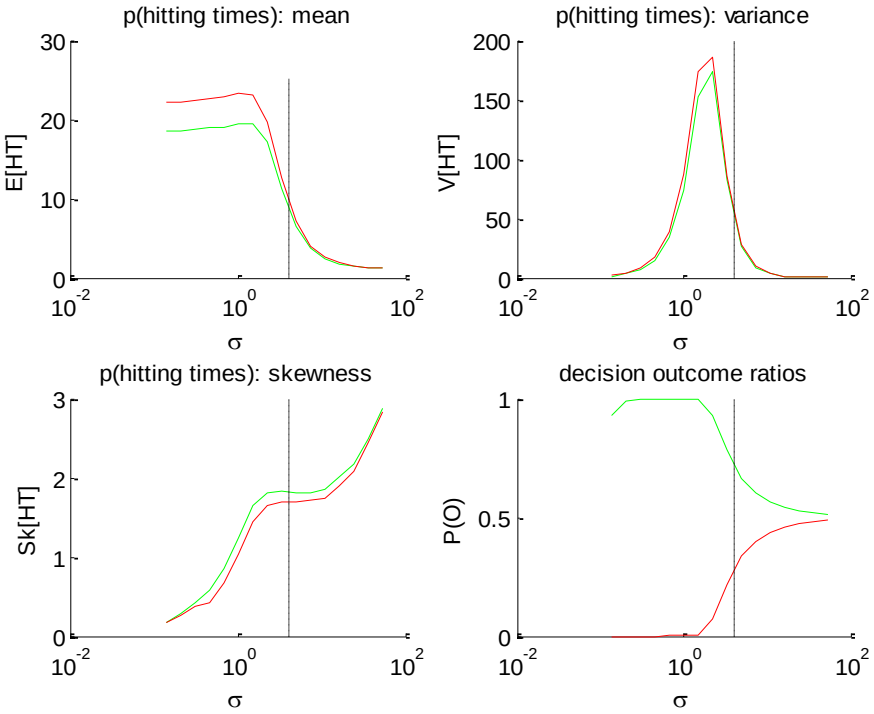

Figure 3: impact of the perturbation' standard deviation σ . Same format as Fig 1 (but the x-axis is now in log-scale).

One can see that increasing the standard deviation decreases the mean HT, and increases its skewness. This is, again, because increasing σ increases the probability of an early bound hit. Its impact on the variance, however, is less trivial. When the standard deviation σ is very low, increasing σ first increases the hitting times' variance. This is because it progressively frees the system from its deterministic fate, therefore enabling HT variability around the mean. Then, it reaches a critical point above which increasing it further now decreases the variance. This is again a consequence of increasing the probability of an early bound hit. The associated HT squeezing effect can be seen on the skewness, which steadily increases beyond the critical point. Note that the standard deviation has the same impact on HT for 'up' and 'down' decisions. Finally, increasing the standard deviation eventually maximizes the entropy of the decision outcomes, i.e. $P(o) \to 1/2$ when $\sigma \to \infty$. This is because the relative contribution of the diffusion term eventually masks the drift.

Figure 4 below shows the impact of the bound's height x^*.

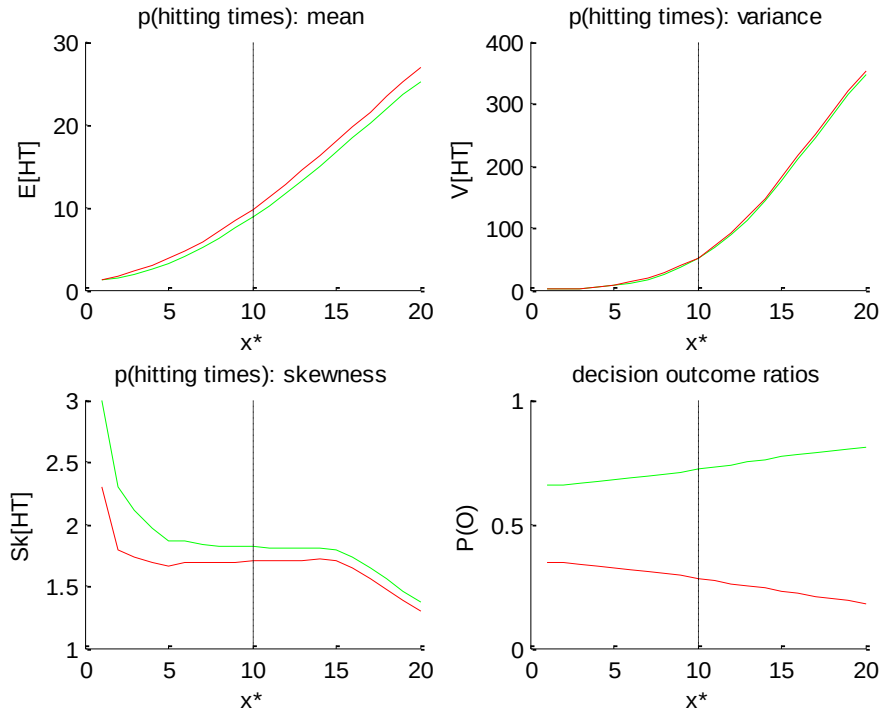

Figure 4: impact of the threshold's height x^*. Same format as Fig 1.

One can see that increasing the bound's height increases both the mean and the variance of HT, and decreases its skewness, identically for 'up' and 'down' decisions. Finally, increasing the threshold's height decreases the entropy of the decision outcomes, i.e. $P(o) \to 0 \text{ or } 1$ when $x^* \to \infty$. This directly derives from the fact that increasing x^* decreases the probability of an early bound hit. This effect basically competes with the effect of the standard deviation σ, which accelerates HTs. This is why one may say that increasing the threshold's height effectively increases the demand for evidence strength in favour of one of the decision outcomes.

Note that the impact of the "non-decision" time T_{ND} simply reduces to shifting the mean of the RT distribution, without any effect on other moments.

In brief, DDM parameters have distinct impacts on the sufficient statistics of hitting times. This means that they could, in principle, be discriminated from each other. Standard DDM fitting procedures rely on adjusting the DDM parameters so that the observed empirical moments (e.g., up to third order) match the model predictions. In what follows, we refer to this as the "method of moments". However, we will see below that the DDM is not perfectly identifiable. One can also see that changing these parameters from trial to trial will most likely induce non-trivial variations in RT data. Below, we suggest a simple way of exploiting such trial-by-trial variations to perform parameter estimation. This follows from fitting a self-consistency equation that derives from a simple approximation to the above DDM.

2. A self-consistency equation for DDM

First, note that Equation 2 can be rewritten as follows:

$$\begin{aligned} x_t &= x_0 + tv + \sigma \sum_{t'=0}^{t-1} \eta_{t'} \\ &= x_0 + tv + \sigma \sqrt{t}\,\tilde{\eta} \end{aligned} \tag{3}$$

where $\tilde{\eta} \triangleq \frac{1}{\sqrt{t}} \sum_{t'=0}^{t-1} \eta_{t'}$ is a rescaled cumulative perturbation. By construction, $\tilde{\eta}$ is a standard normal random variable, with mean 0 and unit variance, i.e. $\tilde{\eta} \sim N(0,1)$, irrespective of the temporal position on the system's trajectory. This is why, in what follows, we refer to it as the "normalized cumulative perturbation"'. Now let τ_i be the decision time of the i^{th} trial. Note that τ_i is trivially related to $\tilde{\eta}_i$ because, by definition, $\left| x^i_{\tau_i} \right| = x^*$. This implies that:

$$x^* = \left| x_0 + \tau_i v + \sigma \sqrt{\tau_i}\,\tilde{\eta}_i \right| \tag{4}$$

where $\tilde{\eta}_i$ denotes the (unknown) cumulative perturbation term of the i^{th} trial.

Information regarding the binary decision outcome $o_i \in \{-1,1\}$ further disambiguates Equation 4 as follows:

$$\begin{aligned} x^* &= \begin{cases} x_0 + \tau_i v + \sigma\sqrt{\tau_i}\tilde{\eta}_i & \text{if } o_i = 1 \quad \text{('up' decision)} \\ -x_0 - \tau_i v - \sigma\sqrt{\tau_i}\tilde{\eta}_i & \text{if } o_i = -1 \text{ ('down' decision)} \end{cases} \\ &= o_i\left(x_0 + \tau_i v + \sigma\sqrt{\tau_i}\tilde{\eta}_i\right) \end{aligned} \tag{5}$$

Equation 5 can be used to express decision time directly as a function of DDM model parameters:

$$\tau_i = \frac{o_i x^* - x_0}{v} - \frac{\sigma\sqrt{\tau_i}}{v}\tilde{\eta}_i \tag{6}$$

This provides an observation equation for empirical response times y_i, which are defined as noisy measures of decision times:

$$y_i \approx \frac{o_i x^* - x_0}{v} - \frac{\sigma\sqrt{y_i - T_{ND}}}{v}\tilde{\eta}_i + T_{ND} + \varepsilon_i \tag{7}$$

where ε_i are unknown i.i.d. noise residuals.

Note that reaction times appear on both the left-hand and the right-hand side of Equation 7. This is a slightly unorthodox feature, but, as we will see, this has effectively no consequence from the perspective of model inversion. In fact, one can think of Equation 7 as a "self-consistency" constraint that the DDM has to fulfil. This is why we refer to Equation 7 as the *self-consistency equation* of DDMs. This, however, prevents us from using Equation 7 to generate data under the DDM. In other terms, Equation 7 is only useful when analyzing empirical RT data.

3. An overcomplete likelihood approach to DDM inversion

Equation 7 can be used directly to fit the DDM parameters to trial-by-trial reaction time data. Note that this induces an "overcomplete" likelihood function, because here cumulative

perturbation terms $\tilde{\eta}_i$ are treated as nuisance model parameters, but model parameters nonetheless. This means that there are more model parameters than there are data points, hence the overcomplete likelihood function. Thus, one may want to constrain parameter estimation as much as possible. This can be done by noticing that one can set prior constraints on the distribution of cumulative perturbation terms $\tilde{\eta}_i$.

Recall that, under the DDM framework, errors can only be due to the stochastic perturbation noise. More precisely, errors are due to those perturbations that are strong enough to deviate the system's trajectory and make it hit the "wrong" bound (e.g., the lower bound if the drift rate is positive). Let $Q_=$ be the proportion of correct responses. For example, if the drift rate is positive, then $Q_=$ corresponds to responses that hit the upper bound. Now let $\tilde{\eta}_=$ be the critical value of $\tilde{\eta}_i$ such that $P\left(\tilde{\eta} > \tilde{\eta}_=\right) = Q_=$ (see Figure 5 below). Then, we know that errors correspond to those perturbations $\tilde{\eta}_i$ that are smaller than $\tilde{\eta}_=$. By construction, the marginal distribution $p\left(\tilde{\eta}\right) = N\left(0,1\right)$ of normalized cumulative is invariant over time. This enables us to derive the conditional expectations $\mu_=^0$ and $\mu_{\neq}^0$ of the perturbation $\tilde{\eta}_i$, given that the decision outcome o_i is correct or erroneous, respectively:

$$\begin{cases} \mu_=^0 \triangleq E\left[\tilde{\eta}_i \middle| o_i = 1\right] = E\left[\tilde{\eta}_i \middle| \tilde{\eta}_i > \tilde{\eta}_=\right] = \dfrac{1}{\left(1-Q_=\right)\sqrt{2\pi}} \exp\left(-\dfrac{1}{2}\tilde{\eta}_=^{\,2}\right) \\ \mu_{\neq}^0 \triangleq E\left[\tilde{\eta}_i \middle| o_i = -1\right] = E\left[\tilde{\eta}_i \middle| \tilde{\eta}_i < \tilde{\eta}_=\right] = -\dfrac{1}{Q_=\sqrt{2\pi}} \exp\left(-\dfrac{1}{2}\tilde{\eta}_=^{\,2}\right) \end{cases} \tag{8}$$

Equation 8 is obtained from the known expression of first-order moments of a truncated normal density. Critically, Equation 8 holds true, irrespective of the other DDM parameters (as long as $v > 0$). Of course, the same logic extends to conditional variances $\Sigma_=^0$ and $\Sigma_{\neq}^0$, whose analytical expressions are given by:

$$\begin{cases} \Sigma_=^0 \triangleq V\left[\tilde{\eta}_i \middle| o_i = 1\right] = V\left[\tilde{\eta}_i \middle| \tilde{\eta}_i > \tilde{\eta}_=\right] = 1 + \tilde{\eta}_=\mu_=^0 - \mu_=^{02} \\ \Sigma_{\neq}^0 \triangleq V\left[\tilde{\eta}_i \middle| o_i = -1\right] = V\left[\tilde{\eta}_i \middle| \tilde{\eta}_i < \tilde{\eta}_=\right] = 1 + \tilde{\eta}_=\mu_{\neq}^0 - \mu_{\neq}^{02} \end{cases} \tag{9}$$

A simple moment-matching approach thus suggests to set conditional prior distributions on each $\tilde{\eta}_i$ as follows:

$$p\left(\tilde{\eta}_i \middle| o_i\right) = \begin{cases} N\left(\mu_{=}^{0}, \Sigma_{=}^{0}\right) & \text{if } o_i = correct \\ N\left(\mu_{\neq}^{0}, \Sigma_{\neq}^{0}\right) & \text{if } o_i = error \end{cases} \quad (10)$$

where the correct/error label depends on the sign of the drift rate.

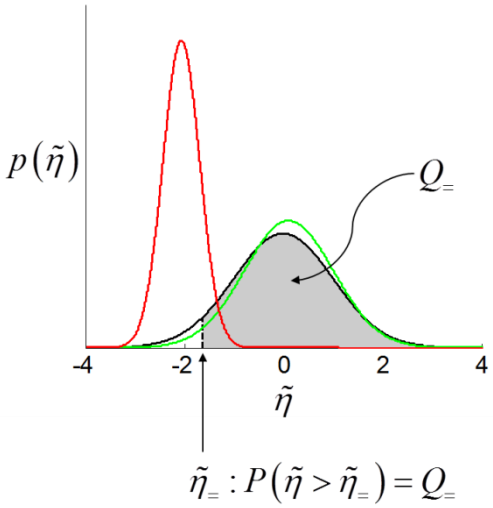

Figure 5: conditional prior distributions of the normalized cumulative perturbations. The black line shows the marginal distribution $p\left(\tilde{\eta}\right) = N\left(0,1\right)$, and the grey area's left bound is the critical value $\tilde{\eta}_{=}$ such that $P\left(\tilde{\eta} > \tilde{\eta}_{=}\right) = Q_{=}$. The green and red lines depict the conditional normal distributions $p\left(\tilde{\eta}_i \middle| o_i = correct\right)$ and $p\left(\tilde{\eta}_i \middle| o_i = error\right)$, respectively, that derive from matching the two first moments of the respective truncated gaussian densities.

When fitting the DDM to empirical RT data, one thus wants to enforce the distributional constraint in Equations 10. A simple way to do this is to replace $\tilde{\eta}$ with a non-constrained set of dummy variables ς, which we map through the following moment-enforcing transform:

$$\tilde{\eta}_i = \begin{cases} \mu_{=}^{0} + \left(\varsigma_i - \dfrac{1}{nQ_{=}} \sum_{i \in i_{=}} \varsigma_i \right) \sqrt{\dfrac{nQ_{=}\Sigma_{=}^{0}}{\sum_{i \in i_{=}} \left(\varsigma_i - \dfrac{1}{nQ_{=}} \sum_{i \in i_{=}} \varsigma_i \right)^2}} & \text{if } i \in i_{=} \\ \mu_{\neq}^{0} + \left(\varsigma_i - \dfrac{1}{n(1-Q_{=})} \sum_{i \in i_{\neq}} \varsigma_i \right) \sqrt{\dfrac{n(1-Q_{=})\Sigma_{\neq}^{0}}{\sum_{i \in i_{\neq}} \left(\varsigma_i - \dfrac{1}{n(1-Q_{=})} \sum_{i \in i_{\neq}} \varsigma_i \right)^2}} & \text{if } i \in i_{\neq} \end{cases} \tag{11}$$

where $i_{=}$ and $i_{\neq}$ are the indices of correct and incorrect trials, respectively (and n is the total number of trials). Equation 11 ensures that the sample moments of the estimated normalized cumulative perturbations match Equations 8-9, irrespective of the dummy variable ς. Thus, replacing the native normalized cumulative perturbations with the dummy variable ς enables one to rely on unconstrained probabilistic parameter estimation schemes. In this work, we use the variational Laplace approach (Daunizeau, 2017; Friston et al., 2007), which was developed to perform approximate bayesian inference on nonlinear generative models.

In addition, one want to introduce the following prior constraints on the native DDM parameters:

- The bound's height x^* is necessarily positive. This positivity constraint can be enforced by replacing x^* with a non-bounded parameter φ_1, which relates to the bound's height through the following mapping: $x^* = \exp(\varphi_1)$.
- The standard deviation σ is necessarily positive. Again, this can be enforced by replacing it with the following mapped parameter φ_2: $\sigma = \exp(\varphi_2)$.
- The non-decision time T_{ND} is necessarily smaller than the minimum observed reaction time. This can be enforced by replacing the native non-decision time with the following mapped parameter φ_3: $T_{ND} = \min(RT)\, s(\varphi_3)$, where $s(\bullet)$ is the standard sigmoid mapping.

- The initial bias x_0 is necessarily constrained between $-x^*$ and x^*. This can be enforced by replacing the native initial condition with the following mapped parameter φ_4: $x_0 = \exp(\varphi_1)\left(2s(\varphi_4)-1\right)$.
- In principle, the drift rate v can be either positive or negative. However, its magnitude is necessarily smaller than $\frac{x^* + |x_0|}{\min(RT) - T_{ND}}$. This can be enforced by replacing the native drift rate with the following mapped parameter φ_5: $v = \frac{\left(1+\left|2s(\varphi_4)-1\right|\right)\exp(\varphi_1)}{\min(RT)\left(1-s(\varphi_3)\right)}\left(2s(\varphi_5)-1\right)$.

Here again, we use the set of dummy variables $\varphi_{1:5}$ in lieu of native DDM parameters.

The statistical efficiency of the ensuing overcomplete approach can be evaluated by simulating RT and choice data under different settings of the DM parameters, and then comparing estimated and simulated parameters. Below, we use such recovery analysis to compare the overcomplete approach with the standard DDM fitting procedure (i.e., the method of moments). Note that both approaches benefit from the above prior constraints on native DDM parameters.

4. Accounting for predictable trial-by-trial RT variability

Critically, the above overcomplete approach can be extended to ask whether trial-by-trial variations in DDM parameters explain trial-by-trial variations in observed RT, above and beyond the impact of the random perturbation term in Equation 7. For example, one may want to assess whether predictable variations in e.g., the drift term, accurately predict variations in RT data. This kind of question underlies many recent empirical studies of human and/or animal decision making. In the context of perceptual decision making, the drift rate is

assumed to derive from the strength of momentary evidence, which is controlled experimentally and varies in a trial-by-trial fashion (Bitzer et al., 2014; Huk and Shadlen, 2005). A straightforward extension of this logic to value-based decisions implies that the drift rate should vary in proportion to the value difference between alternative options (De Martino et al., 2012; Krajbich et al., 2010; Lopez-Persem et al., 2016). In both cases, a prediction for drift rate variations across trials is available, which is likely to induce trial-by-trial variations in choice and RT data. This is most simply done by altering the self-consistency equation such that DDM parameters are treated as affine functions of trial-by-trial predictors. For example, if the drift rate is defined as $v \triangleq v_0 + v_1 D_i$, then Equation 7 becomes:

$$y_i \approx \frac{o_i x^* - x_0}{v_0 + v_1 D_i} - \frac{\sigma\sqrt{y_i - T_{ND}}}{v_0 + v_1 D_i}\tilde{\eta}_i + T_{ND} + \varepsilon_i \qquad (12)$$

where D_i is the predictor variable of drift rate at trial i, v_1 measures the strength of the linear relationship between the drift rate and the predictor , and v_0 is an offset term.

Alternatively, one can simply set the drift rate to the predictor variable (i.e. assume *a priori* $v_0 = 0$ and $v_1 = 1$). As we will see below, this significantly improves model identifiability for the remaining parameters. This is because trial-by-trial variations in the drift rate will only accurately predict trial-by-trial variations in reaction time data if the remaining parameters are correctly set. This is just an example of course, and one can see how easily any prior dependency to a predictor variable could be accounted for. The critical point here is that the overcomplete approach can exploit predictable trial-by-trial variations in RT data to constrain the inference on model parameters.

The situation is quite different for the standard DDM fitting procedure, which relies on the "method of moments". In brief, information regarding trial-by-trial variations cannot be easily conveyed to this type of approach, because it focuses on the RT distribution *over trials*. Practically speaking, applying the method of moments requires to gather trials into subsets according to some binning of the predictor variable. This trial partition can be problematic,

given that trials also have to be separated according to choice outcome. To ensure that all subsets contain a "reasonable" number of trials, one may thus have to rely on a coarse binning of the predictor variable. At the limit, all outcome-specific trials would be considered together, and trial-by-trial variations in the predictor variable would be absorbed in the moments of the corresponding conditional RT distributions. On the other hand, predictable trials-by-trial variations can be accounted for by reducing the model prediction to a single trial-by-trial summary statistics, e.g. the mean or mode of the conditional RT distribution. This does not require any trial partition, but incur an enormous computational cost as the number of trials grows. Given that we consider realistic applications of the DDM, we will not evaluate parameter recovery for such method. Now recall that, strictly speaking, accounting for variations in e.g., the drift rate v, over a given subset of n trials, reduces to marginalizing the conditional RT distribution over v:

$$\begin{aligned} p\left(HT\left|o,x_0,x^*,\sigma\right.\right) &= \int p\left(HT\left|o,v,x_0,x^*,\sigma\right.\right) p\left(v\right) dv \\ &\approx \frac{1}{n}\sum_{i=1}^{n} p\left(HT\left|o,v_i,x_0,x^*,\sigma\right.\right) \end{aligned} \tag{13}$$

where v_i is the predicted drift rate at trial i, and $p\left(HT\left|o,v,x_0,x^*,\sigma\right.\right)$ is the hitting times' distribution that is induced by the corresponding DDM parameter setting. One can see that the resulting marginal distribution $p\left(HT\left|o,x_0,x^*,\sigma\right.\right)$ will be blind to many forms of predictable trial-by-trial variations. To begin with, $p\left(HT\left|o,x_0,x^*,\sigma\right.\right)$ is invariant to any arbitrary permutation of the trials. In fact, in most cases, the marginal distribution $p\left(HT\left|o,x_0,x^*,\sigma\right.\right)$ closely resembles the conditional distribution $p\left(HT\left|o,\bar{v},x_0,x^*,\sigma\right.\right)$, where $\bar{v} = 1/n\sum_{i=1}^{n} v_i$ is the mean drift rate over the trial subset. Figure 6 below shows the impact of the trial-by-trial variance of drift rate on the moments of the marginal distribution $p\left(HT\left|o,x_0,x^*,\sigma\right.\right)$.

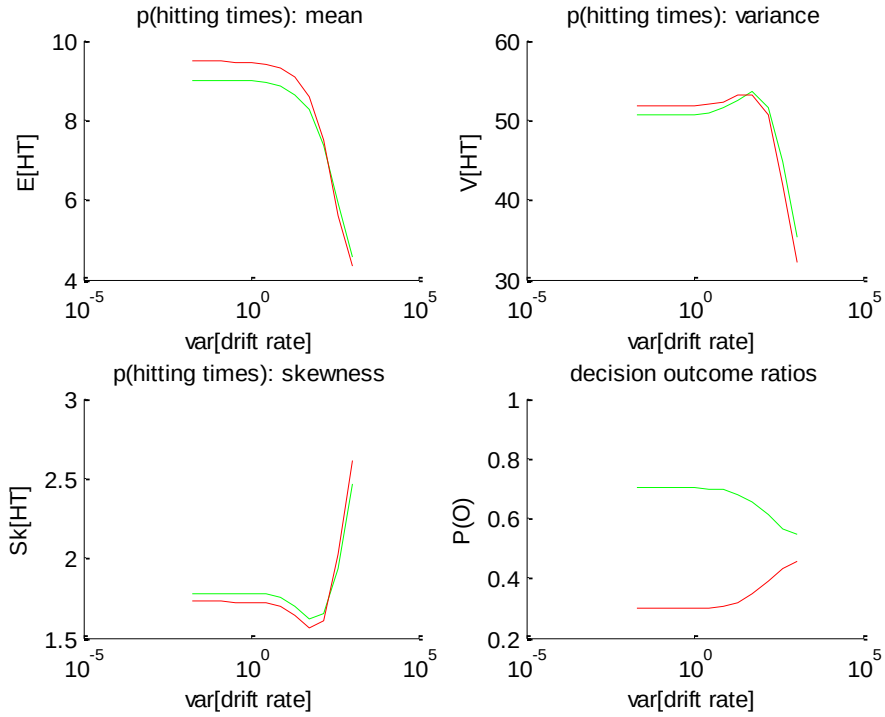

Figure 6: impact of the trial-by-trial variance of drift rate. Same format as Fig 1. Note that the mean drift rate was fixed to its default value, i.e.: $\bar{v} = 1/2$. The moments of the conditional distribution $p\left(HT \middle| o, \bar{v}, x_0, x^*, \sigma\right)$ can be eyeballed at the leftmost part of the panels (i.e. when the variance of the drift rate is close to zero). One can check that it corresponds to the moments shown at the location of the black dotted lines on all panels of Figures 1 to 4.

As one can see, the trial-by-trial variance of drift rate essentially has no impact on the marginal distribution $p\left(HT \middle| o, x_0, x^*, \sigma\right)$ as long as it is below or of the same order of magnitude than the mean drift rate. From this point onwards, it decreases the hitting time's mean and its variance, increases its skewness, and increases the entropy of the outcome probability. This effect is qualitatively identical to the impact of the noise's standard deviation σ (cf. all panels of Figure 3, on the right of the black dotted line). Practically speaking, Figure 6 implies that the method of moments will be blind to most trial-by-trial variations of drift rates that occur within trial subsets.

We note that this argument generalizes to trial-by-trial variations in other model parameters. But the insensitivity to trial-by-trial variations in the drift rate is important, because this is the DDM parameter that is typically assumed to vary in empirical applications of the DDM model.

5. Parameter recovery analysis

First, we compare the overcomplete approach and the method of moments w.r.t. their ability to recover DDM parameters, when no trial-by-trial variations are considered. Model inversion, for both the method of moments and the ovecomplete approach, was performed using the variational Laplace approximation to bayesian inference (Beal, 2003; Daunizeau, 2017; Friston et al., 2007) that is implemented in the opensource VBA toolbox (Daunizeau et al., 2014). In both cases, we inserted the prior constraints on DDM parameters (cf. Section 3), along with standard normal priors on mapped parameters $\varphi_{1:5}$.

In what follows, we summarize the results of parameter recovery analyses. These proceed as follows. First, we sample 1000 sets of model parameters $\varphi_{1:5}$ under some arbitrary distribution. Second, for each of these parameter, we simulate a series of N=200 DDM trials according to Equation 2 above. Third, we fit the DDM to each series of simulated reaction times (200 data points) using both approaches. Last, we ask whether variations in estimated parameters (across the 1000 Monte-Carlo simulations) faithfully capture variations in simulated parameters. We do this by regressing each estimated parameter against all simulated parameters. The ensuing regression coefficients quantify the concurrent influence simulated parameters have on each estimated parameter.

The recovery analysis of the full DDM is summarized in Figure 7 below, in terms of the standardized regression of estimated parameters against simulated parameters.

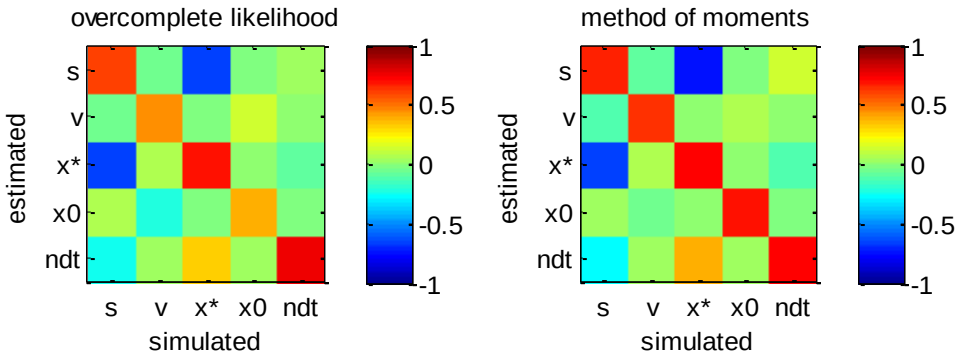

Figure 7: DDM parameter recovery analysis for the overcomplete likelihood method (left panel) and the method of moments (right panel). Each line shows how variations in a given estimated parameter is concurrently explained by variations in simulated parameters (across -here 1000- Monte-Carlo simulations). Note that perfect recovery would exhibit a diagonal structure, where variations in each estimated parameter is only due to variations in the corresponding simulated parameter. Strong non-diagonal elements in the matrix, either negative (blue) or positive (red), thus signal pairwise non-identifiability issues.

One can see that the neither the overcomplete approach nor the method of moments are capable of perfectly identifying DDM parameters. In brief, both methods exhibit strong non-identifiability issues. In particular: variations in $\hat{\sigma}$ are partially confused with variations in $x*$, and reciprocally. In addition, estimates of T_{ND} are also influenced by $x*$.

In fact, we expect non-identifiability issues of this sort, given the form of the DDM's self-consistency equation (Equation 7). In particular, this equation is invariant to a rescaling of all DDM parameters except T_{ND} (ie. predicted reaction times are left unchanged if all these parameters are rescaled by the same amount). Although this problematic invariance would disappear if, eg., the noise's standard deviation σ was fixed rather than fitted, other non-identifiability issues may still hamper DDM parameter recovery. To test this, we re-performed the above parameter recovery analysis, this time fixing the (simulated and estimated) noise's standard deviation σ to unity. We note that this arbitrary reduction of the parameter space was already suggested in seminal empirical applications of the DDM (Ratcliff, 1978). The ensuing results are summarized in Figure 8 below.

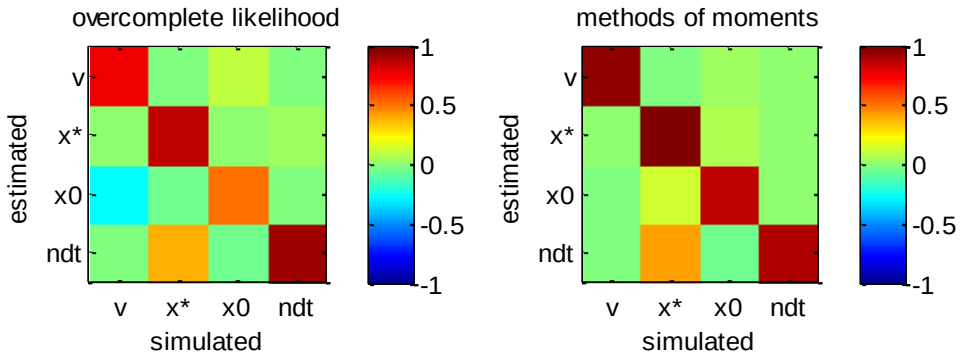

Figure 8: DDM parameter recovery analysis for the overcomplete likelihood method (left panel) and the method of moments (right panel). Here, the noise's standard deviation σ was fixed to unity. Same format as Fig 7.

One can see that some non-identifiability issues still remain, even when fixing the noise's standard deviation. In particular, both methods still partially confuse the non-decision time T_{ND} with the bound's height $x*$.

Now, accounting for predictable trial-by-trial variations in model parameters may, in principle, improve model identifiability. This is due to the fact that the net effect of each DDM parameter depends upon the setting of other parameters. Let us assume, for example, that the drift rate varies across trials according to some predictor variable (e.g., the relative evidence strength of alternative options in the context of decision making). The impact of other DDM parameters will not be the same, depending on whether the drift rate is high or low. In turn, there are fewer settings of these parameters that can predict trial-by-trial variations in RT data from variations in drift rate. To test this, we re-performed the recovery analysis, this time setting the drift rate according to a varying predictor variable, which is supposed to be known. The ensuing recovery results are given in Figure 9 below.

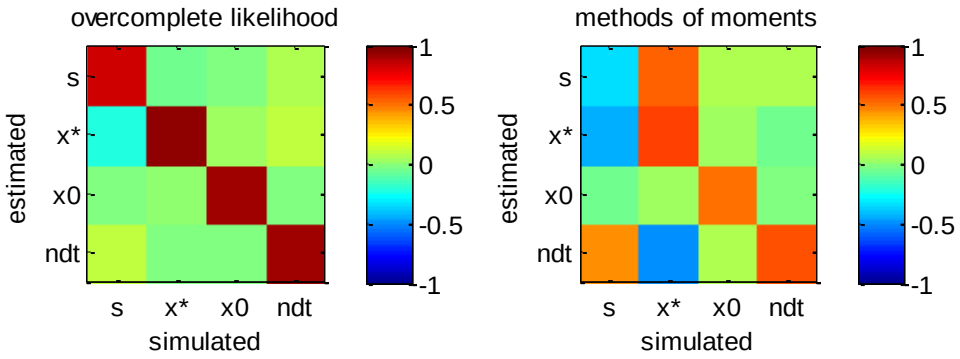

Figure 9: DDM parameter recovery analysis for the overcomplete likelihood method (left panel) and the method of moments (right panel). Here, the drift rate was set to a trial-dependant predictor variable. Same format as Fig 7.

One can see that the parameter recovery of the method of moments has worsened, when compared to the situation above, where σ was set to unity. The only parameter that is now accurately recovered is the initial bias x_0. These parameter recovery issues arise simply because the methods of moments cannot efficiently exploit trial-by-trial variations in drift rates (cf. section 5 above). The situation is quite different for the overcomplete approach. In brief, DDM parameters are now well identifiable (cf. strong diagonal structure of the recovery matrix on Figure 9). In particular, parameter recovery under the overcomplete approach is better than when simply fixing the noise standard deviation to unity (cf. Figure 8). This is a simple example of the gain in statistical efficiency that result from exploiting known trial-by-trial variations in DDM model parameters.

We note that fixing an additional DDM parameter, e.g. the standard deviation σ, would probably improve model identifiability further (for both approaches). This is typically done in most applications of the DDM . But the risk of drawing erroneous conclusions, e.g. blindly interpreting differences due to σ in terms of differences in other DDM parameters, should invite modelers to be cautious with this kind of strategy.

6. Application to a value-based decision making experiment

To demonstrate the above overcomplete likelihood approach, we apply it to data acquired in the context of a value-based decision making experiment (Lopez-Persem et al., 2016). This experiment was designed to understand how option values are compared when making a choice. In particular, it tested whether agents may have prior preferences that create default policies and shape the neural comparison process.

Prior to the choice session, participants (n = 24) rated the likeability of 432 items belonging to three different domains (food, music, magazines). Each domain included four categories of 36 items. At that time, participants were unaware of these categories. During the choice session, subjects performed series of choices between two items, knowing that one choice in each domain would be randomly selected at the end of the experiment and that they would stay in the lab for another 15 min to enjoy their reward (listening to the selected music, eating the selected food and reading the selected magazine). Trials were blocked in a series of nine choices between items belonging to the same two categories within a same domain. The two categories were announced at the beginning of the block, such that subjects could form a prior preference (although they were not explicitly asked to do so). We quantified this prior preference as the difference between mean likeability ratings (across all items within each of the two categories). Figure 10 below summarizes the main effects of a bias toward the default option (i.e. the option belonging to the favored category) in both choice and response time, above and beyond the effect of individual item values.

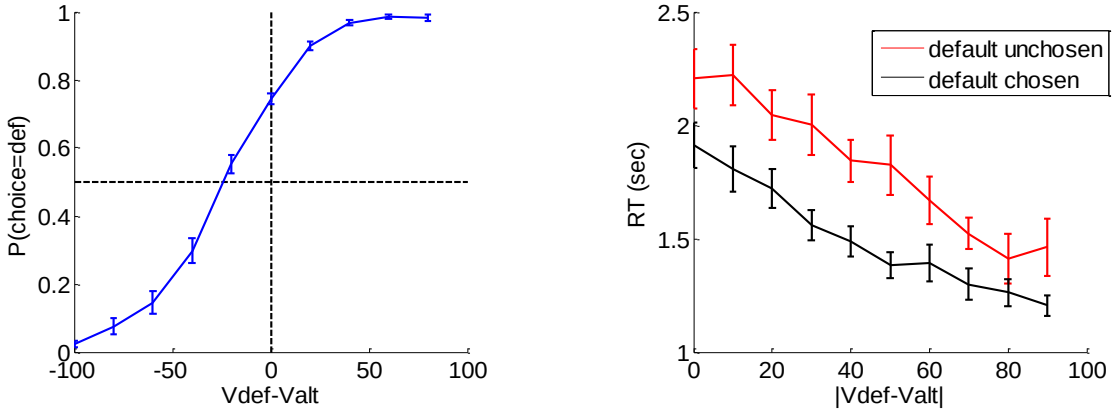

Figure 10: Evidence for choice and RT biases in the default/alternative frame. Left: Probability of choosing the default option (y-axis) is plotted as a function of decision value Vdef-Valt (x-axis), divided into 11 bins. Values correspond to likeability ratings given by the subject prior to choice session. Choice bias was defined as the difference between chance level (50%) and the observed probability of choosing the default option for a null decision value (when Vdef=Valt). Right: Response time RT (y-axis) is plotted as a function of the absolute decision value |Vdef-Valt| (x-axis) divided into 10 bins, separately for trials in which the default option was chosen (black) and unchosen (red). RT bias was defined as the difference between the intercepts (when Vdef=Valt) observed for the two types of choice.

A simple random effect analysis based upon logistic regression shows that the probability of choosing the default option significantly increases with decision value, i.e. the difference Vdef-Valt between the default and alternative option values ($t=8.4$, $dof=23$, $p<10^{-4}$). In addition, choice bias is significant ($t=8.7$, $dof=23$, $p<10^{-4}$). Similarly, RT significantly decreases with absolute decision value |Vdef-Valt| ($t=8.7$, $dof=23$, $p<10^{-4}$), and RT bias is significant ($t=7.4$, $dof=23$, $p<10^{-4}$).

To interpret these results, we fitted the DDM using the above overcomplete approach, when encoding the choice either (i) in terms of default versus alternative option (i.e. as is implicit on Figure 10) or (ii) in terms of right option versus left option. In what follows, we refer to the former choice frame as the "default/alternative" frame, and to the latter as the "native" frame. In both cases, the drift rate of each choice trial was set to the corresponding decision value (either Vdef- Valt or Vright-Vleft). It turns out that within-subject estimates of σ, x^* and T_{ND} do not depend upon the choice frame. More precisely, the cross-subjects correlation of these

estimates between the two choice frames is significant in all three cases (σ : r=0.76, p<10^{-4}; x^*: r=0.82, p<10^{-4}; T_{ND}: r=0.94, p<10^{-4}). This implies that inter-individual differences in σ, x^* and T_{ND} can be robustly identified, irrespective of the choice frame. However, the between-frame correlation is not significant for the initial bias x_0 (r=0.29, p=0.17). In addition, the initial bias is significant at the group level for the default/alternative frame (F=45.2, dof=[1,23], p<10^{-4}) but not for the native frame (F=2.36, dof=[1,23], p=0.14). In brief, the two choice frames only differ in terms of the underlying initial bias, which is only revealed in the default/alternative frame.

Now, we expect, from model simulations, that the presence of an initial bias induces both a choice bias, and a reduction of response times for default choices when compared to alternative choices (cf. upper-left and lower-right panels in Figure 1). The fact that $\hat{x}_0$ is significant in the default/alternative frame thus explains the observed choice and RT biases shown on Figure 10. But do inter-individual differences in $\hat{x}_0$ predict inter-individual differences in observed choice and RT biases? The corresponding statistical relationships are summarized on Figure 11 below.

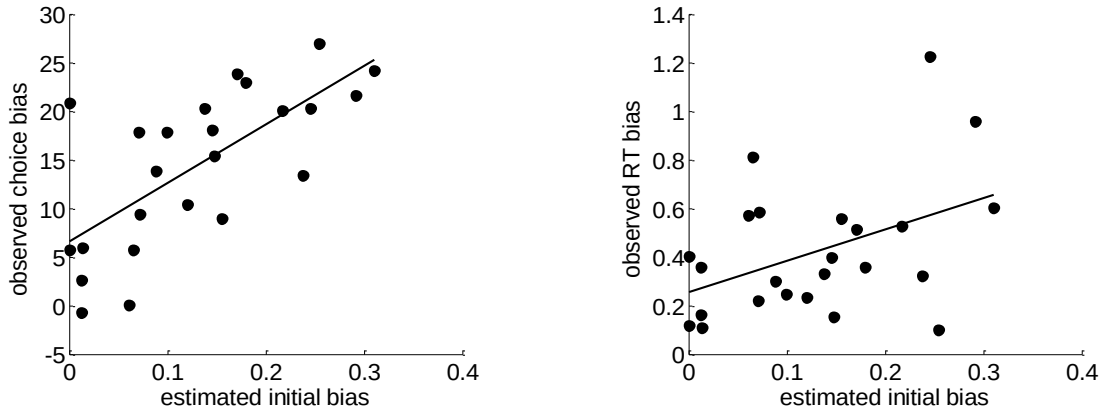

Figure 11: Model-based analyses of choice and RT data. Left: For each participant, the observed choice bias (y-axis) is plotted as a function of the initial bias estimate $\hat{x}_0$ in the default/alternative frame (x-axis). Right: Same for the observed RT bias.

One can see that both pairs of variables are statistically related (choice bias: r=0.70, $p<10^{-4}$; RT bias: r=0.44, p=0.03). This is important, because this provides further evidence in favor of the hypothesis that people's covert decision frame facilitates the default option. Note that this could not be shown using the method of moments, which was not able to capture these inter-individual differences.

Finally, can we exploit model fits to provide a normative argument for why the brain favors a biased choice frame? Recall that the DDM has been shown to be the optimal solution to the problem of optimizing the speed-accuracy tradeoff inherent in making online value-based decisions (Tajima et al., 2016). It may seem that the presence of an initial bias would induce a gain in decision speed that would be overcompensated by the ensuing loss of accuracy. But in fact, the net tradeoff between decision speed and accuracy depends upon how the system sets the bound's height $x*$. This is because $x*$ determines the demand for evidence before the system commits to a decision. More precisely, the system can favor decision accuracy by increasing $x*$, or improve decision speed by decreasing $x*$. We thus defined a measure $\widehat{e}$ of the optimality of each participant's decisions, by comparing the speed-accuracy efficiency of her estimated DDM and the maximum speed-accuracy efficiency that can be achieved over alternative bound heights $x*$ (see Appendix A1 below). This measure of optimality can be obtained either under the default-alternative frame or under the native frame. It turns out that the measured optimality of participants' decisions is significantly higher under the default/alternative frame than under the native frame ($\Delta\widehat{e}$ =0.007 ±0.003, t=2.2, dof=23, p=0.02). In other words, participants' decisions appear more optimal under the default/alternative frame than under the native frame. In this context, the benefit of the initial default bias may simply be to speed up decisions when evidence is congruent with initial preferences, at the expense of slowing down incongruent decisions, without compromising decision accuracy. This is adaptive, if the incongruent decisions are rarer than the congruent ones (as is the case in the experiment).

7. Discussion

In this note, we have described an overcomplete approach to fitting the DDM to trial-by-trial RT data. This approach is based upon a self-consistency equation of DDM models. It can exploit predictable variations of DDM parameters across trials to explain trial-by-trial variations of RT data above and beyond the impact of stochastic perturbations of the underlying decision variable. This eventually improves parameter recovery for the remaining unknown model parameters.

Strictly speaking, the DDM predicts the RT distributions conditional on the two alternative response outcomes. This is why DDM-based data analyses are not optimal when empirical design parameters are varied on a trial-by-trial basis (e.g., evidence strength). More precisely, one would need a few trial repetitions of empirical conditions (e.g., at least a few tens of trials per evidence strength) to estimate the underlying DDM parameters from the observed moments of associated RT distributions (Ratcliff, 2008; Wagenmakers et al., 2007, 2008). In principle, one could also derive the expected response time for each trial, and then find the DDM parameters that best match the expected and observed RTs (see, e.g., Moens and Zenon, 2017). But this becomes computationally cumbersome when the number of trials is high and prevents applications to most decision experiments. Critically however, situations where design parameters are varied on a trial-by-trial basis maximize the statistical power of the overcomplete approach. In other words, the statistical added-value of the overcomplete approach is highest when the experimental design induces predictable trial-by-trial variations in at least one of the DDM parameters (e.g., drift rates, as is exemplified here). This enables one to reliably estimate the other model parameters, which can then be used for between-condition and/or between-subject comparisons. This is because evaluating the underlying self-consistency equation (Equation 7) is much simpler than deriving moments of the conditional RT distributions (Broderick et al., 2009; Navarro and Fuss, 2009). More precisely,

its computational cost scales linearly with the number of trials because no numerical integration is involved.

We note that this feature of the overcomplete approach makes it particularly suited for learning experiments, where sequential decisions are based upon beliefs that are updated on a trial-by-trial basis from systematically varying pieces of evidence. In such contexts, existing modelling studies restricted the number of DDM parameters to deal with parameter recovery issues (Frank et al., 2015; Pedersen et al., 2017). This is problematic, since reducing the set of free DDM parameters can lead to systematic interpretation errors (cf. fitting data under the constraint that $\sigma \triangleq 1$). In contrast, it would be trivial to extend the overcomplete approach to learning experiments without having to simplify the parameter space. We will pursue this in forthcoming publications.

Now what are the limitations of the overcomplete approach?

First and foremost, the self-consistency equation cannot be used to simulate data. This restricts the utility of the approach to data analysis. Note however, that data simulations can still be performed using Equation 2, once the model parameters have been identified. Second, the accuracy of the method depends upon the reliability of reaction time data. In particular, the recovery of the noise's standard deviation depends upon the accuracy of the empirical proxy for decision times (cf. second term in Equation 7). Third, the computational cost of model inversion will scale with the number of trials. This is because each trial has its own nuisance perturbation parameter. Fourth, proper bayesian model comparison may be more difficult. In particular, simulations show that a chance model always has a higher model evidence than the overcomplete model. This is another consequence of the overcompleteness of the likelihood function, which eventually pays a high complexity penalty cost in the context of Bayesian model comparison. Whether different model variants of the self-consistency equation can be discriminated is beyond the scope of the current work.

Is the overcomplete approach limited to the basic DDM model variant?

Recall that recent extensions of vanilla DDMs include e.g., collapsing bounds (Hawkins et al., 2015; Voskuilen et al., 2016) or nonlinear transformations of the state-space (Tajima et al., 2016). Although we did not consider these variants of the DDM model, we would argue that most of these effectively reduce to a simple modification of the self-consistency equation. Let us assume that Equations 2-3 still hold, i.e. the decision process is still somehow based upon a gaussian random walk. However, we now assume that the decision is triggered when an arbitrary transformation $z: x \to z(x)$ of the base random walk x_t has reached a predefined threshold $z^* \triangleq z^*(t)$ that can vary over time (e.g., a collapsing bound). Equation 5 now becomes:

$$z^*(\tau_i) = o_i \; z\left(x_0 + \tau_i v + \sigma\sqrt{\tau_i}\tilde{\eta}_i\right) \tag{14}$$

If the transformation $z: x \to z(x)$ is invertible (i.e. if z^{-1} exists and is unique), then the observation equation for reaction times y_i reduces to the following variant of Equation 7:

$$y_i \approx \frac{z^{-1}\left(o_i \; z^*(y_i - T_{ND})\right) - x_0}{v} - \frac{\sigma\sqrt{y_i - T_{ND}}}{v}\tilde{\eta}_i + T_{ND} + \varepsilon_i \tag{15}$$

where $z^{-1}\left(o_i \; z^*(\bullet)\right)$ is effectively operating a simple change of variable.

If, however, the transformation $z: x \to z(x)$ is not invertible, then it is not possible, in all generality, to derive an observation equation of reaction times equivalent to Equation 7.

Finally, we note that the code that is required to perform a DDM-based data analysis under the overcomplete approach will be made available soon from the VBA academic freeware https://mbb-team.github.io/VBA-toolbox/ (Daunizeau et al., 2014).

References

Beal, M.J. (2003). Variational algorithms for approximate Bayesian inference /. PhD Thesis.

Bitzer, S., Park, H., Blankenburg, F., and Kiebel, S.J. (2014). Perceptual decision making: drift-diffusion model is equivalent to a Bayesian model. Front. Hum. Neurosci. *8*.

Broderick, T., Wong-Lin, K.F., and Holmes, P. (2009). Closed-Form Approximations of First-Passage Distributions for a Stochastic Decision-Making Model. Appl. Math. Res. EXpress *2009*, 123–141.

Daunizeau, J. (2017). The variational Laplace approach to approximate Bayesian inference. ArXiv170302089 Q-Bio Stat.

Daunizeau, J., Adam, V., and Rigoux, L. (2014). VBA: A Probabilistic Treatment of Nonlinear Models for Neurobiological and Behavioural Data. PLoS Comput Biol *10*, e1003441.

De Martino, B., Fleming, S.M., Garrett, N., and Dolan, R.J. (2012). Confidence in value-based choice. Nat. Neurosci. *16*, 105–110.

Drugowitsch, J., Moreno-Bote, R., Churchland, A.K., Shadlen, M.N., and Pouget, A. (2012). The Cost of Accumulating Evidence in Perceptual Decision Making. J. Neurosci. *32*, 3612–3628.

Frank, M.J., Gagne, C., Nyhus, E., Masters, S., Wiecki, T.V., Cavanagh, J.F., and Badre, D. (2015). fMRI and EEG Predictors of Dynamic Decision Parameters during Human Reinforcement Learning. J. Neurosci. *35*, 485–494.

Friston, K., Mattout, J., Trujillo-Barreto, N., Ashburner, J., and Penny, W. (2007). Variational free energy and the Laplace approximation. NeuroImage *34*, 220–234.

Gold, J.I., and Shadlen, M.N. (2007). The neural basis of decision making. Annu. Rev. Neurosci. *30*, 535–574.

Hanks, T., Kiani, R., and Shadlen, M.N. (2014). A neural mechanism of speed-accuracy tradeoff in macaque area LIP. ELife *3*, e02260.

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. J. Neurosci. *35*, 2476–2484.

Huk, A.C., and Shadlen, M.N. (2005). Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. J. Neurosci. Off. J. Soc. Neurosci. *25*, 10420–10436.

Krajbich, I., Armel, C., and Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. Nat. Neurosci. *13*, 1292–1298.

Lopez-Persem, A., Domenech, P., and Pessiglione, M. (2016). How prior preferences determine decision-making frames and biases in the human brain. ELife *5*, e20317.

Milosavljevic, M., Malmaud, J., Huth, A., Koch, C., and Rangel, A. (2010). The drift diffusion model can account for the accuracy and reaction time of value-based choices under high and low time pressure. Judgm. Decis. Mak. *5*, 437–449.

Moens, V., and Zenon, A. (2017). Variational Treatment of Trial-by-Trial Drift-Diffusion Models of Behaviour. BioRxiv 220517.

Navarro, D.J., and Fuss, I.G. (2009). Fast and accurate calculations for first-passage times in Wiener diffusion models. J. Math. Psychol. *53*, 222–230.

Pedersen, M.L., Frank, M.J., and Biele, G. (2017). The drift diffusion model as the choice rule in reinforcement learning. Psychon. Bull. Rev. *24*, 1234–1251.

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

Ratcliff, R. (2008). The EZ diffusion method: too EZ? Psychon. Bull. Rev. *15*, 1218–1228.

Ratcliff, R., and McKoon, G. (2008). The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. Neural Comput. *20*, 873–922.

Ratcliff, R., and Tuerlinckx, F. (2002). Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. Psychon. Bull. Rev. *9*, 438–481.

Ratcliff, R., Smith, P.L., Brown, S.D., and McKoon, G. (2016). Diffusion Decision Model: Current Issues and History. Trends Cogn. Sci. *20*, 260–281.

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

Tajima, S., Drugowitsch, J., and Pouget, A. (2016). Optimal policy for value-based decision-making. Nat. Commun. *7*, 12400.

Voskuilen, C., Ratcliff, R., and Smith, P.L. (2016). Comparing fixed and collapsing boundary versions of the diffusion model. J. Math. Psychol. *73*, 59–79.

Wagenmakers, E.-J., van der Maas, H.L.J., and Grasman, R.P.P.P. (2007). An EZ-diffusion model for response time and accuracy. Psychon. Bull. Rev. *14*, 3–22.

Wagenmakers, E.-J., van der Maas, H.L.J., Dolan, C.V., and Grasman, R.P.P.P. (2008). EZ does it! Extensions of the EZ-diffusion model. Psychon. Bull. Rev. *15*, 1229–1235.

Acknowledgements

We would like to thank Dr. Alizée Lopez-Persem for providing us with the empirical data that serves to demonstrate our approach.

Appendix: evaluating the computational optimality of a DDM

Recall that the DDM has been shown to be the optimal solution to the computational problem of trading speed with accuracy when making online value-based decisions (Tajima et al., 2016). This speed-accuracy tradeoff arises because the decision system needs to accumulate noisy value signals before making a reliable decision. Arguably, the decision system cannot control the upstream decision variables, such as the evidence strength (i.e. the drift rate v), or the reliability of value signals (i.e. the noise's standard deviation σ). However, it can set the bound's height x^* so as to optimize the ensuing speed-accuracy tradeoff. More precisely, the system can favor decision accuracy by increasing x^*, or improve decision speed by decreasing x^*. Now one can score the efficiency $e(x^*)$ of any x^* setting, in terms of the ensuing mean rate of accurate responses:

$$e(x^*) = \frac{P_c(x^*)}{E\left[HT \middle| x^*\right]} \tag{16}$$

where $E\left[HT \middle| x^*\right]$ is the expected hitting time, and $P_c(x^*) \triangleq P\left(o = correct \middle| x^*\right)$ is the ensuing probability of making a correct decision. The optimal bound's height is such that it maximizes the efficiency. Now the system may not have set the bound's height to its optimal value. The optimality $\hat{e}$ of the system can thus be measured in terms of the ratio between the actual efficiency of the system and its optimal efficiency:

$$\hat{e} = \frac{e(\hat{x}^*)}{\max_{x^*} e(x^*)} \tag{17}$$

where $\hat{x}*$ is the estimated bound's height of a given participant. Equation 17 can be extended to situations where drift rate changes in a trial-by-trial fashion, by averaging the optimality score $\hat{e}$ over trials. This makes the average optimality score $\hat{e}$ a measure of how well the system adapts to the global statistics of decisions it has to make.

Figure A1 below shows a typical example of the derivation of the optimality $\hat{e}$ score, based upon the estimated DDM parameters of a study participant (and the corresponding sequence of drift rates over trials) under the default/alternative frame of reference.

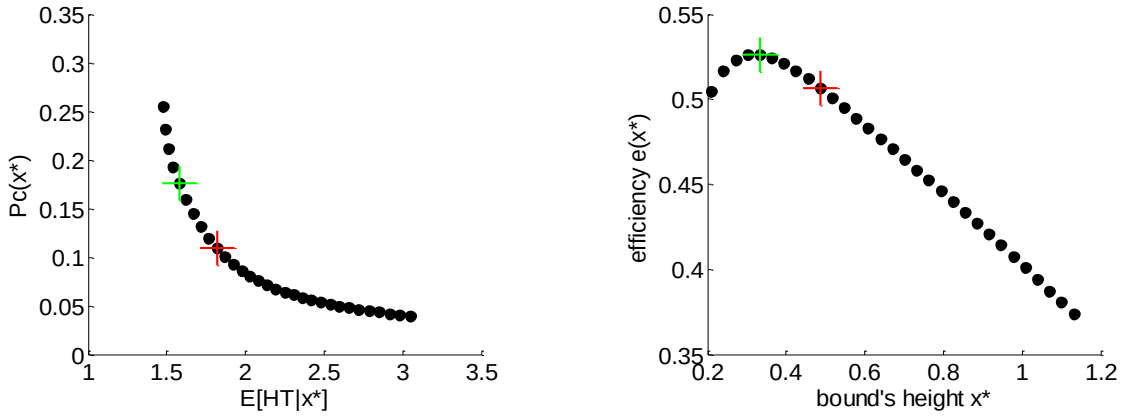

Figure A1: Model-based evaluation of DDM optimality. Left: The probability $P_c(x*)$ of making a correct decision (y-axis) is plotted against the expected decision time $E[HT|x*]$ (y-axis) for each possible bound height. The green and red crosses depict the optimal and actual speed-accuracy tradeoffs, respectively. Right: The efficiency $e(\hat{x}*)$ (y-axis) is plotted for each possible bound height $x*$. The green cross depicts the optimal efficiency and the red cross depicts the actual efficiency achieved for this participant.

The optimality score of this participant is estimated as $\hat{e} = 0.5066/0.5262 \approx 0.96$. In other terms, this participant's decisions seem to arise from a DDM decision system that is close to its optimal setting, in terms of the underlying speed-accuracy tradeoff.