nl-DDM: a non-linear drift-diffusion model accounting for the dynamics of single-trial perceptual decisions

Isabelle Hoxha[1,2*], **Sylvain Chevallier**[3], **Matteo Ciarchi**[4], **Stefan Glasauer**[5], **Arnaud Delorme**[6], **Michel-Ange Amorim**[1,2]

[1]CIAMS, Université Paris-Saclay; [2]CIAMS, Université d'Orléans; [3]LISV, Université Paris-Saclay; [4]Max-Planck Institute for the Physics of Complex Systems, Dresden, Germany; [5]Brandenburgische Technische Universität Cottbus-Senftenberg; [6]CerCo, CNRS, Université Toulouse III - Paul Sabatier, Toulouse, France

isabelle.hoxha@universite-paris-saclay.fr

Abstract The Drift-Diffusion Model (DDM) is widely accepted for two-alternative forced-choice decision paradigms thanks to its simple formalism, straightforward interpretation, and close fit to behavioral and neurophysiological data. However, this formalism presents strong limitations to capture inter-trial dependency and dynamics at the single-trial level. We propose a novel model, the non-linear Drift-Diffusion Model (nl-DDM), that addresses these issues by allowing the existence of several trajectories to the decision boundary. We show that the fitting accuracy of our model is comparable to the accuracy of the DDM, with the non-linear model performing better than the drift-diffusion model for an equivalent complexity. To give better intuition on the meaning of nl-DDM parameters, we compare the DDM and the nl-DDM through correlation analysis. This paper provides evidence of the functioning of our model as an extension of the DDM. Our model paves the way toward more accurately analyzing single-trial dynamics for perceptual decisions and accounts for pre- and post-stimulus influences.

Introduction

Perceptual decision-making has been studied extensively from behavioral (Ratcliff and McKoon, 2008; Ratcliff and Smith, 2004), neurophysiological (Gold and Shadlen, 2001), and computational (Gold and Shadlen, 2007) perspectives, as it is omnipresent in daily activities. When decisions are timed, evidence accumulation models describe human and animal behavior well. They assume that decisions are made when enough sensory evidence from the external world has been gathered. Typically, evidence is accumulated at a given rate (or *drift*) until reaching a decision boundary, triggering an action.

Among them, the Drift-Diffusion Model (DDM) (Ratcliff, 1978) suggests that evidence is accumulated linearly, that is, with a constant drift. The accumulation is additionally subject to Gaussian noise; hence the decision state can be seen as a particle following a Brownian motion. The popularity of this model yields from its intuitive and straightforward formalism and its good fit to behavioral (Ratcliff and McKoon, 2008) and neurophysiological data (Gold and Shadlen, 2001). It has also been shown that the DDM formalizes the optimal strategy for decision-making under time constraints (Bogacz et al., 2006; Moehlis et al., 2004). Interestingly, other forms of decision models such as the Leaky-Competing Accumulator model (Usher and McClelland, 2001), and even attractor models

(Wang, 2002; Ditterich et al., 2003) can be formulated equivalently to the DDM or are similar to it under certain performance constraints (Bogacz et al., 2006).

The initial version of the DDM accounts for global statistics of the behavior. More specifically, it describes the Response Times (RT) distribution and the error rate. A major limitation of this model is that this simple form does not take into account inter-trial variability. However, behavioral studies have shown sequential effects (Abrahamyan et al., 2016, for example) which impact prior expectations on the decisions and the subsequent decision process (Glaze et al., 2015). Traditionally, prior expectations on the decision are modelled through the starting point, or *bias*, of the accumulation process (Ratcliff, 1978). Recent accounts have also suggested that choice history affects subsequent drift rates (Urai et al., 2019). Taken together, these studies suggest that these parameters could be intertwined and that they can vary throughout an experiment, as participants are more experienced in the task. To address this issue, (Ratcliff and Rouder, 1998, 2000) proposed an extended form of the DDM, which uses a uniform distribution of starting points and a Gaussian distribution of drifts without explicit dependence between them. However, this only provides global statistics about perceptual responses, without insight at the single-trial level or on inter-trial interactions. Moreover, the linear dynamics do not describe the variation of the dynamics at the scale of the single decision, which seems inconsistent with the aforementioned physiological and behavioral (empirical) observations.

Linear evidence accumulation also assumes that evidence accumulation is independent of the evidence that has already been gathered, or of the time that passes. While some models take into account the effect of time on the decision parameters (Cisek et al., 2009), or dynamics close to the threshold (Busemeyer and Townsend, 1993; Schurger, 2018), no model to our knowledge allows for an account of initial dynamics. For example, ambiguous stimuli could yield flat initial drifts. This is in part translated into non-decision time, as it is assumed to be a time during which sensory evidence is processed in the brain without contributing to the decision process.

In addition, the DDM also assumes that the response only occurs after a decision has been made. Mathematically speaking, it means that the decision variable has reached a decision boundary. However, paradigms that show spontaneous change of mind indicate that responses can occur before the final decision has been reached and that a decision can change under ambiguous stimuli after enough time (Pleskac and Busemeyer, 2010). This can only occur if decision and motor processes overlap. The DDM, however, assumes that they happen sequentially. In addition, the DDM would explain spontaneous change of mind by the presence of noise in the system. In reality, error-correcting behaviors (Rabbitt, 1966) indicate the existence of more explainable processes underlying these changes.

Previous attempts at single-trial fitting of decisions have been made through attractor models (Wang, 2002; Wong and Wang, 2006; Wong et al., 2007), and it has also been shown, using some simplifying assumptions, that these models can be put in the form of a generalized Drift-Diffusion Model (Shinn et al., 2020b), that is in that case, a Langevin equation with a non-linear drift (Roxin and Ledberg, 2008). It has been shown that this model can be reduced to the DDM in certain cases (Bogacz et al., 2006), but that it dynamics allows for transitions between decision states under fluctuating stimuli (Prat-Ortega et al., 2021). However, the link between each parameter and the dynamics of the model is complicated to interpret. Moreover, the reduction proposed assumes a reflection symmetry of the network to obtain the given form. This, however, seems limiting in particular in the case where each perceptual decision recruits different sensory modalities.

Here we propose a straightforward one-dimensional non-linear form to address these limitations: the non-linear Drift-Diffusion Model (nl-DDM). It recreates double-well-like dynamics from an evidence-accumulation perspective, without assuming reflection symmetry. We show its validity and compare its fitting performances to these of the DDM. We first provide a formal description of the nl-DDM, relating it to the DDM. Then, we fit them on two human behavior datasets: one that was already published (Wagenmakers et al., 2008) where participants classified words into two categories (existing vs. invented), and one that we collected ourselves that consists of a classification task recruiting two different sensory modalities. Last, we compared the parameters of both models to provide an empirical explanation of the effect of each of the nl-DDM parameters with analogies on the DDM by showing correlation on fitted parameters on the same data. We show that it fits data equally well as the DDM while providing drift variability like the extended DDM. The dependency of the drift rate on the current decision state provides a framework for more refined analyses of the decision process. We provide open-source code that is directly pluggable onto the PyDDM toolbox (Shinn et al., 2020b) for reproducibility and easy use of our model.

Results

In this paper we introduce a model, the non-linear Drift-Diffusion Model (nl-DDM), that, similarly to the DDM, can be formulated through a Langevin equation. This model takes the form $dx = -k(x-a)(x-z)(x+a)dt+N(t)$, where the decision variable follows an infinitesimal change of dx during the time interval dt. More details on the formalism of this model can be found in the Methods section of this paper.

We show that the nl-DDM performs better than the DDM in terms of fitting accuracy and theoretical predictions on behavior. To do that, we fitted both models on two datasets: a classification task we designed and a dataset published previously in Wagenmakers et al. (2008). To provide more insight into the empirical meaning of the parameters beyond the formalism, we performed correlation analyses between nl-DDM and DDM parameters. The link between models is hence explicitly exposed.

nl-DDM formalism

Our goal was to propose a simple model in which trajectories are naturally attracted to a boundary. Placing ourselves in the context of two-alternative (forced) choice paradigms, our model needed two attractive states. In one dimension, this forces the existence of an unstable fixed point between the two stable fixed points making the stable states (Strogatz, 2015). These models are widely used in classical and quantum mechanics (Jelic and Marsiglio, 2012). For a simple analogy, we imagine that the decision variable is a ball traveling on valleys and hills. The stable points represent points downhill from which the decision variable cannot escape without a substantial uphill input. Two distinct valleys can exist only if there is a hill separating them. This profile is called a double-well potential profile.

Therefore, the model we propose follows a Langevin equation, as the DDM does, but this time the drift varies with the state of the decision instead of being constant. The drift equation can be written in the following form:

$$dx = -k(x+a)(x-z)(x-a)dt + N(t)\,, \tag{1}$$

where x represents the decision variable and dx its variation in infinitesimal time dt, as previously seen on the DDM (Equation (6)). $N(t)$ is a Gaussian white noise term, characterized in the same way as in the DDM and relates similarly to the accuracy. The term $-k(x+a)(x-z)(x-a)$ represents the drift, and depends itself on several parameters. The parameter k can be seen as a time constant of the system, and a and z determine where the attractors, or decision boundaries, lie. $\pm a$ represent the two attractive states, and we constraint z to the interval $]-a, a[$ to obtain three distinct fixed points to the differential equation with z the unstable fixed point. In this case, the drift corresponds to the deterministic part of the equation, and is dependent on the current decision state. A summary of the parameters of the nl-DDM is given in Figure 1, which can be compared to the description of the DDM we provided in Figure 10. In the following, we provide a formal explanation of the meaning of each parameter.

The interpretation of k as a time constant is straightforward from the equation: as k increases, a decision is reached faster for any given set of parameters. This is the closest parameter to the constant drift v in the DDM.

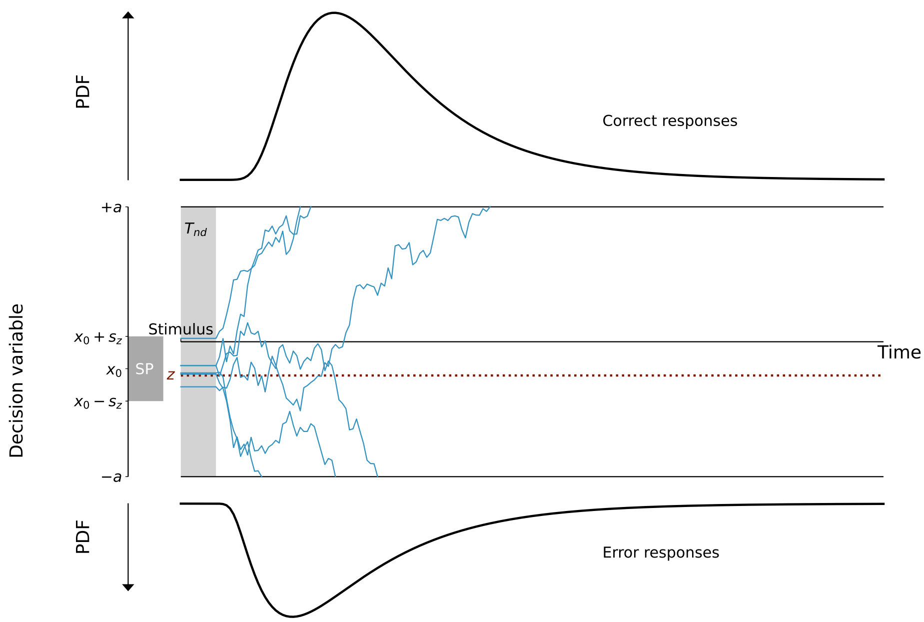

Figure 1. Description of the Non-linear Drift-Diffusion Model (nl-DDM). The decision state is represented by a decision variable x traveling from a starting point (for example, drawn from a uniform distribution, centered around x_0 and of width $2s_z$. It is represented as "SP" on the figure) to a boundary ("Correct boundary" or "Incorrect boundary") under the influence of a drift. Here, the drift depends on the current state of the decision. Depending on the position of x_0 relative to z, the drift will hence have different shapes. The trajectory is also impacted by white noise so that real trajectories are similar to the thin blue lines. From the stimulus onset, the decision process is delayed by a certain non-decision time (T_{nd}). Over an ensemble of decisions, probability density functions of correct and error response times can be created, as displayed here.

In order to provide an intuition for the other parameters, we consider first the potential function derived from the drift term (Figure 2). It is a function $V(x)$ defined from a drift $v(x)$ such that:

$$v(x) = -\frac{\partial V}{\partial x} \,. \tag{2}$$

In our case, we therefore have:

$$V(x) = k\left(\frac{1}{4}x^4 - \frac{z}{3}x^3 - \frac{a^2}{2}x^2 + a^2 zx\right) . \tag{3}$$

The decision variable can be seen as a ball traveling along the potential function.

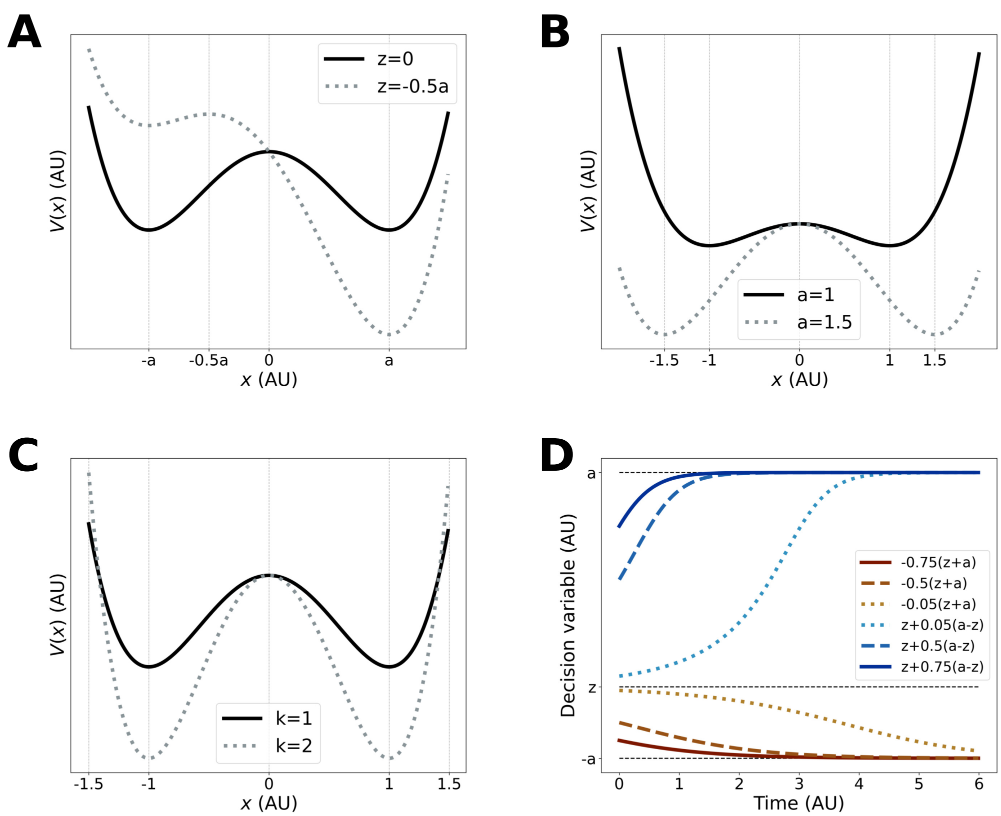

Figure 2. Parameter manipulation on the nl-DDM. A, B, C: Potential functions of the nl-DDM for different z (A. Shifting z changes the relative attractiveness of each boundary, a (B. Shifting a changes the accuracy and the speed of decisions), and k (C. Shifting k changes the speed of decisions). The parameters are always the same for the solid black curve: $a = 1, k = 1, z = 0$, allowing for comparison of the effects of the different parameters. D: Trajectories in the absence of noise for different values of x_0, under $a = 1, k = 1, z = 0$. It becomes clear that the drift range for each trajectory depends on the starting point. The trajectory approach the boundary asymptotically, and will eventually be crossed since noise is omnipresent.

From Figure 2, we can see that there are two potential sinks at a and $-a$, as well as a source at z, which derive directly from the topology of the system. Therefore, $\pm a$ are the decision boundaries and controls along with z the speed-accuracy trade-off. Taking again a as the boundary for correct responses and $-a$ that for incorrect ones, we can see that moving z closer to $-a$ makes the $-a$ well shallower and the well in a deeper (Figure 2A). In other words, the correct decision becomes more attractive than the incorrect one. The gradient becoming more positive on the interval $[z, a]$, the trajectories starting on that interval also reach the correct decision faster.

By reducing the boundary separation, that is, reducing a, both wells become shallower, making decisions slower (Figure 2B). However, for a given noise scale, this also means that any perturbation in the wrong direction is easier to correct because a small perturbation in the other direction can counterbalance that effect. This is not as much the case when the wells are deep because then the decision variable is driven rapidly to the stable fixed point, making perturbations less reversible.

We can also observe the impact of k on the potential function in Figure 2C. Similar to the DDM, fitting of response times can be obtained by solving the Fokker-Planck equation corresponding to the Langevin equation defined above (Shinn et al., 2020b). Then, a non-decision time T_{nd} comes into play in order to shift the resulting distribution to account for biological transmission delays.

To better understand the parameters of our model in comparison to the DDM, it can be useful to define a mean drift rate across all trajectories. Since the deterministic trajectories only approach the decision boundary asymptotically, we define an estimate of the mean drift rate. Considering that the maximum drift for each trajectory causes the largest variation in decision value, we can approximate the mean drift of each trajectory by its maximum, and subsequently average over all the trajectories to get an estimate of the mean drift. Put in equations, we obtain:

$$\bar{v} = \frac{1}{x_- + a}\int_{-a}^{x_-} -k(x_0-a)(x_0+a)(x_0-z)\mathrm{d}x_0 + \frac{1}{z-x_-}\int_{x_-}^{z} v_{min}\mathrm{d}x_0 + \frac{1}{x_+ - z}\int_{z}^{x_+} v_{max}\mathrm{d}x_0 + \frac{1}{a-x_+}\int_{x_+}^{a} -k(x_0-a)(x_0+a)(x_0-z)\mathrm{d}x_0 \tag{4}$$

The noise term does not intervene as we assumed a Gaussian white noise. We observe a discontinuity in z, due to the presence of an unstable fixed point at that location. Trajectories determined by $x_0 = z$ will finish in either well under the influence of noise, and the mean of the noise being zero, the two scenarios are equally likely. Consequently, the mean drift for these trajectories is the average between v_{min} and v_{max}, with v_{min} (respectively v_{max}) is the maximum negative (respectively positive) drift rate achievable by the system. The graph of the max drift as a function of starting point is given in Figure 3.

From Figures 2 and 3 we can see that z and a impact the mean drift (see also Figure 4). It becomes clear that the parameter z has a larger effect on the mean drift than the parameter a. That is explained by the fact that z determines which proportion of the trajectories is attracted to the positive boundary for a given distribution of starting points. In contrast, a determines the scale of the drift.

This model is similar to the Double-Well Model (DWM), which emerges from attractor network models (Prat-Ortega et al., 2021; Roxin and Ledberg, 2008). The potential profile of the DWM indeed takes the form:

$$V(x) = -\mu x - \alpha x^2 + x^4. \tag{5}$$

Comparing this equation to Equation (3), we observe a term in x^3 that is absent from the DWM, because of the reflection symmetry assumption made in the DWM (Strogatz, 2015; Roxin and Ledberg, 2008). However, when $z = 0$ and $\mu = 0$, we observe the equivalence of the systems by having:

$$k = 4$$
$$a^2 = \alpha/2$$

This equivalence is coherent with the interpretation of z and μ as the impact of the stimulus on the decision, and shows that in the absence of a stimulus, the two models follow the same behavior. Because the nl-DDM is not assuming reflection symmetry, the presence of a stimulus impacts the trajectories generated by the two models in different ways.

Behavioral results

For decision-making analysis, it is helpful to obtain each participant's response times and decision accuracy, particularly for decision model fitting.

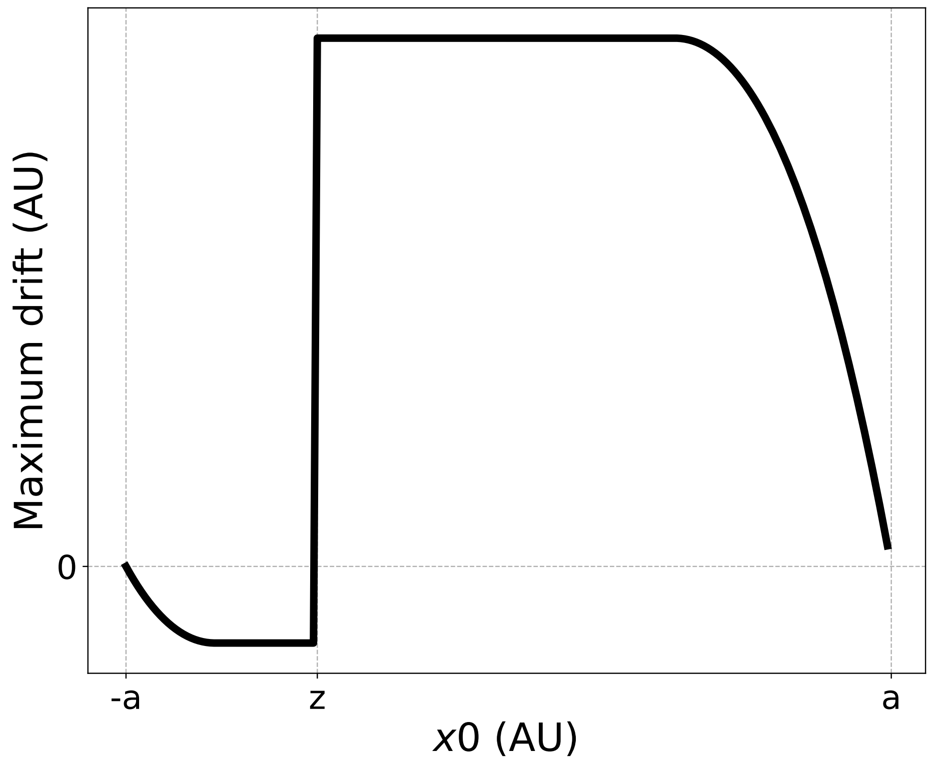

Figure 3. Maximum drift as a function of starting point

We used two datasets in this paper. The section Data collection and processing describes these datasets in detail. They both consist of classification tasks performed by human participants. One of them is a dataset collected by Wagenmakers et al. (2008), in which participants had to assess whether a word presented on screen existed or not. The second one is a dataset not presented before, in which participants were shown visual stimuli on screen and had to classify them according to their type (either "face" or "number").

To ensure the correctness of both datasets in terms of behavioral measurements, we describe here the validation conducted on our dataset. Analyses of the Wagenmakers' dataset are available in Wagenmakers et al. (2008) and are not discussed further here.

First, we ruled out methodological artifacts, as we aimed at providing equiprobable stimuli for each participant. On average, participants were shown $49.82 \pm 2.42\%$ of "number+sound" stimuli, showing the quasi equiprobability of each stimulus. We then tested whether the experiment we designed led to similar responses across all participants by performing mixed-model ANOVAs on their response times and response accuracy for both stimulus-response mapping (between-subject factor) and stimuli (within-subject factor). Across all participants and stimulus types, the mean response time is 535 ± 61 ms (mean $\pm$ standard deviation, $N = 25$), with an accuracy of $98.59 \pm 0.95\%$. For the "face" stimulus, participants responded after 539 ± 56 ms with an average accuracy of $98.51 \pm 1.17\%$. Participants responded to the "number + sound" stimulus after 531 ± 69 ms on average with an accuracy of $98.68 \pm 0.94\%$. The difference in performance between the types of stimuli is not significant in terms of accuracy (Table 1) nor in terms of response times (Table 2).

In the "face is left button" stimulus-response mapping, where participants were instructed to click left upon face stimulus presentation and right when they were presented with a number+sound stimulus, participants responded on average within 531 ± 74 ms with an accuracy of

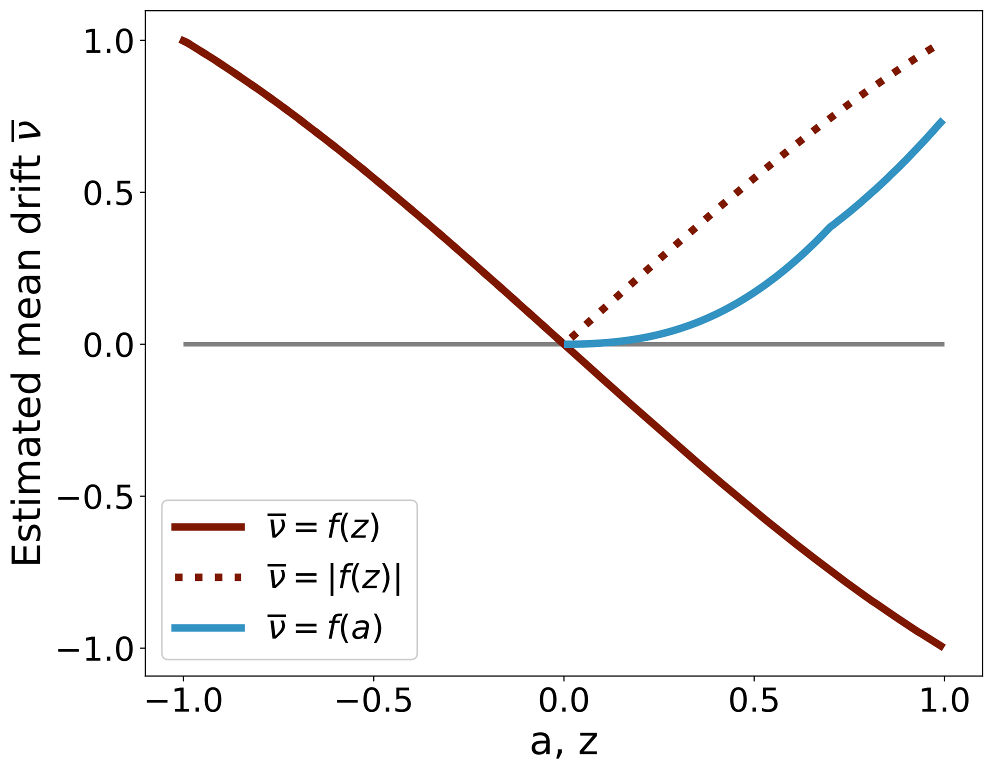

Figure 4. Effect of z and a on the mean drift, estimated as the mean of the maximum drift for each trajectory determined by its starting point. We formulated the nl-DDM drift under the form $dx = -k(x-a)(x+a)(x-az)$, having $-1 < z < 1$, without loss of generality. The mean drift is defined as in Equation (4), which depends both on z and a. The darker line represents the variation of the mean drift thus defined as a function of z, while the pale blue curve is the variation of the mean drift as a function of a. Since a is strictly positive, we also represented the absolute value of the mean drift (dotted line). That allows for comparing the magnitude difference of the mean drift rate when z or a vary. We see that varying z changes the mean drift rate more strongly than similar variations of a at a given value of z.

$98.48 \pm 1.12\%$ ($N = 15$), whereas participants who underwent the "face is right button" stimulus-response mapping, participants ($N = 10$) responded within 541 ± 30 ms and an accuracy of $98.77 \pm 0.60\%$. The effect of the stimulus-response mapping on accuracy and response time was not significant (Tables 3 and 4). We do note however a marginal interaction effect between stimulus-response mapping and stimulus type on the accuracy of participants ($p = 0.052$, Table 1).

These results show the uniformity of participant responses across mappings and stimuli. All participants, mappings and stimuli were considered together in the subsequent analyses.

Table 1. Within Subjects Effects on Accuracy

Cases	Sum of Squares	df	Mean Square	F	p
Stimulus	1.249×10^{-5}	1	1.249×10^{-5}	0.299	0.590
Stimulus * S-R mapping	1.758×10^{-4}	1	1.758×10^{-4}	4.202	0.052
Residuals	9.623×10^{-4}	23	4.184×10^{-5}		

Table 2. Within Subjects Effects on Response Times

Cases	Sum of Squares	df	Mean Square	F	p
Stimulus	1201.903	1	1201.903	2.446	0.132
Stimulus * S-R mapping	370.446	1	370.446	0.754	0.394
Residuals	11303.230	23	491.445		

Table 3. Between Subjects Effects on Accuracy

Cases	Sum of Squares	df	Mean Square	F	p
S-R mapping	8.608×10^{-5}	1	8.608×10^{-5}	0.447	0.510
Residuals	0.004	23	1.926×10^{-4}		

Fitting on data

The fitting of parameters was performed using the PyDDM (Shinn et al., 2020b) Python toolbox for both the nl-DDM and the DDM, minimizing the negative log-likelihood function. As participants in our experiment were shown two types of stimuli, we fitted a model per participant for each model type, resulting in 25 DDM and 25 nl-DDM fitted. In addition, 17×2 models of each type were computed for the Wagenmakers dataset (17 participants $\times$ 2 conditions = 34 models). Since the two datasets did not use the same number of parameters for each model, we performed a pairwise comparison of loss values over the models. To remove the possible effect of outliers, for which fitting would have failed, we removed the models for which the loss values were above the mean loss + standard deviation over all models. This resulted in the rejection of 11 participants $\times$ conditions (7/17 rejected in the Wagenmakers accuracy condition (41%), 2/17 in the Wagenmakers speed condition (12%), and 2/25 in our dataset (8%)), so 81% of all fitted models were kept.

Comparison of loss values

The first metric we used to compare the models is the loss value after fitting. Fitting is done by minimizing the negative log-likelihood, which gives information on how close the curve of theoretical response times is to empirical response times histograms. For a measure that takes into consideration the number of parameters and samples, we also computed the Bayesian Information Criterion (BIC). All the test results on fitting performance are summarized in Table 5.

Table 4. Between Subjects Effects on Response Times

Cases	Sum of Squares	df	Mean Square	F	p
S-R mapping	1438.081	1	1438.081	0.179	0.676
Residuals	184867.754	23	8037.728		

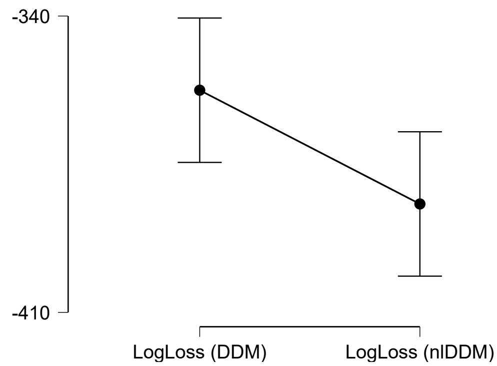

Figure 5. Comparison of fitting loss values between the DDM and the nl-DDM. Error bars show the 95% confidence interval on the mean values.

The comparison of loss values between model types (Figure 5) shows that the nl-DDM fits data significantly better than the DDM for the same number of parameters. Indeed, the loss values are significantly smaller in the nl-DDM compared to the DDM, with a moderate effect size (right-tailed paired t−test, $t(47) = 2.18, t = 2.241, p = 0.015, d = 0.324, N = 48$).

We computed the Bayesian Information Criterion (BIC) for each model to establish a comparison of model performance that takes into account the sample size and number of parameters necessary for each model. This is indeed necessary when comparing performance across datasets, since the number of conditions, and hence of parameters needed, is different. We observed that the nl-DDM fitted response time data significantly better than the DDM even when accounting for the number of parameters (Figure 6, $t(47) = 2.18, t = 2.207, p = 0.016, d = 0.319$).

Speed-accuracy trade-off

We computed the behavior prediction of each model type to ensure that the results are consistent with empirical observations. For that, we used a metric described by Roitman and Shadlen (2002), whereby the loss is computed as the sum of the mean squared error on mean response time and the mean squared error on the predicted accuracy over all conditions. We observe that there is no

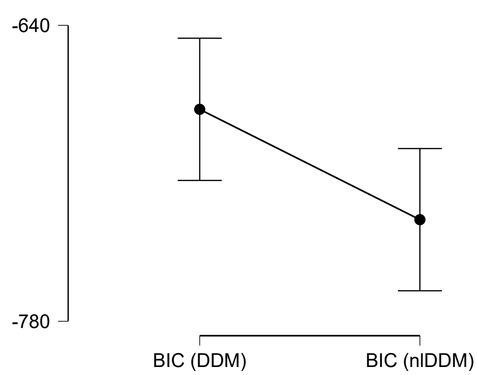

Figure 6. Comparison of the Bayesian Information Criterion (BIC) between the DDM and the nl-DDM. Error bars show the 95% confidence interval on the mean values.

Table 5. Left-tailed paired samples t-test on the quality of the fit between the nl-DDM and the DDM over all fitted models ($N = 52, df = 51, t(51) = 1.675285$). Significant p-values are marked with *.

Measure 1		Measure 2	t	df	p	Cohen's d
LogLoss (DDM)	-	LogLoss (nlDDM)	2.241	47	0.015	0.324
BIC (DDM)	-	BIC (nlDDM)	2.207	47	0.016	0.319
Performance Loss (DDM)	-	Performance Loss (nlDDM)	−0.357	47	0.639	−0.052

significant difference in terms of behavioral prediction capacity between the DDM and the nl-DDM (see Performance Loss, Table 5, and Figure 7).

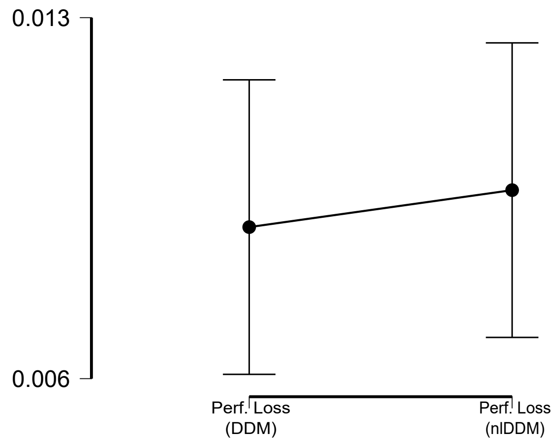

Figure 7. Comparison of loss computed on behavioral performance between the DDM and the nl-DDM. Error bars show the 95% confidence interval on the mean values.

Comparison of parameters

Although we fit the parameters separately for each stimulus type, we merge all the results to build relations between the parameters of the DDM and the parameters of the nl-DDM. Of the resulting fitted models, we rejected the participants that were rejected in our previous analysis (participants 6 and 11). In addition, participant 22 was rejected due to a fitted boundary outside of the other models' range. Hence, 44 models were taken into account.

Given the mathematical formalism described above, we expect to find a negative correlation between the decision boundaries of the two models. Indeed, while an increase in the boundary in the DDM results in increased accuracy and response times, a similar increase in the nl-DDM results in decreased accuracy and response times. Our previous explanation of model parameters showed that both a and k in the nl-DDM impacted the decision boundaries. Therefore, the boundary of the DDM could be negatively correlated to either of these parameters. We also expect the parameter z of the nl-DDM to be correlated negatively with the drift in the DDM. The reason for this can be derived from Figures 2 and 3: if we shift z closer to 0, the negative and positive plateaus of Figure 3 will tend to be at the same level in absolute value. Averaging them, it means that the mean maximum drift will decrease towards zero as z increases closer to the middle of the two boundaries $\pm a$. In other words, increasing z will decrease the drift, hence the negative correlation.

The correlation matrix of the nl-DDM and DDM parameters across all models is given in Figure 8. Note that we only took into account our dataset for parameter comparison, and while fitting the models to this set, we assumed that the starting points followed a uniform distribution spanning the entire decision interval $[-a, a]$ for the nl-DDM, while we took a shorter interval for the DDM to avoid border effects. This difference in assumptions for the starting point distribution was compensated by fitting a non-decision time in the nl-DDM and taking the same value when fitting the

DDM.

We first empirically show the relation between the parameters within the nl-DDM. We observe a strong negative correlation between a the boundary and k the time constant. This corresponds to their similar effect on the attractiveness of the correct response. Increasing either will make the decision more attractive, so to keep the same attractiveness of the correct response, if one increases, the other should decrease. In our data, since the accuracy is similar across all participants and conditions, and the noise term is kept constant, these two terms are strongly correlated. Note that the effect of each parameter is still different, as shown on Figure 2B and C. While increasing k deepens both wells, increasing a will not only deepen the wells but also pull them apart. Effectively, the relation between a and k is not linear, as seen on Figure 9.

We also observed the known relations within DDM parameters in the correlation matrix (Figure 8). The starting point distribution, parameterized by its center x_0 and half-width s_z, correlates to the boundary b and the drift v. It could be expected, as the same response time for correct responses will necessitate faster integration of evidence (that is, accumulation to a bigger drift) if the starting point is further (smaller). Similarly, as the boundaries get stretched, the starting point distribution also needs to become wider for similar response time shapes, hence the positive correlation between b and s_z.

Concerning the cross-model comparison, we observe that k is negatively correlated to the DDM boundary b and positively to x_0. The link with the decision boundary is expected, as k in the nl-DDM regulates the depth of the decision wells, that is, the time necessary to reach each decision. More specifically, increasing k makes the wells more attractive and hence results in fast decisions. Conversely, increasing the DDM boundary will result in longer response times as more time will be needed for the decision variable to reach the boundary, given the linear drift.

We also observe a significant negative correlation between the parameter z of the nl-DDM and the drift parameter of the DDM v. This relationship was also expected, as increasing the drift v in the DDM results in faster correct decisions. Mirroring this effect, z regulates the relative attractiveness of each decision well. As z becomes more negative, the correct decision (corresponding to decision boundary $+a$) becomes more attractive, and hence correct decisions are made faster.

Last, the position of the stable fixed-points, parameterized by a, is positively correlated to the DDM drift v. We have described earlier that v was also related to z. It therefore seems that the drift is related to these two parameters. A formal analysis of the nl-DDM (see Methods) provides the following explanation: all three parameters a, k, and z, impact the speed at which the decision boundary is reached. While z controls the relative attractiveness of the boundaries, both a and k modulate their absolute attractiveness. An increase in either of them will result in faster response times for both correct and incorrect responses. In the DDM, two parameters impact the speed of responses: v and the boundary b. Similarly, k should be more related to the decision boundary given the previous correlation, and a to the drift. Given the correlation between a and v, we explain the observed correlation between a and x_0 due to the correlation between v and x_0.

We observed more precisely the intertwining of parameters through Principal Component Analysis, keeping only the components corresponding to eigenvalues of the correlation matrix greater than 1. For clarity, we displayed for each principal component only the parameters with a loading in absolute value greater than 0.4, that is, parameters sharing at least 16% of variance with the component. We thus obtained 3 principal components, accounting for $82,2\%$ of the total variance, summarized in Table 6. The first principal component (PC1) accounts for 34.4% of the variance. We interpret it as the decision attractiveness, given that a and k share more than 90% of their variance with this component. The second component (PC2) is loaded mainly by v and z, which relate to the decision speed. The fact that x_0 shares its variance between both PC1 and PC2 is consistent with the fact that decision bias x_0 contributes both to the decision attractiveness and to its speed. Finally, the third component (PC3) is strongly loaded ($> 80\%$ of shared variance) by s_z and b, suggesting that it reflects the amount of information that needs to be integrated in order to make a decision, or the decision caution (Ratcliff and McKoon, 2008). Indeed, the greater the uncertainty

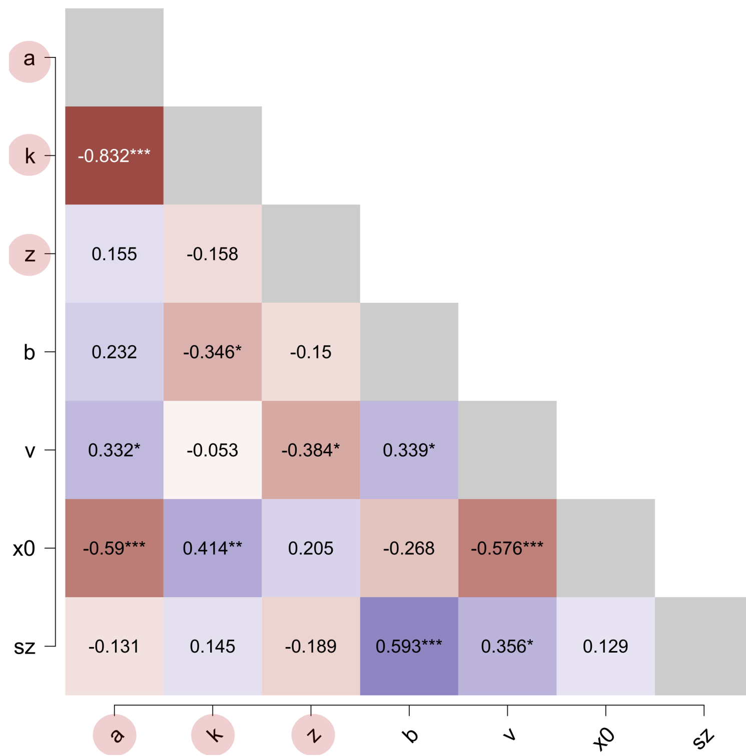

Figure 8. Correlation matrix of all parameters. a, k and z are parameters of the nl-DDM (marked in a circular patch), while b, v, x_0 and s_z are parameters of the DDM. Pearson correlation coefficients were computed over $N = 44$ observations.
$\star : p < 0.05, \star\star : p < 0.01, \star\star\star : p < 0.001$.

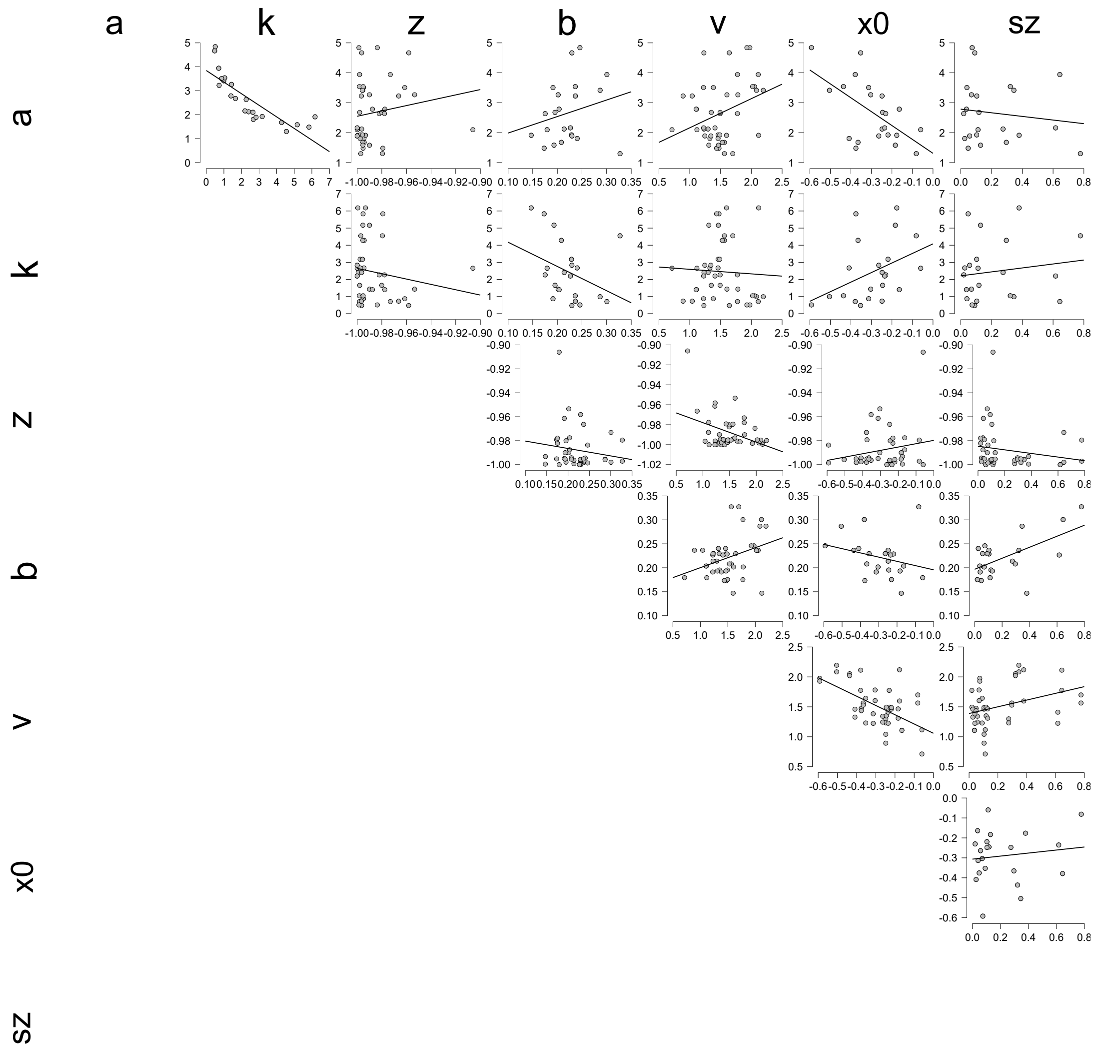

Figure 9. Correlation plots of DDM and nl-DDM parameters. The correlations were computed as Pearson's correlation coefficient.

s_z associated with decision bias x_0, the greater the amount of information needed to reach a decision. In this case, PC3 would reflect this decision caution: the greater the uncertainty s_z associated with decision bias x_0, the more cautious one must be, and in turn greater amounts of evidence one needs to process.

Table 6. Component Loadings

	PC1	PC2	PC3	Uniqueness
a	−0.942			0.104
k	0.912			0.141
x_0	0.649	0.626		0.181
v		−0.795		0.227
z		0.780		0.320
s_z			0.907	0.111
b			0.847	0.166
Variance explained	0.344	0.242	0.236	

Discussion

We have presented in this paper a non-linear model of decision-making. This model is a form of generalized drift-diffusion model (Shinn et al., 2020b), and provides a framework in which individual trajectories of the decision variable can have different shapes under the same global parameters (Figure 2D). Even without considering the single-trial fitting capability, we have shown that this model predicts behavioral data equally well as the DDM, with a slight but significant improvement in the goodness of fit. From the formalism we have described, it becomes clear that inter-trial variability in drift emerges from the dynamics of the system proposed, offering the possibility for further single-trial analyses and modelling.

The interpretation of the nl-DDM parameters may seem counter-intuitive at first, in particular when considering that decisions are made faster when the boundaries are further apart. Indeed, we observe the opposite effect in the DDM. However, our correlation analysis provided insight into bridging the meaning of nl-DDM parameters to DDM parameters. The difference is that in the DDM, the gradient of the drift is constant, whereas it varies in decision space with the nl-DDM. By pulling the boundaries further apart, we effectively reduce the impact of one attractor on the other, making each of them more attractive. Therefore, a decision can be reached faster, at the price of accuracy. Similarly, increasing the drift in the DDM is equivalent to shifting the unstable fixed point towards the negative boundary in the nl-DDM, as they both result in fast correct responses. However, it must be noted that these parameters are not entirely equivalent as we did not find a perfect mapping between them, meaning that the nl-DDM is conceptually different from the DDM.

We have shown that, while similar to the DWM (Prat-Ortega et al., 2021) derived from attractor models (Roxin and Ledberg, 2008), the nl-DDM is equivalent to it only in the absence of input. A question that remains open is that of the mechanism underlying this equation. From the reduction computed in the paper by Roxin and Ledberg (2008), it would seem that a network of three populations could produce the dynamics we have described. However, the main assumption of the reduction was that the network was invariant through reflection. We argue that the mechanisms described by the nl-DDM are in fact similar to these of the DWM, but offer a broader range of application beyond the case of symmetrical models.

A question that arises from our analyses is the different assumptions made on the starting point. We took in both cases a uniform distribution, but while fitting our dataset, we assumed that this distribution spanned the whole decision space for the nl-DDM, while it was an interval $[x_0 - s_z, x_0 + s_z] \subseteq [-b, b]$ for the DDM, with $\pm b$ the decision boundaries of the DDM. By doing so, we wanted to show that with fewer degrees of freedom, our model could fit behavioral data better than the DDM. The DDM assumes a global bias towards either boundary transcribed in the position of the starting point distribution within the decision interval. Variability in the starting point enables faster error responses (Ratcliff and Rouder, 1998). In our model, that could also be achieved by including starting point variability similarly to the DDM, that is by defining an interval $[x_0 - s_z, x_0 + s_z] \subseteq [-a, a]$, with $\pm a$ the decision boundaries of the nl-DDM. It is the strategy that we have implemented while fitting the Wagenmakers' dataset. However, this would have been an issue when comparing the two models as the starting point distributions fitted by the DDM and the nl-DDM did not necessarily match. The most striking consequence of such a mismatch is the difference between the non-decision times of the two models. Indeed, the non-decision time of the nl-DDM would be close to the minimum response time displayed by a participant. At the same time, it would be smaller in the DDM, as the displacement of the decision variable from the bias to either decision boundary in the absence of noise is not instantaneous. To minimize this effect, we have also fitted the variability of the non-decision time for the Wagenmakers' dataset, although introducing such variability in one model and not the other made the two models less comparable. One could also argue that we could have chosen a uniform starting point distribution for both models. The problem with this solution is that due to the linearity of the drift in the DDM, the resulting response time distributions would have had sharp edges, which are a direct consequence of starting close to the decision boundary. We thus fitted x_0 and s_z for the DDM and not for the nl-DDM when comparing the two models. That is, we fitted the starting point distribution for the DDM, but not for the nl-DDM.

We argue that drift and starting point variability are not independent, which is transcribed in the system's dynamics we created. EEG research has shown a matching between pre-stimulus activity and confidence ratings in human participants (Wöstmann et al., 2019; Samaha et al., 2017). Pre-stimulus states are modeled by the starting point and its variability, and in the DDM the drift relates to the quality of the stimulus being integrated (Gold and Shadlen, 2007), with more ambiguous stimuli corresponding to lower drift rates. Translated to the single-trial level, drift variability relates to the variation of how well the brain perceives and processes the stimulus (Ratcliff and McKoon, 2008). In our model, the starting point directly impacts the evidence accumulation, allowing for a more uniform theory of decision-making than the DDM that includes explicit co-dependency of certain parameters. Some general forms of the DDM include a variance of the drift, which we have never considered here. In the current nl-DDM, we have not implemented such a possibility, as we assumed that the inter-trial variability of the drift simply emerged from the variability of the starting point. In neurophysiological terms, we assumed that the pre-stimulus arousal and expectations on the stimulus led to differences in the rate of evidence accumulation. This is supported by past observations, according to which pre-stimulus brain activation impact response times (Petro et al., 2019; Chen et al., 2020). Pre-stimulus brain activity also modifies perceptual (van Dijk et al., 2008) and pain (Taesler and Rose, 2016) thresholds. Therefore, depending on the pre-stimulus activity, decisions can be made, even in the absence of actual evidence (Barik et al., 2019; Wöstmann et al., 2019), or under ambiguous evidence (Rassi et al., 2019; Railo et al., 2021). Along the same lines, Kloosterman et al. (2019) have shown that biases were implemented through local changes in accumulation rate, which supports the intertwining of accumulation rate and pre-stimulus states. However, (Benwell et al., 2021; Samaha et al., 2017; Iemi et al., 2017; Lange et al., 2013) argue that pre-stimulus brain states should only affect the decision criterion, not how well participants could perceive the stimuli. Translating the signal-detection theory to the evidence-accumulation scheme (Ratcliff and Rouder, 2000), it means that pre-stimulus states should only be changing the decision boundary, or equivalently, changing the starting point, and not the drift rate. For example, Samaha et al. (2017) found that pre-stimulus alpha power did not impact the accuracy of visual evidence accumulated, but only the confidence in the decision. Wöstmann et al. (2019) found similar results with the auditory modality. Although these observations seem to contradict our assumption that the starting point should impact the evidence-accumulation process, both phenomena could co-exist, as indeed more extreme starting points are more attracted to the closer attractor. This results in fast and confident observations, although little evidence has been accumulated (we would be located at a plateau in our model), that is, even if the stimulus was not well perceived.

The dynamics that we propose here are not the sole product of mathematical formalism and constraints, but have deep roots in empirical observations made in neurophysiological studies. More specifically, three phases can be identified in the decision trajectories: an initial inertia stage, a roughly linear evidence accumulation stage, and a plateau stage. The initial inertia relates directly to the brain activation needed to integrate sensory evidence. Petro et al. (2019) and Chen et al. (2020) have shown in human EEG studies that depending on the brain activity prior to stimulus presentation changed the speed of responses. More specifically, they showed that the more pre-activated the required sensory area, the faster the decision. The nl-DDM mimics this behavior at the single-trial level: for trials starting close to the unstable fixed-point (that is, further from the correct decision well), the trajectories start with a plateau-like stage, whereby little evidence is accumulated because the brain would need to process the stimulus more intensively in order to extract information from it, before integrating evidence faster. This initial inertia is circumvented by shifting the starting point closer to the decision well, resulting in faster and more accurate responses. The initial inertia in the DDM is referred to as the non-decision time and encompasses both sensory processing and motor planning and execution processes. The nl-DDM assumes therefore that part of these processes participate in the decision process, which goes beyond the conceptualization of decision-making as a sequential process of sensation, perception and motion.

A recent review from Evans and Wagenmakers (2020) shows the limitations of existing evidence-accumulation models. We try to address several of them with the present model, including the possibility for analyses beyond the global description of response times and the formulation of initial and final dynamic changes during the decision process. In particular, our formal description has shown that different shapes of decision trajectories can co-exist within the same framework, not solely because of noise, but because of meaningful variability. We expect this model to be further analyzed and used to gain insight into the single-trial dynamics of decisions.

The present work did not include trials where the response is missing, which sometimes occur when participants need more time to decide. However, it could easily be implemented with a timeout. Typically, with a stronger constraint on response time in the experimental paradigm, it is likely that participants do not have time to give a response. One could imagine that the decision has not settled in either well at timeout, and this parameter could be taken into account in future works with different experimental paradigms.

The current study only addressed the case where the input was presented at the beginning of the trial and affected the decision in a constant fashion. We could also imagine more dynamic cases, where the input is processed over a finite amount of time and participants accumulate evidence solely during stimulus presentation, as has been done in past DDM analyses (Huk and Shadlen, 2005; Shinn et al., 2020a). In non-stationary contexts, the input can be considered as a variation of z in time. By shifting that parameter to either boundary, we make more trajectories attracted to the opposite boundary, hence increasing the likelihood of correct answers. In addition, it can be inferred from our formal analysis that changing z means changing the drift rate. This change in input could also explain error-correcting behaviors (Rabbitt, 1966) and spontaneous changes of mind (Pleskac and Busemeyer, 2010). When the stimulus ends, the DDM is modified so that the drift is null, i.e. evidence is no longer accumulated. Therefore, changes of mind are the result of noise in the system. Conversely, stimulus termination could be modelled through shifting z in the nl-DDM, which effectively modifies the drift-rate of the current decision, in a way that the decision variable could toggle towards the opposite boundary upon stimulus disappearance. Conceptually, the drift in the nl-DDM not only relates to the accumulation of evidence but also encompasses decision processes related to the post-processing of evidence.

Ideas and Speculations

As mentioned above, the presence of two attractors offers an interesting perspective for when participants are asked to alter their perceptual responses during the trials. Indeed, not all types of evidence can extract the decision variable from the region of attraction of a fixed point with enough strength for a participant to change their mind. This has already been conceptualized in studies of perception (Hafemeister et al., 2010), whereby different representations can emerge, and participants can switch, consciously or not, from one representation to the other (see Rolls and Treves, 1999, Chapters 4-6). The size of evidence, modeled by the position of the unstable fixed point z, can be estimated quantitatively from experimental parameters fixed by the experimenter, such as sound level and luminosity.

In the past, attempts at single-trial fitting have been debated (Latimer et al., 2015; Zylberberg and Shadlen, 2016; Latimer et al., 2017). Maybe this model could explain the observations made by Latimer et al. (2015). Fitting of neural data (spiking data, similar to analyses in Latimer et al. (2015) and Gold and Shadlen (2001)) could give insight into the goodness of fit of this model with regards to the choice of the starting point and drift variance: the DDM assumes that they are two separate phenomena, and it is hard to extract any baseline excitation from behavioral data only. This could give us an indication on whether there is a link between the starting point and inter-trial variability of the drift, and whether the nl-DDM captures this interaction.

Extending this model to multiple-choice situations is another interesting ground of research. The DDM is not easily applicable in such situations, whereas models such as the Linear Ballistic Accumulator model (Brown and Heathcote, 2008) are. We argue that the current model would require structural changes in its formulation, without however changing its essence, for such situations to be implemented. Indeed, the trajectory of the decision variable is here modelled in a one-dimensional space, where the possible alternatives are represented as attractors. Its multiple-choice variant would require several other attractors. In $1D$ space however, adding more stable fixed-points will result in two issues. First and foremost, travelling from one alternative to another may require passing through other decision wells, which seems incoherent with behavior. It seems counter-intuitive that a person has to make a decision before travelling to another decision state. Second, adding more stable fixed-points requires the implementation of as many unstable fixed-points between two stable fixed-points (see nl-DDM formalism), which would mean that the number of parameters to fit increases by 2 when adding one choice. A simpler solution would be to switch to a $2D$ space, so there could still be a central unstable fixed-point, and the position of each stable fixed-point in $2D$ space would be determined by the subjective preference of each alternative.

Methods and Materials

Drift-Diffusion Model

The Drift-diffusion model (Ratcliff, 1978) is characterized by a linear accumulation disturbed by additive noise. Formally, this can be written as the following Langevin equation (Equation (6)):

$$dx = v dt + N(t), \tag{6}$$

where x represents the decision variable, an abstract quantity representing the state of the decision, dx its infinitesimal variation in time dt, and $N(t)$ is a Gaussian white noise, parameterized by its standard deviation σ. Figure 10 gives a representation of this model.

Evidence is accumulated following Equation (6) until a decision boundary $A > 0$ or $-A$ is reached. Typically, the positive boundary corresponds to correct decisions and the negative one to incorrect responses.

Finally, the starting point of accumulation is called the bias and is defined as a single point within the two boundaries. In general forms of this model, it is also possible to consider that the starting point is drawn from a uniform distribution centered around the bias x_0 and of width $2s_z$,

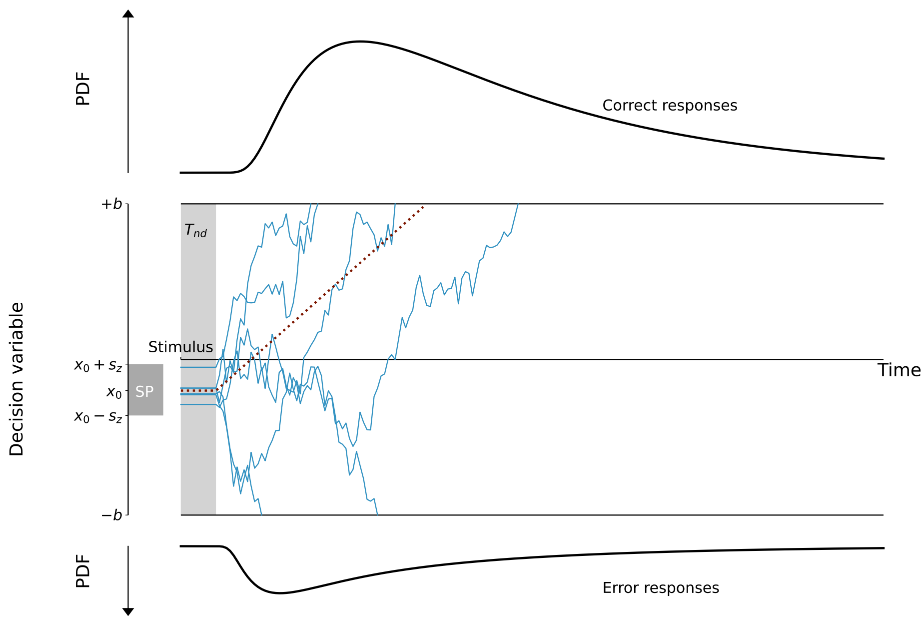

Figure 10. Description of the Drift-Diffusion Model (DDM). The decision state is represented through a decision variable that travels from a starting point that can be drawn for example from a uniform distribution, centered around x_0 and of width $2s_z$. The decision state is represented through a decision variable x traveling from a starting point (for example, drawn from a uniform distribution, centered around x_0 and of width $2s_z$. It is represented as "SP" on the figure) to a boundary ("Correct boundary" or "Incorrect boundary") under the influence of a constant drift (dotted line). The trajectory is also impacted by white noise so that real trajectories are similar to the thin blue lines. From the stimulus onset, the decision process is delayed by a certain non-decision time (T_{nd}). Over an ensemble of decisions, response time distributions of correct and error responses can be estimated, as displayed here.

such that $[x_0 - s_z, x_0 + s_z] \subseteq]-A, A[$ (Laming, 1968), or from other parametric distributions (Ratcliff and Rouder, 1998). We will consider uniformly distributed starting points in our fitting to provide a fair comparison of the two models without loss of generality.

The boundary separation represents the speed-accuracy trade-off. Indeed, if this separation is bigger, decisions are less impacted by noise and hence more accurate, but at the same time, they will take longer to reach from a given starting point. In contrast, the drift mainly impacts the speed of response, as a higher drift will lead to faster correct responses and longer incorrect responses.

Fitting is typically done globally over response times. In fact, the trajectories defined by the equation cross the decision boundaries, forming a response time distribution usually compared to an exponentially modified Gaussian. In order to obtain a close fit, it is necessary to define a non-decision time (noted T_{nd}), which corresponds to the time necessary for sensory processing of the stimulus, motor planning and execution, independently of the decision process.

Data collection and processing

In order to test the quality of the fitting of the proposed model, we use response times from a classification task performed by humans described thereafter. The paradigm was initially implemented to assess the relation between response times and emotion valence of visual stimuli.

Classification task with different sensory modalities

We first tested the quality of the nl-DDM by fitting it to data we collected. 25 (11 female, 14 male) healthy right-handed participants aged 27.72 ± 8.96 (mean $\pm$ standard deviation) with normal or corrected-to-normal vision and hearing took part in a perceptual classification task experiment. EEG brain activity was also recorded (not reported here). The experiment was performed under the local ethics committee approval of the Comité d'Ethique de la Recherche Paris-Saclay (CER-Paris-Saclay, invoice notice nb. 102). An interview preceded the experiment to check with the participants for non-inclusion criteria (existing neurological and psychiatric disorders, uncorrected visual and hearing deficiencies). Participants were presented at each trial with images of faces or images of numbers, and had to respond with mouse clicks to report what stimulus they perceived. A sound accompanied images of numbers to suppress any ambiguity. Participants were instructed to respond using their right hand. To control for possible differences in motor response speeds between the two fingers, one group of participants ($N = 15$) was instructed to report faces with a left click and numbers with a right click ("face is left button" stimulus-response mapping), while the other ($N = 10$) was given the opposite instruction ("face is right button" stimulus-response mapping). Responses were constrained to two seconds after stimulus onset. No feedback on the performance was given to participants. At each trial, each stimulus had a 50% chance of occurring.

Each participant performed 480 classification trials, split into 8 blocks of 60 trials each. Between each block, participants were offered a break of free duration. Each trial followed the sequence described in Figure 11. First, a central red cross appeared on the screen, indicating a pause period. After 1.5 second, the cross became white as a signal for trial start. The white cross stayed for 1.5 second, after which a video clip of visual noise appeared: 9 frames of noise of 100 ms each were displayed. After the noise clip, a last frame of random visual noise was presented, and the stimulus appeared on top of it. The last frame stayed intact until the end of the trial, and the stimulus was displayed over it for $200ms$. The trial was terminated upon participant response or timed out after 2 seconds. A trial lasted for about 5 seconds, resulting in blocks of about 5 minutes each.

We used face sketches as used in Yang et al. (2020), which were generated from the Radboud Face Dataset (Langner et al., 2010). Number stimuli were generated at the beginning of the session for each participant, under the constraint that they were 3-digit integers. In total, 10 different face stimuli and 10 different number stimuli were used for each participant.

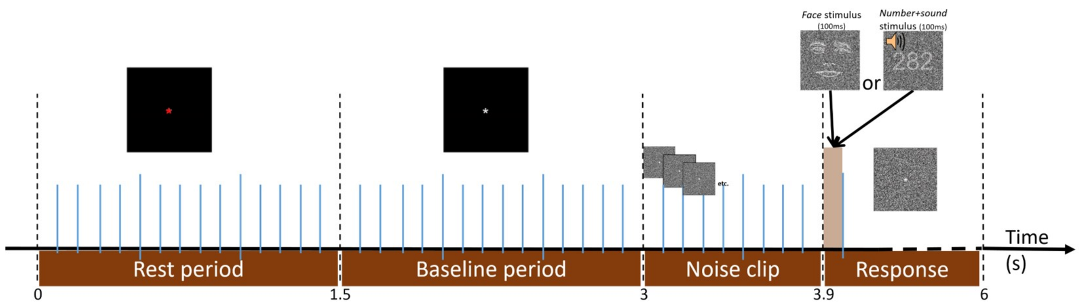

Figure 11. Timeline of a single trial. Each trial is preceded by a rest period, followed by a baseline period (necessary for EEG processing, not reported here), each lasting 1.5 seconds. A noise clip consisting of 9 random-dot frames of 100 ms each indicates the arrival of the stimulus in a non-stimulus-specific fashion. The stimulus then appears on a noisy visual background for 100 ms. The same noisy background frame then lasts until the participant's response and times out after 2 seconds otherwise.

Pre-existing dataset from Wagenmakers et al. (2008)

To discard the possibility of better performances emerging from the fitting algorithm or data acquisition, we also lead our analyses on a pre-existing dataset taken from Wagenmakers et al. (2008). 17 human participants performed a classification task, as they were randomly presented with real or invented words. The invented words were generated from real words by changing a vowel, and the real words were labeled in three categories depending on their frequency (frequent, rare, or very rare). In total, stimuli were split into 4 categories of interest. Each participant performed 20 blocks of 96 trials each, with as many invented words as real words in each block. Participants were given the additional instruction to define the speed-accuracy trade-off in each block: they alternated between blocks where speed was emphasized and blocks where accuracy was more important. Responses were limited to 3 seconds, and trials with response times below 180 ms were discarded to avoid anticipatory responses. More details can be found in Wagenmakers et al. (2008), and the dataset can be accessed from here.

Behavioral analyses

We are interested in comparing model parameters between the DDM and the nl-DDM. It is important to check whether participants' performance across stimulus-response mappings and stimuli is coherent in terms of response times and accuracy. Indeed, the experimental paradigm we defined entails two types of stimuli and two motor commands for the choices. In addition, we have created two experimental groups, which were instructed to respond with opposite motor commands. First, we computed the percentage of stimuli in each class to verify that the stimuli were globally equiprobable for each participant. Since we designed the experiment to display each stimulus with the same probability at each trial, we expect this number to be close to 50%. Otherwise, participants could opt for a strategy that prioritizes one response against the other. Then, we performed two mixed-model ANOVAs, testing response times and accuracy respectively. The stimulus-response mapping was considered a between-subject factor and the stimulus type a within-subject factor.

Data fitting

The classical way of fitting evidence-accumulation models is by fitting one drift for each stimulus category separately. In that case, the positive and negative boundaries still correspond to correct and incorrect responses respectively, and the starting points are taken from the same distribution regardless of the stimulus. Consequently, one pair of boundaries $\pm B$, the middle of the starting point distribution x_0 and its half-width s_z, and two drifts v_0 and v_1 (corresponding respectively to "face" and "number+sound" trials) have to be fitted in the DDM. Similarly, one pair of stable fixed points (attractors, also corresponding to the decision boundaries) $\pm a$, one time scale k and two unstable fixed points z_0 and z_1 (repellers, that will tune the drift in the "face" and "number+sound" stimuli respectively) are needed for the nl-DDM. In both cases we fix the noise parameter to $\sigma = 0.3$. As explained by Ratcliff (1978), since the speed-accuracy trade-off is determined by the boundary separation, fitting two parameters among drift, boundary, and noise is constraining enough. Hence, 5 parameters have to be fitted per participant for the DDM, against 4 for the nl-DDM. In addition, fitting requires one non-decision time T_{nd} per stimulus type. The non-decision time and the starting point distribution are intertwined in the case of the DDM. Therefore, as trajectories starting closer to the boundary will reach it faster than trajectories starting further away, it is necessary to constrain either of these to provide comparison grounds between the nl-DDM and the DDM. For this reason, we first fit the nl-DDM and use the computed non-decision times as fixed parameters in the DDM.

We used the PyDDM toolbox (Shinn et al., 2020b, see: pyddm.readthedocs.io) for the fitting, minimizing the log-loss function and an implicit resolution. The explicit resolution is indeed impractical with the nl-DDM, which does not allow for explicit solutions when z is not centered. The log-likelihood is such that the more negative, the closer the modeled distribution of response times is to the empirical response time histogram.

Fitting Wagenmakers et al. (2008)

For this dataset, we reproduced the methods of Wagenmakers et al. (2008) by fitting the same parameters as in that paper for the DDM: the "accuracy" condition was first fitted globally for all participants, with a single boundary, starting point, non-decision time and noise term. In addition, the starting point and non-decision time variability were fit. One drift was computed per stimulus type, resulting in 4 drift terms: v_1, v_2, v_3, v_{NW}, corresponding respectively to frequent, rare, very rare and non-existent word stimuli. Hence, each model consisted of 10 parameters. Then, the same drifts, non-decision times (with its variability), and starting point variability were kept to fit the boundary and starting point in the "speed" condition.

We performed this analysis for each participant separately for more comparison grounds, while the original paper fitted all participants' response times together.

Given our formal analysis, we fitted a single a, k, noise, and starting point interval (centered around zero) parameters and one z per stimulus type (z_1, z_2, z_3, z_{NW}), resulting in 8 parameters, on the "accuracy" trials. Then, we fit again a, all other parameters fixed, on the "speed" condition. We did not fit the middle point of the starting point interval because z should fulfill this role, and not the non-decision time variability because the dynamics of the trajectories account for delayed onsets of the maximum drift rate depending on the starting point.

As previously, we used PyDDM (Shinn et al., 2020b) with log-loss minimization and implicit resolution.

Performance comparison

We used two metrics to compare the fitting performance of both models. First, we compared pairwise the loss scores, here the Negative Log-Likelihood, obtained after fitting. For our hypothesis to be validated, we expected the nl-DDM losses to be lower than these of the DDM. This metric assesses the shape of the predicted distribution of response times.

Since the fitting on both datasets was performed using a different number of parameters and samples, we also computed the Bayesian Information Criterion for each model, defined as:

$$BIC = \log(\text{sample size}) \times n_{\text{parameters}} + 2 \times (\text{Negative Log-Likelihood})\,.$$

That way, a penalty for more samples and parameters is considered.

Another interesting metric to compare decision models is their capacity to predict behavior. Indeed, one goal of the decision models we consider is to provide a theoretical description of individual speed-accuracy trade-offs. A good model predicts mean response times and error rates as close as possible to the empirical quantities. The metric we used to quantify the speed-accuracy trade-off is described in Roitman and Shadlen (2002), which associates a squared error to any deviation from the empirical mean response time and accuracy rate, summed over all conditions. Mathematically, this translates into a loss of the form:

$$L = \sum_{\text{conditions}} (\overline{\text{RT}} - \widehat{\overline{\text{RT}}})^2 + (\text{accuracy} - \widehat{\text{accuracy}})^2 \tag{7}$$

Note that the accuracy is computed as the ratio of correct responses over all responses, and lies within $[0 : 1]$, while the response times are provided in seconds. Since the mean response time is shorter than 1 s and of the order of a few hundred milliseconds, this metric scales speed and accuracy similarly.

Hence, we compare each loss pairwise, using three repeated-measure one-sided paired-sample t-tests. Indeed, we want to test whether the nl-DDM is better than the DDM with these three metrics, hence testing the hypothesis $\text{loss}_{\text{nl-DDM}} < \text{loss}_{\text{DDM}}$. Since we are comparing 3 losses, we set the threshold for significance to $\alpha = 0.017$, corresponding to the Bonferroni-corrected 5% threshold.

Comparison of parameters

For a better empirical understanding of the parameters of the nl-DDM, we computed the Pearson's correlation coefficients of the nl-DDM parameters over all conditions and participants, using only our dataset, that is, over $N = 50$ observations. This allows supporting the observations we have noted in the formalism part. Indeed, since fewer parameters were fitted in this case than for the Wagenmakers' dataset, the comparison becomes more straightforward. From the 25 participants, we obtained 50 fits per model type by duplicating for each stimulus type the boundaries and time constant terms, hence separating the stimulus types and obtaining 25×2 fits per model type. The models were filtered as previously based on the quality of the fit over all models. 6 models were thus rejected (12% of the total), limiting the comparison to 44 fits of each model type.

First, we computed the correlation matrix between all the parameters of both models. This allows for a first look into first-order interactions between model parameters, within and across model types. The correlation coefficients were computed using Pearson's ρ, defined as:

$$\rho_{x,y} = \frac{\text{cov}(x, y)}{\sigma_x \sigma_y} .$$

Next, to compare parameters of the DDM to parameters of the nl-DDM more quantitatively, we performed principal component analysis on the correlation matrix of DDM and nl-DDM fitted parameters. The goal is indeed to find how parameters relate to each other. This becomes possible by observing the coefficients of the decomposition matrix.

References

Abrahamyan, A., Silva, L. L., Dakin, S. C., Carandini, M., and Gardner, J. L. (2016). Adaptable history biases in human perceptual decisions. *Proc Natl Acad Sci USA*, 113(25):E3548–E3557.

Barik, K., Daimi, S. N., Jones, R., Bhattacharya, J., and Saha, G. (2019). A machine learning approach to predict perceptual decisions: an insight into face pareidolia. *Brain Inf.*, 6(1):2.

Benwell, C. S. Y., Coldea, A., Harvey, M., and Thut, G. (2021). Low pre-stimulus EEG alpha power amplifies visual awareness but not visual sensitivity. *Eur J Neurosci*, page ejn.15166.

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. *Psychological Review*, 113(4):700–765.

Brown, S. D. and Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, 57(3):153–178.

Busemeyer, J. R. and Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, 100(3):432–459.

Chen, Y., He, H., Xu, P., Wang, J., Qiu, Y., Feng, W., Luo, Y., Hu, L., and Guan, Q. (2020). The Weakened Relationship Between Prestimulus Alpha Oscillations and Response Time in Older Adults With Mild Cognitive Impairment. *Front. Hum. Neurosci.*, 14:48.

Cisek, P., Puskas, G. A., and El-Murr, S. (2009). Decisions in Changing Conditions: The Urgency-Gating Model. *Journal of Neuroscience*, 29(37):11560–11571.

Ditterich, J., Mazurek, M. E., and Shadlen, M. N. (2003). Microstimulation of visual cortex affects the speed of perceptual decisions. *Nat Neurosci*, 6(8):891–898.

Evans, N. J. and Wagenmakers, E.-J. (2020). Evidence Accumulation Models: Current Limitations and Future Directions. *TQMP*, 16(2):73–90.

Glaze, C. M., Kable, J. W., and Gold, J. I. (2015). Normative evidence accumulation in unpredictable environments. *eLife*, 4:e08825.

Gold, J. I. and Shadlen, M. N. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends in Cognitive Sciences*, 5(1):10–16.

Gold, J. I. and Shadlen, M. N. (2007). The Neural Basis of Decision Making. *Annu. Rev. Neurosci.*, 30(1):535–574.

Hafemeister, L., Gaussier, P., Maillard, M., Boucenna, S., and Giovannangeli, C. (2010). Perception: Insights from the sensori-motor approach. In *2010 2nd European Workshop on Visual Information Processing (EUVIP)*, pages 261–268, Paris, France. IEEE.

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. *Journal of Neuroscience*, 25(45):10420–10436.

Iemi, L., Chaumon, M., Crouzet, S. M., and Busch, N. A. (2017). Spontaneous Neural Oscillations Bias Perception by Modulating Baseline Excitability. *J. Neurosci.*, 37(4):807–819.

Jelic, V. and Marsiglio, F. (2012). The double-well potential in quantum mechanics: a simple, numerically exact formulation. *Eur. J. Phys.*, 33(6):1651–1666.

Kloosterman, N. A., de Gee, J. W., Werkle-Bergner, M., Lindenberger, U., Garrett, D. D., and Fahrenfort, J. J. (2019). Humans strategically shift decision bias by flexibly adjusting sensory evidence accumulation. *eLife*, 8:e37321.

Laming, D. R. (1968). *Information Theory of Choice-Reaction Times*. New York: Academic Press, academic press edition.

Lange, J., Oostenveld, R., and Fries, P. (2013). Reduced Occipital Alpha Power Indexes Enhanced Excitability Rather than Improved Visual Perception. *Journal of Neuroscience*, 33(7):3212–3220.

Langner, O., Dotsch, R., Bijlstra, G., Wigboldus, D. H. J., Hawk, S. T., and van Knippenberg, A. (2010). Presentation and validation of the Radboud Faces Database. *Cognition & Emotion*, 24(8):1377–1388.

Latimer, K. W., Huk, A. C., and Pillow, J. W. (2017). No cause for pause: new analyses of ramping and stepping dynamics in LIP (Rebuttal to Response to Reply to Comment on Latimer et al 2015). preprint, Neuroscience.

Latimer, K. W., Yates, J. L., Meister, M. L. R., Huk, A. C., and Pillow, J. W. (2015). Single-trial spike trains in parietal cortex reveal discrete steps during decision-making. *Science*, 349(6244):184–187.

Moehlis, J., Brown, E., Holmes, P., and Cohen, J. D. (2004). Optimizing reward rate in two alternative choice tasks: Mathematical formalism. Tech. Rep. 04-01, Princeton University, Ceneter for the Study of Brain, Mind and Behavior, Princeton, NJ.

Petro, N. M., Thigpen, N. N., Garcia, S., Boylan, M. R., and Keil, A. (2019). Pre-target alpha power predicts the speed of cued target discrimination. *NeuroImage*, 189:878–885.

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

Prat-Ortega, G., Wimmer, K., Roxin, A., and de la Rocha, J. (2021). Flexible categorization in perceptual decision making. *Nat Commun*, 12(1):1283.

Rabbitt, P. M. (1966). Errors and error correction in choice-response tasks. *Journal of Experimental Psychology*, 71(2):264–272.

Railo, H., Piccin, R., and Lukasik, K. M. (2021). Subliminal perception is continuous with conscious vision and can be predicted from prestimulus electroencephalographic activity. *Eur J Neurosci*, 54(3):4985–4999.

Rassi, E., Wutz, A., Müller-Voggel, N., and Weisz, N. (2019). Prestimulus feedback connectivity biases the content of visual experiences. *Proc Natl Acad Sci USA*, 116(32):16056–16061.

Ratcliff, R. (1978). A Theory of Memory Retrieval. *Psychological Review*, 85:59–108.

Ratcliff, R. and McKoon, G. (2008). The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. *Neural Computation*, 20(4):873–922.

Ratcliff, R. and Rouder, J. N. (1998). Modeling Response Times for Two-Choice Decisions. *Psychol Sci*, 9(5):347–356.

Ratcliff, R. and Rouder, J. N. (2000). A diffusion model account of masking in two-choice letter identification. *Journal of Experimental Psychology: Human Perception and Performance*, 26(1):127–140.

Ratcliff, R. and Smith, P. L. (2004). A Comparison of Sequential Sampling Models for Two-Choice Reaction Time. *Psychological Review*, 111(2):333–367.

Roitman, J. D. and Shadlen, M. N. (2002). Response of Neurons in the Lateral Intraparietal Area during a Combined Visual Discrimination Reaction Time Task. *J. Neurosci.*, 22(21):9475–9489.

Rolls, E. T. and Treves, A. (1999). *Neural networks and brain function*. Oxford Univ. Press, Oxford, reprint edition.

Roxin, A. and Ledberg, A. (2008). Neurobiological Models of Two-Choice Decision Making Can Be Reduced to a One-Dimensional Nonlinear Diffusion Equation. *PLoS Comput Biol*, 4(3):e1000046.

Samaha, J., Iemi, L., and Postle, B. R. (2017). Prestimulus alpha-band power biases visual discrimination confidence, but not accuracy. *Consciousness and Cognition*, 54:47–55.

Schurger, A. (2018). Specific Relationship between the Shape of the Readiness Potential, Subjective Decision Time, and Waiting Time Predicted by an Accumulator Model with Temporally Autocorrelated Input Noise. *eNeuro*, 5(1):ENEURO.0302–17.2018.

Shinn, M., Ehrlich, D. B., Lee, D., Murray, J. D., and Seo, H. (2020a). Confluence of Timing and Reward Biases in Perceptual Decision-Making Dynamics. *J. Neurosci.*, 40(38):7326–7342.

Shinn, M., Lam, N. H., and Murray, J. D. (2020b). A flexible framework for simulating and fitting generalized drift-diffusion models. *eLife*, 9:e56938.

Strogatz, S. H. (2015). *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering*. Westview Press, a member of the Perseus Books Group, Boulder, CO, second edition edition. OCLC: ocn842877119.

Taesler, P. and Rose, M. (2016). Prestimulus Theta Oscillations and Connectivity Modulate Pain Perception. *J. Neurosci.*, 36(18):5026–5033.

Urai, A. E., de Gee, J. W., Tsetsos, K., and Donner, T. H. (2019). Choice history biases subsequent evidence accumulation. *eLife*, 8:e46331.

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

van Dijk, H., Schoffelen, J.-M., Oostenveld, R., and Jensen, O. (2008). Prestimulus Oscillatory Activity in the Alpha Band Predicts Visual Discrimination Ability. *Journal of Neuroscience*, 28(8):1816–1823.

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

Wang, X.-J. (2002). Probabilistic Decision Making by Slow Reverberation in Cortical Circuits. *Neuron*, 36(5):955–968.

Wong, K.-F., Huk, A. C., Shadlen, M. N., and Wang, X.-J. (2007). Neural circuit dynamics underlying accumulation of time-varying evidence during perceptual decision making. *Front. Comput. Neurosci.*, 1.

Wong, K.-F. and Wang, X.-J. (2006). A Recurrent Network Mechanism of Time Integration in Perceptual Decisions. *Journal of Neuroscience*, 26(4):1314–1328.

Wöstmann, M., Waschke, L., and Obleser, J. (2019). Prestimulus neural alpha power predicts confidence in discriminating identical auditory stimuli. *Eur J Neurosci*, 49(1):94–105.

Yang, Y.-F., Brunet-Gouet, E., Burca, M., Kalunga, E. K., and Amorim, M.-A. (2020). Brain Processes While Struggling With Evidence Accumulation During Facial Emotion Recognition: An ERP Study. *Front. Hum. Neurosci.*, 14:340.

Zylberberg, A. and Shadlen, M. N. (2016). Cause for pause before leaping to conclusions about stepping. preprint, Neuroscience.