An Overcomplete Approach to Fitting Drift-Diffusion Decision Models to Trial-By-Trial Data

Q. Feltgen, J. Daunizeau

► To cite this version:

Q. Feltgen, J. Daunizeau. An Overcomplete Approach to Fitting Drift-Diffusion Decision Models to Trial-By-Trial Data. *Frontiers in Artificial Intelligence and Applications*, 2021, 4, pp.531316. 10.3389/frai.2021.531316 . hal-03212563

HAL Id: hal-03212563

<https://hal.sorbonne-universite.fr/hal-03212563>

Submitted on 29 Apr 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

An Overcomplete Approach to Fitting Drift-Diffusion Decision Models to Trial-By-Trial Data

Q. Feltgen¹ and J. Daunizeau^{1,2*}

¹Paris Brain Institute (ICM), Sorbonne Université, Inserm, CNRS, Hôpital Pitié-Salpêtrière, Paris, France, ²ETH, Zurich, Switzerland

OPEN ACCESS

Edited by:

Thomas Parr,
University College London,
United Kingdom

Reviewed by:

Sebastian Gluth,
University of Hamburg, Germany
Vincent Moens,
Catholic University of Louvain,
Belgium

*Correspondence:

J. Daunizeau
jean.daunizeau@gmail.com

Specialty section:

This article was submitted to
Machine Learning and Artificial
Intelligence,
a section of the journal
Frontiers in Artificial Intelligence

Received: 31 January 2020

Accepted: 17 February 2021

Published: 09 April 2021

Citation:

Feltgen Q and Daunizeau J (2021) An
Overcomplete Approach to Fitting
Drift-Diffusion Decision Models to Trial-
By-Trial Data.
Front. Artif. Intell. 4:531316.
doi: 10.3389/frai.2021.531316

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 only yield reliable parameter estimates in specific situations (c.f. fixed drift rates vs drift rates varying over trials). In addition, they become computationally unfeasible when more general DDM variants are considered (e.g., with collapsing bounds). In this note, we propose a fast and efficient approach to parameter estimation that relies on fitting a “self-consistency” equation that RT fulfill under the DDM. This effectively bypasses the computational bottleneck of standard DDM parameter estimation approaches, at the cost of estimating the trial-specific neural noise variables that perturb the underlying evidence accumulation process. For the purpose of behavioral data analysis, these act as nuisance variables and render the model “overcomplete,” which is finessed using a variational Bayesian system identification scheme. However, for the purpose of neural data analysis, estimates of neural noise perturbation terms are a desirable (and unique) feature of the approach. Using numerical simulations, we show that this “overcomplete” approach matches the performance of current parameter estimation approaches for simple DDM variants, and outperforms them for more complex DDM variants. Finally, we demonstrate the added-value of the approach, when applied to a recent value-based decision making experiment.

Keywords: DDM, decision making, computational modeling, variational bayes, neural noise

INTRODUCTION

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 response 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 (Gold and Shadlen, 2007; Resulaj et al., 2009; Milosavljevic et al., 2010; De Martino et al., 2012; Hanks et al., 2014; Pedersen

et al., 2017), and to theoretical work showing that DDMs can be thought of as optimal decision problem solvers (Bogacz et al., 2006; Balci et al., 2011; Drugowitsch et al., 2012; Zhang, 2012; Tajima et al., 2016). Critically, mathematical analyses of the DDM soon demonstrated that it suffers from inherent non-identifiability issues, e.g., predicted choices and RTs are invariant under any arbitrary rescaling of DDM parameters (Ratcliff and Tuerlinckx, 2002; Ratcliff et al., 2016). This is important because, in principle, this precludes proper, quantitative, DDM-based data analysis. Nevertheless, over the past decade, many statistical approaches to DDM parameter estimation have been proposed, which yield efficient parameter estimation under simplifying assumptions (Voss and Voss, 2007; Wagenmakers et al., 2007, 2008; Vandekerckhove and Tuerlinckx, 2008; Grasman et al., 2009; Zhang, 2012; Wiecki et al., 2013; Zhang et al., 2014; Hawkins et al., 2015; Voskuilen et al., 2016; Pedersen and Frank, 2020). Typically, these techniques essentially fit the choice-conditional distribution of observed RT (or moments thereof), having arbitrarily fixed at least one of the DDM parameters. They are now established statistical tools for experimental designs where the observed RT variability is mostly induced by internal (e.g., neural) stochasticity in the decision process (Boehm et al., 2018).

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) may be systematically 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) effectively vary over trials. Classical DDM parameter estimation approaches do not optimally handle this kind of experimental designs, because these lack the trial repetitions that would be necessary to provide empirical estimates of RT moments in each condition. In turn, alternative statistical approaches to parameter estimation have been proposed, which can exploit predictable inter-trial variations of DDM variables to fit the model to RT data (Wabersich and Vandekerckhove, 2014; Moens and Zenon, 2017; Pedersen et al., 2017; Fontanesi et al., 2019a; Fontanesi et al., 2019b; Gluth and Meiran, 2019). In brief, they directly compare raw RT data with expected RTs, which vary over trials in response to known variations in internal variables. Although close to optimal from a statistical perspective, they suffer from a challenging computational bottleneck that lies in the trial-by-trial derivation of RT first-order moments. This is why they are typically constrained to simple DDM variants, for which analytical solutions exist (Navarro and Fuss, 2009; Srivastava et al., 2016; Fengler et al., 2020; Shinn et al., 2020).

This note is concerned with the issue of obtaining reliable DDM parameter estimates from concurrent trial-by-trial choice and response time data, for a broad class of DDM variants. We propose a fast and efficient approach that relies on fitting a *self-consistency* equation, which RTs necessarily fulfill under the DDM. This provides a simple and elegant way to bypassing the common computational bottleneck of existing approaches, at the cost of considering additional trial-specific nuisance model

variables. These are the cumulated “neural” noise that perturbs the evidence accumulation process at the corresponding trial. Including these variables in the model makes it “overcomplete,” the identification of which is finessed with a dedicated variational Bayesian scheme. In turn, the ensuing overcomplete approach generalizes to a large class of DDM model variants, without any additional computational and/or implementational burden.

In *Model Formulation and Impact of DDM Parameters* 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 *A Self-Consistency Equation for DDMs* and *An Overcomplete Likelihood Approach to DDM Inversion* sections, we derive the DDM’s self-consistency equation and describe the ensuing overcomplete approach to DDM-based data analysis. In *Parameter Recovery Analysis* section, we perform parameter recovery analyses for standard DDM fitting procedures and the overcomplete approach. In *Application to a Value-Based Decision Making Experiment* section, we demonstrate the added-value of the overcomplete approach, when applied to a value-based decision making experiment. Finally, in *Discussion* section, we discuss our results in the context of the existing literature. In particular, we comment on the potential utility of neural noise perturbation estimates for concurrent neuroimaging data analysis.

MODEL FORMULATION AND IMPACT OF DDM PARAMETERS

First, let us recall the simplest form of a drift-diffusion decision model or DDM (in what follows, we will refer to this variant as the “vanilla” DDM). Let $x(t)$ be a decision variable that captures the accumulated evidence (up to time t) in favor of a given option in a binary choice set. Under the vanilla DDM, a decision is triggered whenever $x(t)$ hits either of two bounds, which are positioned at $x = b$ and $x = -b$, respectively. When a bound hit occurs defines the decision time, and which bound is hit determines the (binary) decision outcome o . By assumption, the decision variable $x(t)$ is supposed to follow the following stochastic differential equation:

$$dx = \tilde{\nu} \times dt + \tilde{\sigma} \times d\eta \quad (1)$$

where ν is the drift rate, $d\eta \sim N(0, dt)$ is a standard Wiener process, and $\tilde{\sigma}$ is the standard deviation of the stochastic (diffusion) perturbation term.

Equation 1 can be discretized using an Euler-Maruyama scheme (Kloeden and Platen, 1992), yielding the following discrete form of the decision variable dynamics:

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

where t indexes time on a temporal grid with resolution Δt , $\nu = \tilde{\nu} \Delta t$ is the discrete-time drift rate, $\sigma = \tilde{\sigma} \sqrt{\Delta t}$ is the discrete-time standard deviation of the perturbation term 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”.

FIGURE 1 | Impact of initial bias x_0 . In all panels, the color 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).

The joint distribution of response times and decision outcomes depends upon the DDM parameters, which include: the drift rate ν , the bound’s height b , the noise’s standard deviation σ and the initial condition x_0 . DDMs also typically include a so-called “non-decision” time parameter T_{ND} , which captures systematic latencies between covert bound hit times and overt response times. Under such simple DDM variant, variability in response times and decision outcomes derive from stochastic terms η . These are typically thought of as neural noise that perturb the evidence accumulation process within the brain’s decision system (Gold and Shadlen, 2007; Turner et al., 2015; Guevara Erra et al., 2019).

Under such simple DDM variant, analytical expressions exist for the first two moments of RT distributions (Srivastava et al., 2016). Higher-order moments can also be derived from efficient semi-analytical solutions to the issue of deriving the joint choice/RT distribution (Navarro and Fuss, 2009). However, more complex variants of the DDM (including, e.g., collapsing bounds) are much more difficult to simulate, and require either sampling schemes or numerical solvers of the underlying Fokker-Planck equation (Fengler et al., 2020; Shinn et al., 2020).

Figures 1–4 below demonstrate the impact of model parameters on the decision outcome ratios $P(o|\nu, x_0, b, \sigma)$ and the first three moments of conditional hitting time (HT) distributions, namely: their mean $E[HT|o, \nu, x_0, b, \sigma]$, variance $V[HT|o, \nu, x_0, b, \sigma]$ and skewness $S_k[HT|o, \nu, x_0, b, \sigma]$. 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 let first us summarize the impact of DDM parameters. To do this, we first set model parameters to the following “default” values: $\nu = 1/2$, $x_0 = 1$, $b = 10$ and $\sigma = 4$. This parameter setting yields about 30% decision errors, which we take as a valid reference point for typical studies of decision making. In what follows, 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 .

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 ν is positive, and therefore favors

FIGURE 2 | Impact of drift rate v . Same format as **Figure 1**.

FIGURE 3 | Impact of the perturbation's standard deviation σ . Same format as **Figure 1** (but the x-axis is now in log-scale).

“up” decisions. Most importantly, the initial bias is the only DDM parameter that has opposite effects on mean HT for “up” and “down” decision outcomes.

Figure 2 below shows the impact of drift rate v .

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 $v \rightarrow 0$. Also,

the drift rate has little impact on the HT skewness. Note that, in contrast to the initial bias, the impact of the drift rate on mean 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 4 | Impact of the threshold's height b . Same format as **Figure 1**.

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 mean HT for “up” and “down” decisions. Finally, increasing the standard deviation eventually maximizes the entropy of the decision outcomes, i.e., $P(o) \rightarrow 1/2$ when $\sigma \rightarrow \infty$. 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 b .

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., $P(o) \rightarrow 0$ or 1 when $b \rightarrow \infty$. This directly derives from the fact that increasing b 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 favor 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 response 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 RT moments (e.g., up to third order) match model predictions. In what follows, we refer to this as the “method of moments” (see **Supplementary Appendix S2**). However, we will see below that the DDM is not perfectly identifiable. One can also see that changing any of these parameters from trial to trial will most likely induce non-trivial variations in RT data. Here, the method of moments may not be optimal, because predictable trial-by-trial variations in DDM parameters will be lumped together with stochastic perturbation-induced variations. One may thus rather attempt to match the trial-by-trial series of raw response times directly with their corresponding first-order moments. In what follows, we refer to this as the “method of trial means” (see **Supplementary Appendix S3**). Given the computational cost of deriving expected response times for each trial, this type of approach is typically restricted to the vanilla DDM, since there is no known analytical expression for response time moments under more complex DDM variants.

Below, we suggest a simple and efficient way of performing DDM parameter estimation, which applies to a broad class of DDM variants without significant additional computational burden. This follows from fitting a self-consistency equation that, under a broad class of DDM variants, response times have to obey.

A SELF-CONSISTENCY EQUATION FOR DDMs

First, note that **Eq. 2** can be rewritten as follows:

FIGURE 5 | Self-consistency equation. Monte-Carlo simulation of 200 trials of a DDM, with arbitrary parameters (in this example, the drift rate is positive). In all panels, the color code indicates the decision outcomes, which depends upon the sign of the drift rate (green: correct decisions, red: incorrect decisions). Upper-left panel: simulated trajectories of the decision variable (y -axis) as a function of time (x -axis). Upper-right panel: response times' distribution for both correct and incorrect choice outcomes over the 200 Monte-Carlo simulations. Lower-left panel: outcome ratios. Lower-right panel: the left-hand side of **Eq. 7** (y -axis) is plotted against the right-hand side of **Eq. 7** (x -axis), for each of the 200 trials.

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

where we coin $\tilde{\eta}_t \triangleq 1/\sqrt{t} \sum_{t'=0}^{t-1} \eta_{t'}$ the “normalized cumulative perturbation”. Now let τ_i be the decision time of the i th trial. Note that decision times are trivially related to cumulative perturbations because, by definition, $|x_{\tau_i}| = b$. This implies that:

$$b = |x_0 + \tau_i v + \sigma \sqrt{\tau_i} \tilde{\eta}_{\tau_i}| \quad (4)$$

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

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

$$b = \begin{cases} x_0 + \tau_i v + \sigma \sqrt{\tau_i} \tilde{\eta}_{\tau_i} & \text{if } o_i = 1 \text{ ('up' decision)} \\ -x_0 - \tau_i v - \sigma \sqrt{\tau_i} \tilde{\eta}_{\tau_i} & \text{if } o_i = -1 \text{ ('down' decision)} \end{cases} \quad (5)$$

$$= o_i (x_0 + \tau_i v + \sigma \sqrt{\tau_i} \tilde{\eta}_{\tau_i})$$

where o_i can only take two possible values (-1 or 1). **Eq. 5** can then be used to relate decision times directly to DDM model parameters (and cumulative perturbations):

$$\tau_i = \frac{o_i b - x_0}{v} - \frac{\sigma \sqrt{\tau_i}}{v} \tilde{\eta}_{\tau_i} \quad (6)$$

From **Eq. 6**, one can express observed trial-by-trial empirical response times y_i as follows:

$$y_i \approx \frac{o_i b - x_0}{v} - \frac{\sigma \sqrt{y_i - T_{ND}}}{v} \tilde{\eta}_{\tau_i} + T_{ND} + \varepsilon_i \quad (7)$$

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

Note that decision times effectively appear on both the left-hand and the right-hand side of **Eqs 6, 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 **Eq. 7** as a “self-consistency” constraint that response times have to fulfill under the DDM. This is why we refer to **Eq. 7** as the *self-consistency equation* of DDMs. This, however, prevents us from using **Eq. 7** to generate data under the DDM. In other terms, **Eq. 7** is only useful when analyzing empirical RT data.

Figure 5 below exemplifies the accuracy of DDM's self-consistency equation, using a Monte-Carlo simulation of 200 trials under the vanilla DDM.

One can see that the DDM's self-consistency equation is valid, i.e., simulated response times almost always equate their

theoretical prediction. The few (small) deviations that can be eyeballed on the lower-right panel of **Figure 5** actually correspond to simulation artifacts where the decision variable exceeds the bound by some relatively small amount. This happens when the discretization step Δt (cf. **Eq. 2**) is too large when compared to the relative magnitude of the stochastic component of the system's dynamics. In effect, these artifactual errors grow when σ/ν increases. Nevertheless, in principle, these and other errors would be absorbed in the model residuals ε_i of **Eq. 7**.

Now recall that recent extensions of vanilla DDMs include e.g., collapsing bounds (Hawkins et al., 2015; Voskuilen et al., 2016) and/or nonlinear transformations of the state-space (Tajima et al., 2016). As the astute reader may have already guessed, the self-consistency equation can be generalized to such DDM variants. Let us assume that **Eqs 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 $z : x \rightarrow z(x)$ of the base random walk x_t has reached a predefined threshold $\hat{b}(t)$ that can vary over time (e.g., a collapsing bound). **Eq. 5** now becomes:

$$\hat{b}(t_i) = o_i z(x_0 + \tau_i \nu + \sigma \sqrt{\tau_i} \tilde{\eta}_i) \quad (8)$$

If the transformation $z : x \rightarrow z(x)$ is invertible (i.e., if z^{-1} exists and is unique), then the self-consistency equation for reaction times y_i now generalizes as follows:

$$y_i \approx \underbrace{\frac{z^{-1} \left[o_i \hat{b}(y_i - T_{ND}) \right] - x_0}{\nu}}_{g(\nu, x_0, \sigma, T_{ND}, \tilde{\eta}_i)} - \frac{\sigma \sqrt{y_i - T_{ND}}}{\nu} \tilde{\eta}_i + T_{ND} + \varepsilon_i \quad (9)$$

where $g(\nu, x_0, \sigma, T_{ND}, \tilde{\eta}_i)$ is the “expected” (or rather, “self-consistent”) response time at trial i , which depends nonlinearly on DDM parameters (and on response times). Note that one recovers the self-consistency equation of “vanilla” DDM (**Eq. 7**) when setting $z(x) = z^{-1}(x) = x$ and $\hat{b}(t) = b \forall t$.

Importantly, inverting **Eq. 9** can be used to estimate parameters γ and ω that control the transformation $z_\gamma : x \rightarrow z_\gamma(x)$ or the collapsing bounds $\hat{b}_\omega : t \rightarrow \hat{b}_\omega(t)$, respectively. We will see examples of this in the Results section below. This implies that the self-consistency equation can be used, in conjunction with adequate statistical parameter estimation approaches (see below), for estimating DDM parameters under many different variants of DDM, including those for which no analytical result exists for the response time distribution.

AN OVERCOMPLETE LIKELIHOOD APPROACH TO DDM INVERSION

Fitting **Eq. 9** to response time data reduces to finding the set of parameters that renders the DDM self-consistent. In doing so, normalized cumulative perturbation terms $\tilde{\eta}$ are treated as nuisance model parameters, but model parameters nonetheless. This means that there are more model parameters

than there are data points. In other words, **Eq. 9** induces an “overcomplete” likelihood function $p(y|\nu, x_0, \sigma, \omega, \gamma, T_{ND}, \tilde{\eta}, \lambda)$:

$$\begin{aligned} p(y|\nu, x_0, \sigma, \omega, \gamma, T_{ND}, \tilde{\eta}, \lambda) &= \prod_{i=1}^n p(y_i|\nu, x_0, \sigma, \omega, \gamma, T_{ND}, \tilde{\eta}_i, \lambda) \\ &= \prod_{i=1}^n N(g(\nu, x_0, \sigma, \omega, \gamma, T_{ND}, \tilde{\eta}_i), \lambda) \end{aligned} \quad (10)$$

where λ is the variance of the model residuals ε_i of **Eq. 9**, $g(\cdot)$ is the “self-consistent” response time given in **Eq. 9**, and we have used the (convenient but slightly abusive) notation $\tilde{\eta}_i$ to reference cumulative perturbations w.r.t. to their corresponding trial index.

Dealing with such overcomplete likelihood function requires additional constraints on model parameters: this is easily done within a Bayesian framework. Therefore, we rely on the variational Laplace approach (Friston et al., 2007; Daunizeau, 2017), which was developed to perform approximate bayesian inference on nonlinear generative models (see **Supplementary Appendix S1** for mathematical details). In what follows, we propose a simple set of prior constraints that help regularizing the inference.

- a. Prior moments of the cumulative perturbations: the “no barrier” approximation

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 $\tilde{\eta}$ such that $P(\tilde{\eta} \geq \tilde{\eta}_+) = Q_+$ (see **Figure 6** below). Then, we know that errors correspond to those perturbations $\tilde{\eta}_i$ that are smaller than $\tilde{\eta}_+$. But what do we know about the distribution of perturbations? Importantly, if the DDM's stochastic evidence accumulation process had no decision bound, then the distribution of normalized cumulative perturbations would be invariant over time and such that $\tilde{\eta}_t \sim N(0, 1) \forall t$. This, in fact, is the very reason why we introduced normalized cumulative perturbations in **Eq. 3**. Under this “no barrier” approximation, one can now derive the conditional expectations $\tilde{\mu}_+$ and $\tilde{\mu}_-$ of the perturbation $\tilde{\eta}_i$, given that the decision outcome o_i is correct or erroneous, respectively:

$$\begin{cases} \tilde{\mu}_+ \triangleq E[\tilde{\eta}_i | o_i = 1] = E[\tilde{\eta}_i | \tilde{\eta}_i > \tilde{\eta}_+] = \frac{1}{(1 - Q_+) \sqrt{2\pi}} \exp\left(-\frac{1}{2}\tilde{\eta}_+^2\right) \\ \tilde{\mu}_- \triangleq E[\tilde{\eta}_i | o_i = -1] = E[\tilde{\eta}_i | \tilde{\eta}_i < \tilde{\eta}_+] = -\frac{1}{Q_+ \sqrt{2\pi}} \exp\left(-\frac{1}{2}\tilde{\eta}_+^2\right) \end{cases} \quad (11)$$

Equation 11 is obtained from the known expression of first-order moments of a truncated normal density $N(0, 1)$. Critically, **Eq. 11** does not depend upon DDM parameters. Of course, the

FIGURE 6 | Approximate conditional distributions of the normalized cumulative perturbations. Upper-left panel: The black line shows the “no barrier” standard normal distribution of normalized cumulative perturbations. The shaded gray area has size Q_+ , and its left bound (dashed black line) is the critical value $\tilde{\eta}_+$ below which cumulative perturbations eventually induce errors. The green and red lines depict the ensuing approximate conditional distributions given in **Eq. 13**. Upper-right panel: a Representative monte-carlo simulation. The green and red bars show the sample histogram of normalized cumulative perturbations for correct and erroneous decisions, respectively (over 200 trials, same simulation as in **Figure 5**). The green and red lines depict the corresponding approximate conditional normal distributions (cf. **Eq. 13**). Lower-left panel: The sample mean estimates of conditional perturbations (y-axis) are plotted against their “no barrier” approximation (x-axis, **Eq. 11**). Monte-carlo simulations are split according to the sign of the drift rate, and then binned according to deciles of approximate conditional means of normalized cumulative perturbations (green: Correct, red: error, errorbars: Within-decile means $\pm$ standard deviations). The black dotted line shows the identity mapping (perfect approximation). Lower-right panel: The sample variance estimates of normalized cumulative perturbations (y-axis) are plotted against their “no barrier” approximation (x-axis, **Eq. 12**). Same format as lower-left panel.

same logic extends to conditional variances $\tilde{\Sigma}_+$ and $\tilde{\Sigma}_-$, whose analytical expressions are given by:

$$\begin{cases} \tilde{\Sigma}_+ \triangleq V[\tilde{\eta}_i | o_i = 1] = V[\tilde{\eta}_i | \tilde{\eta}_i > \tilde{\eta}_+] = 1 + \tilde{\eta}_+ \tilde{\mu}_+ - \tilde{\mu}_+^2 \\ \tilde{\Sigma}_- \triangleq V[\tilde{\eta}_i | o_i = -1] = V[\tilde{\eta}_i | \tilde{\eta}_i < \tilde{\eta}_+] = 1 + \tilde{\eta}_+ \tilde{\mu}_- - \tilde{\mu}_-^2 \end{cases} \quad (12)$$

A simple moment-matching approach thus suggests to approximate the conditional distribution $p(\tilde{\eta}_i | o_i)$ of normalized cumulative perturbations as follows:

$$p(\tilde{\eta}_i | o_i) = \begin{cases} N(\tilde{\mu}_+, \tilde{\Sigma}_+) & \text{if } o_i = \text{correct} \\ N(\tilde{\mu}_-, \tilde{\Sigma}_-) & \text{if } o_i = \text{error} \end{cases} \quad (13)$$

where the correct/error label depends on the sign of the drift rate. This concludes the derivation of our simple “no barrier” approximation to the conditional moments of cumulative perturbations.

Note that we derived this approximation without accounting for the (only) mathematical subtlety of the DDM: namely, the fact that decision bounds formally act as “absorbing barriers” for the system (Broderick et al., 2009). Critically, absorbing barriers

induce some non-trivial forms of dynamical degeneracy. In particular, they eventually favor paths that are made of extreme samples of the perturbation noise. This is because these have a higher chance of crossing the boundary, despite being comparatively less likely than near-zero samples under the corresponding “no barrier” distribution. One may thus wonder whether ignoring absorbing barriers may invalidate the moment-matching approximation given in **Eqs 11–13**. To address this concern, we conducted a series of 1000 Monte-Carlo simulations, where DDM parameters were randomly drawn (each simulation consisted of 200 trials of the same decision system). We use these to compare the sample estimates of first- and second-order moments of normalized cumulative perturbations and their analytical approximations (as given in **Eqs. 11, 12**). The results are given in **Figure 6** below.

One can see on the upper-right panel of **Figure 6** that the distribution of normalized cumulative perturbations may strongly deviate from the standard normal density. In particular, this distribution clearly exhibits two modes, which correspond to correct and incorrect decisions, respectively. We have observed this bimodal shape across

almost all Monte-Carlo simulations. This means that bound hits are less likely to be caused by perturbations of small magnitude than expected under the “no-barrier” distribution (cf. lack of probability mass around zero). Nevertheless, the ensuing approximate conditional distributions seem to be reasonably matched with their sample estimates. In fact, lower panels of **Figure 6** demonstrate that sample means and variances of normalized cumulative perturbations are well approximated by **Eqs 11, 12** for a broad range of DDM parameters. We note that the “no-barrier” approximation tends to slightly underestimate first-order moments, and overestimate second-order moments. This bias is negligible however, when compared to the overall range of variations of conditional moments. In brief, the effect of absorbing barriers on system dynamics has little impact on the conditional moments of normalized cumulative perturbations.

When fitting the DDM to empirical RT data, one thus wants to enforce the distributional constraint in **Eqs 11–13** onto the perturbation term in **Eq. 9**. This can be done using a change of variable $\tilde{\eta}_i = h(\varsigma_i)$, where ς are non-constrained dummy variables and $h: \varsigma_i \rightarrow h(\varsigma_i)$ is the following moment-enforcing mapping:

$$h(\varsigma_i) = \begin{cases} \tilde{\mu}_+ + \left(\varsigma_i - \frac{1}{nQ_+} \sum_{i \in I_+} \varsigma_i \right) \sqrt{\frac{nQ_+ \tilde{\Sigma}_+}{\sum_{i \in I_+} \left(\varsigma_i - \frac{1}{nQ_+} \sum_{i \in I_+} \varsigma_i \right)^2}} & \text{if } i \in I_+ \\ \tilde{\mu}_\# + \left(\varsigma_i - \frac{1}{n(1-Q_+)} \sum_{i \in I_\#} \varsigma_i \right) \sqrt{\frac{n(1-Q_+) \tilde{\Sigma}_\#}{\sum_{i \in I_\#} \left(\varsigma_i - \frac{1}{n(1-Q_+)} \sum_{i \in I_\#} \varsigma_i \right)^2}} & \text{if } i \in I_\# \end{cases} \quad (14)$$

where I_+ and $I_\#$ are the indices of correct and incorrect trials, respectively (and n is the total number of trials). **Eq. 14** ensures that the sample moments of the estimated normalized cumulative perturbations $\tilde{\eta}_i = h(\varsigma_i)$ match **Eqs 11, 12**, irrespective of the dummy variable ς . This also implies that the effective degrees of freedom of the constrained model are in fact lower than what the native self-consistency function would suggest.

b. Prior constraints on native DDM parameters.

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

- The bound's height b is necessarily positive. This positivity constraint can be enforced by replacing b with a non-bounded parameter φ_1 , which relates to the bound's height through the following mapping: $b = \exp(\varphi_1)$. We note that parameters ω of collapsing bounds $b_\omega(t)$ may not have to obey such positivity constraint.
- 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 positive and 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 $s(\cdot)$ is the standard sigmoid mapping.

- The initial bias x_0 is necessarily constrained between $-b$ and b . This can be enforced by replacing the native initial condition with the following mapped parameter φ_4 : $x_0 = \exp(\varphi_1)[2s(\varphi_4) - 1]$.
- In principle, the drift rate v can be either positive or negative. However, its magnitude is necessarily smaller than $\frac{b+|x_0|}{\min(RT)-T_{ND}}$, which corresponds to its “ballistic” limit (see **Supplementary Appendix S6** for more details). This can be enforced by replacing the native drift rate with the following mapped parameter φ_5 : $v = \frac{[1+2s(\varphi_4)-1]\exp(\varphi_1)}{\min(RT)[1-s(\varphi_3)]} [2s(\varphi_5) - 1]$.

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 DDM parameters, and then comparing estimated and simulated parameters. Below, we use such recovery analysis to compare the overcomplete approach with standard DDM fitting procedures.

c. 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 **Eq. 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 questions 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 (Huk and Shadlen, 2005; Bitzer et al., 2014). 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 (Krajbich et al., 2010; De Martino et al., 2012; 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. Let D be a known predictor variable, which is expected to capture trial-by-trial variations in some DDM parameter (e.g., the drift rate). One may then alter the self-consistency equation such that DDM parameters are treated as affine functions of trial-by-trial predictors (e.g., $v_i \triangleq v_0 + v_1 D_i$), and exploit trial-by-trial variations in response times to fit the ensuing offset and slope parameters (here, v_0 and v_1). Alternatively, one can simply set the drift rate to the predictor variable (i.e., assume *a priori* $v_0 = 0$ and $v_1 = 1$), which is currently the favorite approach in the field. 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 response 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 improve the inference on model parameters.

PARAMETER RECOVERY ANALYSIS

In what follows, we use numerical simulations to evaluate the approach's ability to recover DDM parameters. Our parameter recovery analyses proceed as follows. First, we sample 1,000 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 [Eq. 2](#) above. Third, we fit the DDM to each series of simulated reaction times (200 data points) and extract parameter estimates. Last, we compare simulated and estimated parameters to each other. In particular, we measure the relative estimation error for each DDM parameter. We also quantify potential non-identifiability issues using so-called recovery matrices and the ensuing identifiability index. We note that details regarding parameter recovery analyses can be found in [Supplementary Appendix S4](#) of this manuscript (along with definitions of the relative estimation error REE , recovery matrices and identifiability index ΔV).

To begin with, we will focus on “vanilla” DDMs, because they provide a fair benchmark for parameter estimation methods. In this context, we will compare the overcomplete approach with two established methods ([Moens and Zenon, 2017](#); [Boehm et al., 2018](#)), namely: the “method of moments” and the “method of trial means”. These methods are summarized in [Supplementary Appendixes S2, S3](#), respectively. In brief, the former attempts to match empirical and theoretical moments of RT data. We expect this method to perform best when DDM parameters are fixed across trials. The latter rather attempts to match raw trial-by-trial RT data to trial-by-trial theoretical RT means. This will be most reliable when DDM parameters (e.g., the drift rate) vary over trials. Note that, in all cases, we inserted the prior constraints on DDM parameters given in *An Overcomplete Likelihood Approach to DDM Inversion* (section b) above, along with standard normal priors on unmapped parameters $\varphi_{1:5}$. We will therefore compare the ability of these methods to recover DDM parameters (i) when no parameter is fixed (full parameter set), (ii) when the drift rate is fixed, and (iii) when drift rates vary over trials.

Finally, we perform a parameter recovery analysis in the context of a generalized DDM, which includes collapsing bounds. This will serve to demonstrate the flexibility and robustness of the overcomplete approach.

a. Vanilla DDM: recovery analysis for the full parameter set.

First, we compare the three approaches when all DDM parameters have to be estimated. This essentially serves as a reference point for the other recovery analyses. The ensuing recovery analysis is summarized in [Figure 7](#) below, in terms of the comparison between simulated and estimated parameters.

Unsurprisingly, parameter estimates depend on the chosen estimation method, i.e. different methods exhibit distinct estimation errors structures. In addition, estimated and

simulated parameters vary with similar magnitudes, and no systematic estimation bias is noticeable. It turns out that, in this setting, estimation error is minimal for the method of moments, which exhibits lower error than both the overcomplete approach (mean error difference: $\Delta \log(REE) = 0.27 \pm 0.03$, $p < 10^{-4}$, two-sided F-test) and the method of moments (mean error difference: $\Delta \log(REE) = 0.26 \pm 0.02$, $p < 10^{-4}$, two-sided F-test). However, the overcomplete approach and the method of trial means yield comparable estimation errors (mean error difference: $\Delta \log(REE) = 0.006 \pm 0.04$, $p = 0.88$, two-sided F-test).

Now, although estimation errors enable a coarse comparison of methods, it does not provide any quantitative insight regarding potential non-identifiability issues. We address this using recovery matrices (see [Supplementary Appendix S4](#)), which are shown on [Figure 8](#) below.

None of the estimation methods is capable of perfectly identifying DDM parameters (except T_{ND}), i.e., all methods exhibit strong non-identifiability issues. In particular, variations in the perturbations' standard deviation σ are partially confused with variations in the bound's height b , and reciprocally. This is because increasing both at the same time leaves RT trial-by-trial variability unchanged. Therefore, RT produced under strong neural perturbations can be equally well explained with a small bound height (and reciprocally). Interestingly, drift rate estimates are the least reliable: though their amount of “correct variability” is decent for the method of moments (45.3%), it is very low for both the overcomplete approach (5.3%) and the method of trial means (7.5%). If anything, non-identifiability issues are strongest for the overcomplete approach, which also exhibits weak “correct variability” for initial conditions (5.1%).

b. Vanilla DDM: recovery analysis with a fixed drift rate.

In fact, we expect non-identifiability issues of this sort, which were already highlighted in early DDM studies ([Ratcliff, 1978](#)). Note that this basic form of non-identifiability is easily disclosed from the self-consistency equation, which is invariant to a rescaling of all DDM parameters (except T_{ND}). In other terms, response times are left unchanged if all these parameters are rescaled by the same amount. Although this problematic invariance would disappear if a single DDM parameter 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, but this time informing estimation methods about the drift rate, which was set to its simulated value. We note that such arbitrary reduction of the parameter space is routinely performed, as it was already suggested in seminal empirical applications of the DDM ([Ratcliff, 1978](#)). [Figure 9](#) below summarizes the ensuing comparison between simulated and estimated parameters.

Comparing [Figures 7, 9](#) provides a clear insight regarding the impact of reducing the DDM's parameter space. In brief, estimation errors decrease for all methods, which seem to provide much more reliable parameter estimates. The method of moments still yields the most reliable parameter estimates,

FIGURE 7 | Comparison of simulated and estimated DDM parameters (full parameter set). Left panel: Estimated parameters using the overcomplete approach (y -axis) are plotted against simulated parameters (x -axis). Each dot is a monte-carlo simulation and different colors indicate distinct parameters (blue: σ , red: v , yellow: b , purple: x_0 , green: T_{ND}). The black dotted line indicate the identity line (perfect estimation). Middle panel: Method of moments, same format as left panel. Right panel: Method of trial means, same format as left panel.

FIGURE 8 | DDM parameter recovery matrices (full parameter set). Left panel: overcomplete approach. Middle panel: method of moments. Right panel: Method of trial means. Each line shows the squared partial correlation coefficient between a given estimated parameter and each simulated parameter (across 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. Diagonal elements of the recovery matrix measure “correct estimation variability”, i.e., variations in the estimated parameters that are due to variations in the corresponding simulated parameter. In contrast, non-diagonal elements of the recovery matrix measure “incorrect estimation variability”, i.e., variations in the estimated parameters that are due to variations in other parameters. Strong non-diagonal elements in recovery matrices thus signal pairwise non-identifiability issues.

FIGURE 9 | Comparison of simulated and estimated DDM parameters (fixed drift rates). Same format as **Figure 7**, except for the color code in upper panels (blue: σ , red: b , yellow: x_0 , purple: T_{ND}).

FIGURE 10 | DDM parameter recovery matrices (fixed drift rates). Same format as **Figure 8**, except that recovery matrices do not include the line that corresponds to the drift rate estimates. Note, however, that we still account for variations in the remaining estimated parameters that are attributable to variations in simulated drift rates.

eventually exhibiting lower error than the overcomplete approach (mean error difference: $\Delta \log(REE) = 0.21 \pm 0.03$, $p = 0.04$, two-sided F-test) and the method of trial means (mean error difference: $\Delta \log(REE) = 0.53 \pm 0.03$, $p < 10^{-4}$, two-sided F-test). In addition, the overcomplete approach yields lower estimation error than the method of trial means (mean error difference: $\Delta \log(REE) = 0.33 \pm 0.04$, $p < 10^{-4}$, two-sided F-test). The reason why the methods of trial means performs worst here is that it is blind to trial-by-trial variability in the data (beyond mean RT differences between the two decision outcomes). This is not the case however, for the two other methods.

We then evaluated non-identifiability issues using recovery matrices, which are summarized in **Figure 10** below.

Figure 10 clearly demonstrates an overall improvement in parameter identifiability (compare to **Figure 8**). In brief, most parameters are now identifiable, at least for the method of moments (which clearly performs best) and the overcomplete approach. Nevertheless, some weaker non-identifiability issues still remain, even when fixing the drift rate to its simulated value. For example, the overcomplete approach and the method of trial means still somehow confuse bound's heights with perturbations' standard deviations. More precisely, $\hat{\sigma}$ shows unacceptably weak "correct variations" (overcomplete approach: 12.3%, method of trial means: 2.7%), when compared to "incorrect variations" due to the bound's height (overcomplete approach: 12.4%, method of trial means: 14.3%). Note that this does not hold for the method of moments, for which $\hat{\sigma}$ shows strong "correct variations" (30.2%). Having said this, even the method of moments exhibit partial non-identifiability issues, in particular between perturbations' standard deviations and drift rates (incorrect variations: 4.1%).

We note that fixing another DDM parameter, e.g., the noise's standard deviation σ (instead of ν), would not change the relative merits of estimation methods in terms of parameter recovery. In other words, the above results are representative of the impact of fixing any DDM parameter. But situations where the drift rate is fixed can be directly compared with situations where one is attempting to exploit predictable drift rates trial-by-trial variations, which is the focus of the next section.

c. Vanilla DDM: recovery analysis with varying drift rates.

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 perceptual 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 comparison between simulated and estimated parameters is summarized in **Figure 11** below.

On the one hand, the estimation error has now been strongly reduced, at least for the overcomplete approach and the method of trial means. On the other hand, estimation error has increased for the method of moments. This is because the method of moments confuses trial-by-trial variations that are caused by variations in drift rates with those that arise from the DDM's stochastic "neural" perturbation term. This is not the case for the overcomplete approach and the method of trial means. In turn, the method of moments now shows much higher estimation error than the overcomplete approach (mean error difference: $\Delta \log(REE) = 0.55 \pm 0.03$, $p < 10^{-4}$, two-sided F-test) or the method of trial means (mean error difference: $\Delta \log(REE) = 0.83 \pm 0.04$, $p < 10^{-4}$, two-sided F-test). Note that the latter eventually performs slightly better than the overcomplete approach (mean error difference: $\Delta \log(REE) = 0.28 \pm 0.03$, $p = 0.04$, two-sided F-test).

Figure 12 below then summarizes the evaluation of non-identifiability issues, in terms of recovery matrices.

For the overcomplete approach and the method of trial means, **Figure 12** shows a further improvement in parameter identifiability (compare to **Figures 8, 10**). For these two methods, all parameters are now well identifiable ("correct variations" are always greater than 67.2% for all parameters), and no parameter estimate is strongly influenced by other simulated parameters. This is a simple example of the gain in

FIGURE 11 | Comparison of simulated and estimated DDM parameters (varying drift rates). Same format as **Figure 9**.

FIGURE 12 | DDM parameter recovery matrices (varying drift rates). Same format as **Figure 10**, except that fixed drift rates are replaced by their average across DDM trials.

statistical efficiency that result from exploiting known trial-by-trial variations in DDM model parameters. The situation is quite different for the method of moments, which exhibits clear non-identifiability issues for all parameters except the non-decision time. In particular, the bound's height is frequently confused with the perturbations' standard deviation (20.3% of “incorrect variations”), the estimate of which has become unreliable (only 17.6% of “correct variations”).

We note that the gain in parameter recovery that obtains from exploiting predictable trial-by-trial variations in drift rates (with either the method of trial means or the overcomplete approach) does not generalize to situations where drift rates are defined in term of an affine transformation of some predictor variable (see *An Overcomplete Likelihood Approach to DDM Inversion* section. c above). This is because the ensuing offset and slope parameters would then need to be estimated along with other native DDM parameters. In turn, this would reintroduce identifiability issues similar or worse than when the full set of parameters have to be estimated (cf. *An Overcomplete Likelihood Approach to DDM Inversion* section.a). This is why people then typically fix another DDM parameter, e.g., the standard deviation σ (Ratcliff et al., 2016). 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.

d. Generalized DDM: recovery analysis with collapsing bounds.

We now consider generalized DDMs that include collapsing bounds. More precisely, we will consider a DDM where the bound $\hat{b}_w(t)$ is exponentially decaying in time, i.e.: $\hat{b}_w(t) = \exp(\omega_0 - \omega_1 t)$, where ω_0 and ω_1 control the bound's initial height and decay rate, respectively. This DDM variant reduces to the vanilla DDM when $\omega_1 \approx 0$, in which case the parameter ω_0 is formally identical to the vanilla bound's height b . When $\omega_1 \neq 0$ however, collapsing bounds induce a causal dependency between choice accuracy and response times that cannot be captured by the vanilla DDM (Zhang, 2012; Zhang et al., 2014; Hawkins et al., 2015; Tajima et al., 2016; Voskuilen et al., 2016).

In what follows, we report the results of a recovery analysis, in which data was simulated under the above generalized DDM (with drift rates varying across trials). We note that, under such generalized DDM variant, no analytical solution is available to derive RT moments. Applying the method of moments or the method of trial means to such generalized DDM variant thus involves either sampling schemes or numerical solvers for the underlying Fokker-Planck equation (Shinn et al., 2020). However, the computational cost of deriving trial-by-estimates of RT moments precludes routine data analysis using these methods,

FIGURE 13 | Comparison of simulated and estimated DDM parameters (collapsing bounds). Same format as **Figure 9**, except that the left panel includes an additional parameter (w_1 : green color), which controls the decay rate of DDM bounds.

FIGURE 14 | DDM parameter recovery matrices (collapsing bounds). Same format as **Figure 12**, except that recovery matrices now also include the bound's decay rate parameter (w_1), in addition to the bound's initial height (w_0).

which is why most model-based studies are currently restricted to the vanilla DDM (Fengler et al., 2020). In turn, we do not consider here such computationally intensive extensions of the method of moments and/or method of trial means. In this setting, they thus do not rely on the correct generative model. The ensuing estimation errors and related potential identifiability issues should thus be interpreted in terms of the (lack of) robustness against simplifying modeling assumptions. This is not the case for the overcomplete approach, which bypasses this computational bottleneck and hence generalizes without computational harm to such DDM variants.

Figure 13 below summarizes the ensuing comparison between simulated and estimated parameters.

In brief, the overcomplete approach seems to perform as well as for non-collapsing bounds (see **Figure 11**). Expectedly however, the method of moments and the method of trial means do incur some reliability loss. Quantitatively, the overcomplete approach shows much smaller estimation error than the method of moments (mean error difference: $\Delta \log(REE) = 0.88 \pm 0.05$, $p < 10^{-4}$, two-sided F-test) or the method of trial means (mean error difference: $\Delta \log(REE) = 0.61 \pm 0.05$, $p < 10^{-4}$, two-sided F-test).

Figure 14 below then summarizes the ensuing evaluation of non-identifiability issues, in terms of recovery matrices.

For the overcomplete approach, **Figure 14** shows a similar parameter identifiability than **Figure 12**. In brief, all parameters of the generalized DDM are identifiable from each other (the amount of “correct variations” is 33.8% for the bound's decay parameter, and greater than 75.5% for all other parameters). This implies that including collapsing bounds does not impact parameter recovery with this method. This is not the case for the two other methods, however. In particular, the method of moments confuses the perturbations' standard deviation with the bound's decay rate (7.2% “correct variations” against 20.8% “incorrect variations”). This is also true, though to a lesser extent, for the method of trial means (31.6% “correct variations” against 5.4% “incorrect variations”). Again, these identifiability issues are expected, given that neither the method of moments nor the method of trial means (or, more properly, the variant that we use here) rely on the correct generative model. Maybe more surprising is the fact that these methods now exhibit non-identifiability issues w.r.t. parameters that they can, in principle, estimate. This exemplifies the sorts of interpretation issues that arise when relying on methods that neglect decision-relevant mechanisms. We will comment on this and related issues further in the Discussion section below.

e. Summary of recovery analyses.

FIGURE 15 | Summary of DDM parameter recovery analyses. Left panel: The mean log relative estimation error RRE (y-axis) is shown for all methods (OcA: Overcomplete approach, MoM: Method of moments, MoTM: Method of trial means), and all simulation series (black: Full parameter set, blue: fixed drift rate, red: varying drift rates, green: Collapsing bounds). Right panel: The mean identifiability index ΔV (y-axis) is shown for all methods and all simulation series (same format as left panel). Note that the situation in which the full parameter set has to be estimated serves as a Reference point. To enable a fair comparison, both the estimation error and the identifiability index are computed for the parameter subset that is common to all simulation series (i.e.: The perturbations 'standard deviation σ , the bound's height b , the initial condition x_0 , and the non-decision time T_{ND}).

Figure 15 below summarizes all our recovery analyses above, in terms of the average (log-) relative estimation error REE and the parameter identifiability index ΔV (cf. **Supplementary Appendix S4**).

Figure 15 enables a visual comparison of the impact of simulation series on parameter estimation methods. As expected, for the method of moments and the method of trial means, the most favorable situation (in terms of estimation error and identifiability) is when the drift rate is fixed and varying over trials, respectively. This is also when these methods perform best in relation to each other. All other situations are detrimental, and eventually yield estimation error and identifiability issues similar or worse than when the full parameter set has to be estimated. This is not the case for the overcomplete approach, which exhibits comparable estimation error and/or identifiability than the best method in all situations, except for collapsing bounds, where it strongly outperforms the two other methods. Here again, we note that parameter recovery for generalized DDMs may, in principle, be improved for the method of moments and/or the method of trial means. But extending these methods to generalized DDMs is beyond the scope of the current work.

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 or "default" 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). In what follows, we refer to the "default" option as the choice options that belonged to the favored category. Each choice can then be described in terms of choosing between the default and the alternative option.

Figure 16 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.

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 $V_{\text{def}} - V_{\text{alt}}$ between the default and alternative option values ($t = 8.4$, $\text{dof} = 23$, $p < 10^{-4}$). In addition, choice bias is significant at the group-level ($t = 8.7$, $\text{dof} = 23$, $p < 10^{-4}$). Similarly, RT significantly decreases with absolute decision value $|V_{\text{def}} - V_{\text{alt}}|$ ($t = 8.7$, $\text{dof} = 23$, $p < 10^{-4}$), and RT bias is significant at the group-level ($t = 7.4$, $\text{dof} = 23$, $p < 10^{-4}$).

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 $V_{\text{def}} - V_{\text{alt}}$ or $V_{\text{right}} - V_{\text{left}}$). It turns out that within-subject estimates of σ , b and T_{ND} do not

FIGURE 16 | 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 $V_{\text{def}} - V_{\text{alt}}$ (x-axis), divided into 10 bins. Values correspond to likeability ratings given by the subject prior to choice session. For each participant, the 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 (i.e., when $V_{\text{def}} = V_{\text{alt}}$). Right: Response time RT (y-axis) is plotted as a function of the absolute decision value $|V_{\text{def}} - V_{\text{alt}}|$ (x-axis) divided into 10 bins, separately for trials in which the default option was chosen (black) or not (red). For each participant, the RT bias was defined as the difference between the RT intercepts (when $V_{\text{def}} = V_{\text{alt}}$) observed for each choice outcome.

FIGURE 17 | 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.

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 : $r = 0.76$, $p < 10^{-4}$; b : $r = 0.82$, $p < 10^{-4}$; T_{ND} : $r = 0.94$, $p < 10^{-4}$). This implies that inter-individual differences in σ , b 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 ($F = 45.2$, $\text{dof} = [1, 23]$, $p < 10^{-4}$) but not for the native frame ($F = 2.36$, $\text{dof} = [1, 23]$, $p = 0.14$). 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 17 below.

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 or the method of trial means, which were not able to capture these inter-individual differences (see **Supplementary Appendix S7** for details).

Finally, can we exploit model fits to provide a normative argument for why the brain favors a biased choice frame? Recall that, if properly set, the DDM can implement the optimal speed-accuracy tradeoff inherent in making online value-based decisions (Tajima et al., 2016). Here, 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 b . This is because b determines the demand for evidence before the system commits to a decision. More precisely, the system can favor decision accuracy by increasing b , or improve decision speed by decreasing b . We thus defined a measure $\hat{e}$ 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 b (see **Supplementary Appendix SA5** 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 ($\Delta\hat{e} = 0.007 \pm 0.003$, $t = 2.2$, $\text{dof} = 23$, $p = 0.02$). In other words, participants' decisions appear more optimal under the default/alternative frame than under the native frame. We comment on possible interpretations of this result in the Discussion section below.

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 that response times obey under DDM models. It bypasses the computational bottleneck of existing DDM parameter estimation approaches, at the cost of augmenting the model with stochastic neural noise variables that perturb the underlying decision process. This makes it suitable for generalized variants of the DDM, which would not otherwise be considered for behavioral data analysis.

Strictly speaking, the DDM predicts the RT distribution conditional on choice outcomes. This is why variants of the method of moments are not optimal when empirical design parameters (e.g., evidence strength) are varied on a trial-by-trial basis. 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 (Boehm et al., 2018; Ratcliff, 2008; Srivastava et al., 2016). Alternatively, one could rely on variants of the method of trial means to find the DDM parameters that best match expected and observed RTs (Fontanesi et al., 2019a; Fontanesi et al., 2019b; Gluth and Meiran, 2019; Moens and Zenon, 2017; Pedersen et al., 2017; Wabersich and Vandekerckhove, 2014). But this becomes computationally cumbersome when the number of trials is high and one wishes to use generalized variants of the DDM. This however, is not the case for the overcomplete approach. As with the method of trial means, its statistical power is maximal when design parameters are varied on a trial-by-trial basis. But the overcomplete approach does not suffer from the same computational bottleneck. This is because evaluating the underlying self-consistency equation (Eqs. 7–9) is much simpler than deriving moments of the conditional RT distributions (Broderick et al., 2009; Navarro and Fuss, 2009). In turn, the statistical added-value of the overcomplete approach is probably highest for analyzing data acquired with such designs, under generalized DDM variants.

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 modeling studies restrict 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. 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?

In brief, the overcomplete approach effectively reduces to adjusting DDM parameters such that RT become self-consistent. Interestingly, we derived the self-consistency equation without regard to the subtle dynamical degeneracies that (absorbing) bounds induce on stochastic processes (Broderick et al., 2009). It simply follows from noting that if a decision is triggered at time τ , then the underlying stochastic process has reached the bound (i.e., $x_\tau = \pm b$). This serves to identify the cumulative perturbation that eventually drove the system toward the bound. But a bound hit event at time τ is more informative about the history of the stochastic process than just its fate: it also tells us that the path did not cross the barrier before (i.e., $|x_t| < b \quad \forall t < \tau$). This disqualifies those sample paths whose first-passage time happens sooner, even though all barrier crossings are (by definition) “self-consistent”. In retrospect, one may thus wonder whether the self-consistency equation may be suboptimal, in the sense of incurring some loss of information. Critically however, no information is lost about cumulative perturbations (or about DDM parameters). Although these are not sufficient to discriminate between the many sample paths that are compatible with a given RT, this is essentially irrelevant to the objective of the overcomplete approach. In turn, the existing limitations of the overcomplete approach lie elsewhere.

First and foremost, the self-consistency equation cannot be used to simulate data (recall that RTs appear on both the left- and right-hand sides of the equation). This restricts the utility of the approach to data analysis. Note however, that data simulations can still be performed using Eq. 2, once the model parameters have been identified. This enables all forms of posterior predictive checks and/or other types of model fit diagnostics (Palminteri et al., 2017). Second, the accuracy of the method depends upon the reliability of response 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 Eq. 7). In addition, the method inherits the potential limitations of its underlying parameter estimation technique: namely, the variational Laplace approach (Friston et al., 2007; Daunizeau, 2017). In particular, and as is the case for any numerical optimization scheme, it is not immune to multimodal likelihood landscapes. We note that this may result in non-identifiability issues of the sort that we have demonstrated here (cf., e.g., **Figures 8, 10**). One cannot guarantee that this will not happen for some generalized DDM variant of interest. A possible diagnostic to this problem is to perform a systematic fit/sample/refit analysis to evaluate the stability of parameter estimates. In any case, we would advise to re-evaluate (and

report) parameter recovery for any novel DDM variant. Third, the computational cost of model inversion scales with the number of trials. This is because each trial has its own nuisance perturbation parameter. Note however, that the ensuing computational cost is many orders of magnitude lower than that of standard methods for generalized DDM variants. 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 DDM variants can be discriminated using the overcomplete approach is beyond the scope of the current work.

Let us now discuss the results of our model-based data analysis from the value-based decision making experiment (Lopez-Persem et al., 2016). Recall that we eventually provided evidence that peoples' decisions are more optimal under the default/alternative frame than under the native frame. Recall that this efficiency gain is inherited from the initial condition parameter x_0 , which turns out to be significant under the default/alternative frame. The implicit interpretation here is that the brain's decision system starts with a prior bias toward the default option. Critically however, we would have obtained the exact same results, would we have fixed the initial condition to zero but allowed upper and lower decision bounds to be asymmetrical. This is interesting, because it highlights a slightly different interpretation of our results. Under this alternative scenario, one would state that the brain's decision system is comparatively less demanding regarding the evidence that is required for committing to the default option. In turn, the benefit of lowering the bound for the default option may simply be to speed up decisions when evidence is congruent with default preferences, at the expense of slowing down incongruent decisions. Importantly, this strategy does not compromise decision accuracy if the incongruent decisions are rarer than the congruent ones (as is effectively the case in this experiment).

At this point, we would like to discuss potential neuroscientific applications of trial-by-trial estimates of "neural" perturbation terms. Recall that the self-consistency equation makes it possible to infer these neural noise variables from response times (cf. Eq. 7 or 9). For the purpose of behavioral data analysis, where one is mostly interested in native DDM parameters, these are treated as nuisance variables. However, should one acquire neuroimaging data concurrently with behavioral data, one may want to exploit this unique feature of the overcomplete approach. In brief, estimates of "neural" perturbation terms moves the DDM one step closer to neural data. This is because DDM-based analysis of behavioral data now provides quantitative trial-by-trial predictions of an underlying neural variable. This becomes particularly interesting when internal variables (e.g., drift rates) are systematically varied over trials, hence de-correlating the neural predictor from response times. For example, in the context of fMRI investigations of value-based decisions, one may search for brain regions whose activity eventually

perturbs the computation and/or comparison of options' values. This would extend the portfolio of recent empirical studies of neural noise perturbations to learning-relevant computations (Drugowitsch et al., 2016; Wyart and Koechlin, 2016; Findling et al., 2019). Reciprocally, using some variant of mediation analysis (MacKinnon et al., 2007; Lindquist, 2012; Brochard and Daunizeau, 2020), one may extract neuroimaging estimates of neural noise that can inform DDM-based behavioral data analysis. Alternatively, one may model neural and behavioral data in a joint and symmetrical manner, with the purpose of testing some predefined DDM variant (Rigoux and Daunizeau, 2015; Turner et al., 2015).

Finally, one may ask how generalizable the overcomplete approach is? Strictly speaking, one can evaluate the self-consistency equation under any DDM variant, as long as the mapping $z : x \rightarrow z(x)$ from the base random walk to the bound subspace is invertible (cf. Eqs. 8, 9). No such formal constraint exists for the dynamical form of the collapsing bound. This spans a family of DDM variants that is much broader than what is currently being used in the field (Fengler et al., 2020; Shinn et al., 2020). For example, this family includes decision models that trigger a decision when decision *confidence* reaches a bound (Tajima et al., 2016; Lee and Daunizeau, 2020). To the best of our knowledge, there is not a single example of existing DDM variants that does not belong to this class. Having said this, future extensions of the DDM framework may render the current overcomplete approach obsolete. Our guess is that such DDM improvements may then need to be informed with additional behavioral data, such as decision confidence (De Martino et al., 2012) and/or mental effort (Lee and Daunizeau, 2020), for which other kinds of self-consistency equations may be derived.

To conclude, 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).

DATA AVAILABILITY STATEMENT

The datasets presented in this article are not readily available because they were not acquired by the authors. Requests to access the datasets should be directed to jean.daunizeau@inserm.fr.

ETHICS STATEMENT

Ethical review and approval was not required for the reuse of data from human participants in accordance with the local legislation. The patients/participants provided their written informed consent to participate in this study.

AUTHOR CONTRIBUTIONS

All authors listed have made a substantial, direct, and intellectual contribution to the work and approved it for publication.

ACKNOWLEDGMENTS

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

REFERENCES

- Balci, F., Simen, P., Niyogi, R., Saxe, A., Hughes, J. A., Holmes, P., et al. (2011). Acquisition of decision making criteria: reward rate ultimately beats accuracy. *Atten. Percept. Psychophys.* 73, 640–657. doi:10.3758/s13414-010-0049-7
- Beal, M. J. (2003). Variational algorithms for approximate Bayesian inference/. PhD Thesis. London: . University College London.
- 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, 102. doi:10.3389/fnhum.2014.00102
- Boehm, U., Annis, J., Frank, M. J., Hawkins, G. E., Heathcote, A., Kellen, D., et al. (2018). Estimating across-trial variability parameters of the diffusion decision model: expert advice and recommendations. *J. Math. Psychol.* 87, 46–75. doi:10.1016/j.jmp.2018.09.004
- Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and Cohen, J. D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol. Rev.* 113, 700–765. doi:10.1037/0033-295x.113.4.700
- Brochard, J., and Daunizeau, J. (2020). Blaming blunders on the brain: can indifferent choices be driven by range adaptation or synaptic plasticity? *BioRxiv*, 287714. doi:10.1101/2020.09.08.287714
- 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. doi:10.1093/amrx/abp008
- Daunizeau, J., Adam, V., and Rigoux, L. (2014). VBA: a probabilistic treatment of nonlinear models for neurobiological and behavioral data. *Plos Comput. Biol.* 10, e1003441. doi:10.1371/journal.pcbi.1003441
- Daunizeau, J. (2017). The variational Laplace approach to approximate Bayesian inference. arXiv:1703.02089.
- Daunizeau, J. (2019). Variational Bayesian modeling of mixed-effects. arXiv:1903.09003.
- De Martino, B., Fleming, S. M., Garrett, N., and Dolan, R. J. (2012). Confidence in value-based choice. *Nat. Neurosci.* 16, 105–110. doi:10.1038/nn.3279
- 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. doi:10.1523/jneurosci.4010-11.2012
- Drugowitsch, J., Wyart, V., Devauchelle, A.-D., and Koechlin, E. (2016). Computational precision of mental inference as critical source of human choice suboptimality. *Neuron* 92, 1398–1411. doi:10.1016/j.neuron.2016.11.005
- Fengler, A., Govindarajan, L. N., Chen, T., and Frank, M. J. (2020). Likelihood approximation networks (LANs) for fast inference of simulation models in cognitive neuroscience. *BioRxiv*. doi:10.1101/2020.11.20.392274
- Findling, C., Chopin, N., and Koechlin, E. (2019). Imprecise neural computations as source of human adaptive behavior in volatile environments. *BioRxiv*, 799239.
- Fontanesi, L., Gluth, S., Spektor, M. S., and Rieskamp, J. (2019a). A reinforcement learning diffusion decision model for value-based decisions. *Psychon. Bull. Rev.* 26, 1099–1121. doi:10.3758/s13423-018-1554-2
- Fontanesi, L., Palminteri, S., and Lebreton, M. (2019b). Decomposing the effects of context valence and feedback information on speed and accuracy during reinforcement learning: a meta-analytical approach using diffusion decision modeling. *Cogn. Affect. Behav. Neurosci.* 19, 490–502. doi:10.3758/s13415-019-00723-1
- Frank, M. J., Gagne, C., Nyhus, E., Masters, S., Wiecki, T. V., Cavanagh, J. F., et al. (2015). fMRI and EEG predictors of dynamic decision parameters during human reinforcement learning. *J. Neurosci.* 35, 485–494. doi:10.1523/jneurosci.2036-14.2015
- Friston, K., Mattout, J., Trujillo-Barreto, N., Ashburner, J., and Penny, W. (2007). Variational free energy and the Laplace approximation. *NeuroImage* 34, 220–234. doi:10.1016/j.neuroimage.2006.08.035
- Gluth, S., and Meiran, N. (2019). Leave-One-Trial-Out, LOTO, a general approach to link single-trial parameters of cognitive models to neural data. *eLife Sciences* 8. doi:10.7554/eLife.42607
- Gold, J. I., and Shadlen, M. N. (2007). The neural basis of decision making. *Annu. Rev. Neurosci.* 30, 535–574. doi:10.1146/annurev.neuro.29.051605.113038
- Goldfarb, S., Leonard, N. E., Simen, P., Caicedo-Núñez, C. H., and Holmes, P. (2014). A comparative study of drift diffusion and linear ballistic accumulator models in a reward maximization perceptual choice task. *Front. Neurosci.* 8, 148. doi:10.3389/fnins.2014.00148
- Grasman, R. P. P., Wagenmakers, E.-J., and van der Maas, H. L. J. (2009). On the mean and variance of response times under the diffusion model with an application to parameter estimation. *J. Math. Psychol.* 53, 55–68. doi:10.1016/j.jmp.2009.01.006
- Guevara Erra, R., Arbotto, M., and Schurger, A. (2019). An integration-to-bound model of decision-making that accounts for the spectral properties of neural data. *Sci. Rep.* 9, 8365. doi:10.1038/s41598-019-44197-0
- Hanks, T., Kiani, R., and Shadlen, M. N. (2014). A neural mechanism of speed-accuracy tradeoff in macaque area LIP. *ELife* 3, e02260. doi:10.7554/elifelife.02260
- 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. doi:10.1523/jneurosci.2410-14.2015
- 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.* 25, 10420–10436. doi:10.1523/jneurosci.4684-04.2005
- Kloeden, P. E., and Platen, E. (1992). *Numerical solution of stochastic differential equations*. Berlin Heidelberg: Springer-Verlag.
- Krajchich, I., Armel, C., and Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nat. Neurosci.* 13, 1292–1298. doi:10.1038/nn.2635
- Lee, D., and Daunizeau, J. (2020). Trading mental effort for confidence: the metacognitive control of value-based decision-making. *BioRxiv* 837054.
- Lindquist, M. A. (2012). Functional causal mediation analysis with an application to brain connectivity. *J. Am. Stat. Assoc.* 107, 1297–1309. doi:10.1080/01621459.2012.695640
- 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. doi:10.7554/elifelife.20317
- MacKinnon, D. P., Fairchild, A. J., and Fritz, M. S. (2007). Mediation analysis. *Annu. Rev. Psychol.* 58, 593. doi:10.1146/annurev.psych.58.110405.085542
- 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.
- Moen, V., and Zenon, A. (2017). Variational treatment of trial-by-trial drift-diffusion models of behavior. *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. doi:10.1016/j.jmp.2009.02.003
- Newey, W. K., and West, K. D. (1987). Hypothesis testing with efficient method of moments estimation. *Int. Econ. Rev.* 28, 777–787. doi:10.2307/2526578
- Osth, A. F., Bora, B., Dennis, S., and Heathcote, A. (2017). Diffusion vs. linear ballistic accumulation: different models, different conclusions about the slope of the zROC in recognition memory. *J. Mem. Lang.* 96, 36–61. doi:10.1016/j.jml.2017.04.003
- Palminteri, S., Wyart, V., and Koechlin, E. (2017). The importance of falsification in computational cognitive modeling. *Trends Cogn. Sci.* 21, 425–433. doi:10.1016/j.tics.2017.03.011

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/frai.2021.531316/full#supplementary-material>.

- 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. doi:10.3758/s13423-016-1199-y
- Pedersen, M. L., and Frank, M. J. (2020). Simultaneous hierarchical bayesian parameter estimation for reinforcement learning and drift diffusion models: a tutorial and links to neural data. *Comput. Brain Behav.* 3, 458–471. doi:10.1007/s42113-020-00084-w
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychol. Rev.* 85, 59–108. doi:10.1037/0033-295x.85.2.59
- Ratcliff, R., and McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Comput.* 20, 873–922. doi:10.1162/neco.2008.12.06-420
- 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. doi:10.1016/j.tics.2016.01.007
- Ratcliff, R. (2008). The EZ diffusion method: too EZ? *Psychon. Bull. Rev.* 15, 1218–1228. doi:10.3758/pbr.15.6.1218
- 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. doi:10.3758/bf03196302
- Resulaj, A., Kiani, R., Wolpert, D. M., and Shadlen, M. N. (2009). Changes of mind in decision-making. *Nature* 461, 263–266. doi:10.1038/nature08275
- Rigoux, L., and Daunizeau, J. (2015). Dynamic causal modeling of brain-behavior relationships. *NeuroImage* 117, 202–221. doi:10.1016/j.neuroimage.2015.05.041
- Shinn, M., Lam, N. H., and Murray, J. D. (2020). A flexible framework for simulating and fitting generalized drift-diffusion models. *ELife* 9, e56938. doi:10.7554/elife.56938
- Srivastava, V., Holmes, P., and Simen, P. (2016). Explicit moments of decision times for single- and double-threshold drift-diffusion processes. *J. Math. Psychol.* 75, 96–109. doi:10.1016/j.jmp.2016.03.005
- Tajima, S., Drugowitsch, J., and Pouget, A. (2016). Optimal policy for value-based decision-making. *Nat. Commun.* 7, 12400. doi:10.1038/ncomms12400
- Turner, B. M., van Maanen, L., and Forstmann, B. U. (2015). Informing cognitive abstractions through neuroimaging: the neural drift diffusion model. *Psychol. Rev.* 122, 312–336. doi:10.1037/a0038894
- Vandekerckhove, J., and Tuerlinckx, F. (2008). Diffusion model analysis with MATLAB: a DMAT primer. *Behav. Res.* 40, 61–72. doi:10.3758/brm.40.1.61
- 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. doi:10.1016/j.jmp.2016.04.008
- Voss, A., and Voss, J. (2007). Fast-dm: a free program for efficient diffusion model analysis. *Behav. Res. Methods* 39, 767–775. doi:10.3758/bf03192967
- Wabersich, D., and Vandekerckhove, J. (2014). Extending JAGS: a tutorial on adding custom distributions to JAGS (with a diffusion model example). *Behav. Res.* 46, 15–28. doi:10.3758/s13428-013-0369-3
- Wagenmakers, E.-J., van der Maas, H. L. J., Dolan, C. V., and Grasman, R. P. P. (2008). EZ does it! Extensions of the EZ-diffusion model. *Psychon. Bull. Rev.* 15, 1229–1235. doi:10.3758/pbr.15.6.1229
- Wagenmakers, E.-J., van der Maas, H. L. J., and Grasman, R. P. P. (2007). An EZ-diffusion model for response time and accuracy. *Psychon. Bull. Rev.* 14, 3–22. doi:10.3758/bf03194023
- Wiecki, T. V., Sofer, I., and Frank, M. J. (2013). HDDM: hierarchical bayesian estimation of the drift-diffusion model in Python. *Front. Neuroinformatics* 7, 14. doi:10.3389/fninf.2013.00014
- Wyart, V., and Koechlin, E. (2016). Choice variability and suboptimality in uncertain environments. *Curr. Opin. Behav. Sci.* 11, 109–115. doi:10.1016/j.cobeha.2016.07.003
- Zhang, J. (2012). The effects of evidence bounds on decision-making: theoretical and empirical developments. *Front. Psychol.* 3, 263. doi:10.3389/fpsyg.2012.00263
- Zhang, S., Lee, M. D., Vandekerckhove, J., Maris, G., and Wagenmakers, E.-J. (2014). Time-varying boundaries for diffusion models of decision making and response time. *Front. Psychol.* 5, 1364. doi:10.3389/fpsyg.2014.01364

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2021 Feltgen and Daunizeau. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.