Double responding: A new constraint for models of speeded decision making

Evans, N.J.; Dutilh, G.; Wagenmakers, E.-J.; Maas, H.L.J. van der

DOI
10.1016/j.cogpsych.2020.101292

Publication date
2020

Document Version
Final published version

Published in
Cognitive Psychology

Cognitive Psychology 121 (2020) 101292

Contents lists available at ScienceDirect

Cognitive Psychology

journal homepage: www.elsevier.com/locate/cogpsych

Double responding: A new constraint for models of speeded decision making

Nathan J. Evans[a,*], Gilles Dutilh[b], Eric-Jan Wagenmakers[a], Han L.J. van der Maas[a]

[a] *Department of Psychology, University of Amsterdam, the Netherlands*
[b] *Department of Clinical Research, University of Basel Hospital, Switzerland*

ARTICLE INFO

Keywords:
Evidence accumulation models
Model mimicry
Double responding
Response time models

ABSTRACT

Evidence accumulation models (EAMs) have become the dominant models of speeded decision making, which are able to decompose choices and response times into cognitive parameters that drive the decision process. Several models within the EAM framework contain fundamentally different ideas of how the decision making process operates, though previous assessments have found that these models display a high level of mimicry, which has hindered the ability of researchers to contrast these different theoretical viewpoints. Our study introduces a neglected phenomenon that we term "double responding", which can help to further constrain these models. We show that double responding produces several interesting benchmarks, and that the predictions of different EAMs can be distinguished in standard experiment paradigms when they are constrained to account for the choice response time distributions and double responding behaviour in unison. Our findings suggest that lateral inhibition (e.g., the leaky-competing accumulator) provides models with a universal ability to make accurate predictions for these data. Furthermore, only models containing feed-forward inhibition (e.g., the diffusion model) performed poorly under both of our proposed extensions of the standard EAM framework to double responding, suggesting a general inability of feed-forward inhibition to accurately predict these data. We believe that our study provides an important step forward in further constraining models of speeded decision making, though additional research on double responding is required before broad conclusions are made about which models provide the best explanation of the underlying decision-making process.

1. Introduction

Understanding how humans make decisions has been a topic of great theoretical and applied interest within psychology, neuroscience, and economics, ranging from decisions as simple as what direction a cloud of dots is moving towards (Pilly & Seitz, 2009), to decisions as complex as what risks to take in life or death situations (Tversky & Kahneman, 1981). The area of speeded decision making attempts to use quick, simple decisions to better understand the fundamental parts of the decision process. These tasks usually involve paradigms where participants make large numbers of quick decisions to simple stimuli and tasks (for examples, see Dutilh et al., 2018; Evans, Bennett, & Brown, 2018; Evans & Brown, 2017; Evans & Hawkins, 2019; Evans, Rae, Bushmakin, Rubin, & Brown, 2017; Ratcliff & Rouder, 1998; Van Ravenzwaaij, Dutilh, & Wagenmakers, 2012), and researchers are most commonly interested in the response given for each decision and the time taken to respond, with response speed and accuracy often known to trade off with

* Corresponding author.
E-mail address: nathan.j.evans@uon.edu.au (N.J. Evans).

https://doi.org/10.1016/j.cogpsych.2020.101292
Received 14 January 2019; Received in revised form 5 February 2020; Accepted 26 February 2020
Available online 24 March 2020

one another (e.g., Luce, 1986).

Evidence accumulation models (EAMs; Laming, 1968; Ratcliff, 1978; Stone, 1960) have become the dominant models of speeded decision making, which are able to decompose the variables of response time and response choice into cognitive parameters that drive the decision process. Specifically, EAMs propose that evidence accumulates for each alternative at some rate (known as the "drift rate") until one alternative reaches some level of evidence (known as the "threshold") where a response is triggered. Several different models have been developed within the EAM framework that differ in the specifics of the proposed process, such as whether competition occurs between alternatives in accumulation (i.e., inhibition; Ratcliff, 1978; Usher & McClelland, 2001), whether evidence decays over time (i.e., leakage; Busemeyer & Townsend, 1993; Usher & McClelland, 2001), and how these specific parts of the process might operate (e.g., lateral inhibition [Usher & McClelland, 2001] vs. feed-forward inhibition [Ratcliff, 1978]). Importantly, each of these models reflect fundamentally different ideas of how the decision making process operates, with theoretically distinct components (Teodorescu & Usher, 2013).

Despite the theoretical differences between the models of speeded decision making, previous assessments have found that the models display a high level of mimicry: that is, that the predictions of the models can be difficult to distinguish from one another (Ditterich, 2010; Donkin, Brown, Heathcote, & Wagenmakers, 2011). Specifically, each of these models have been designed to account for several benchmark choice response time phenomena (Brown & Heathcote, 2008; Ratcliff, 1978; Ratcliff & Rouder, 1998; Ratcliff & Tuerlinckx, 2002; Usher & McClelland, 2001), and in the majority of cases make accurate predictions about the choice response time distributions. Crucially, the joint success of these models in explaining empirical data – due to their mimicry in choice response time data from standard experimental paradigms – has "hinder[ed] the use of models as instruments for testing theories" (Teodorescu & Usher, 2013, p.2), with researchers suggesting that they "cannot decide among competitive models that have the ability to mimic one" as "small differences in data may rank order them differently" (Leite & Ratcliff, 2010, p.264). One of the key goals of our study is to introduce a neglected phenomenon, "double responding" (explained in more detail below), which can help to further constrain these models. We show that double responding produces several interesting benchmarks, and that the predictions of different EAMs can be distinguished when they are constrained to account for the choice response time distributions and double responding behaviour in unison (without the need for novel experimental paradigms, such as in Teodorescu & Usher, 2013), with models that contain lateral inhibition generally making the most accurate predictions. However, this distinguishability did somewhat depend on the assumption about how double responses are triggered, with models that do not contain any inhibition providing almost equally accurate predictions as models that contain lateral inhibition under an alternative definition of the double responding process. However, these predictions were always distinguishable from models with feed-forward inhibition – such as the well-known diffusion model (Ratcliff, 1978) – which universally failed to provide accurate predictions for these data.

Previous research within speeded decision making has been largely focused on only the variables of response choice (i.e., accuracy) and response time. However, the number of variables that could potentially be measured within rapid decision making is limitless, and attempting to measure and understand these variables can further our understanding of the decision making process. For example, another variable that has been previously investigated is choice confidence (Pleskac & Busemeyer, 2010; Ratcliff & Starns, 2009, 2013; Vickers & Lee, 1998), where researchers are not only interested in what choice a person makes, but also how confident they are in that choice. Several interesting phenomena have been found within the choice confidence literature, such as the "hard-easy effect" (Gigerenzer, Hoffrage, & Kleinbölting, 1991; Rahnev & Denison, 2018), where participants tend to be over-confident when the task is hard, and under-confident when the task is easy. Interest in confidence has resulted in the development of models to explain these phenomena, which attempt to account for both response time and confidence, and mostly involve extensions of the standard EAM framework (Moran, Teodorescu, & Usher, 2015; Pleskac & Busemeyer, 2010; Ratcliff & Starns, 2009, 2013). However, choice confidence has some limitations that prevent it from being a focus in "standard" rapid decision making investigations, as it requires the use of altered paradigms where participants are able to give a range of responses based on their confidence level, and there is no universal agreement on the best ways to make and model these measurements (e.g., one-stage [Ratcliff & Starns, 2009, 2013] vs. two-stage [Pleskac & Busemeyer, 2010; Van Zandt & Maldonado-Molina, 2004]).

Although some additional variables require altered paradigms to measure, this is not necessarily the case for all potential variables of interest. One variable that can be easily measured within standard decision making paradigms is what we term "double responding" (also see "error correcting responses"; Rabbitt & Rodgers, 1977), where after making a response in a speeded decision making task, participants make a second rapid response for another alternative. For example, in the classic two-alternative-forced-choice (2AFC) random dot motion task, where a participant has to decide whether a cloud of dots is moving towards the left or right of the screen, the participant may respond "left", but then straight afterwards also respond "right". This second "right" response is what we would term a "double response", and can be thought of as a "spillover" of information from the decision-making process. More specifically, in the context of EAMs, double responses reflect situations where different competing response alternatives each accumulate enough evidence to trigger a response in quick succession, where the evidence accumulation process continues despite a response having already been triggered for one alternative (i.e., the "spillover"). Double-responding provides additional information beyond response choice and response time on how the brain processes evidence in decision-making, which could be fundamental to better understanding the decision-making process as a whole; however, despite double responding being anecdotally observed, it is rarely studied.

One exception are the studies of Rabbitt (1968), Rabbitt and Rodgers (1977), and Rabbitt (2002), who assessed participants' "error correcting responses" (i.e., a subset of double responses where the initial response was an error), finding that these occurred regardless of whether or not participants were instructed to make them, and that these responses were substantially quicker than other responses used to signify an error. Another similar area of investigation has been two-stage confidence judgments, where participants make an initial response in favour of an alternative, and then a second response for their level of confidence in their choice. In paradigms where the second confidence judgment allows participants to show greater confidence for the alternative that was not previously selected, researchers have attempted to model this change in preference in some cases (see Van Zandt & Maldonado-Molina, 2004 for "response reversals", though note that their investigation did not model the choice response time distributions). However, these studies have focused on testing the ability of the models to explain the possibility of a change in preference based on the confidence calculation, and have not attempted to use this additional source of information to better understand the decision-making process as a whole. One other similar area of investigation has been "partial errors", where the motor movements of participants are recorded using electromyography (EMG), allowing researchers to detect trials where participants make initial movements towards the error response alternative, which they inhibit before making a response and then respond to correctly (Burle, Spieser, Servant, & Hasbroucq, 2014; Coles, Gratton, Bashore, Eriksen, & Donchin, 1985; Servant, White, Montagnini, & Burle, 2015). Although these partial errors have mostly been assessed at a descriptive level, Servant et al. (2015) also performed simulations to compare the mean partial error latency observed in their flanker data to the predictions of EAMs variants adapted to conflict tasks. However, these simulations were only based on a single set of parameter values – those that provided the best fit to the regular choice response time distributions – and therefore, did not attempt to jointly assess partial errors and the choice response time distributions. Importantly, the assessment of double responding does not require altered experimental paradigms, psychophysiological data, or the large number of response alternatives used in many of the error correction studies of Rabbitt and colleagues, and is simply additional information that can be measured within any standard decision making task of interest.

From a theoretical perspective, double responding appears to fit within the broader "change of mind" literature, where researchers are interested in how people change their preferences during and/or after decisions. Initially, Rabbitt and colleagues (Maylor & Rabbitt, 1987; Rabbitt, 1966, 1967, 1968, 1969, 2002; Rabbitt, Cumming, & Vyas, 1978; Rabbitt & Rodgers, 1977; Rabbitt & Vyas, 1981) were interested in people's ability to detect and correct errors, and how errors effected subsequent responses (e.g., post-error slowing; Rabbitt, 1969). Specifically, Rabbitt and colleagues found that people were able to detect most of their errors by making an "error detecting" response that was separate from other task responses (Rabbitt et al., 1978), though error detection was improved (Rabbitt, 2002) and faster (Rabbitt, 1968) when participants were required to "correct" their error by making the responses that they should have made (i.e., error correcting responses), and that these correcting responses were less effected by reductions in the amount of time allowed for error detections to be made (Rabbitt, 2002). Post-error responses were also found to be faster when they were identical to the error correcting response, suggesting that post-error slowing may be due to a conflict between a desire to correct the previous response and to make the correct response to the current stimulus (Rabbitt & Rodgers, 1977; Vickers & Lee, 2000). Recent research has taken neural approaches to assessing post-error behaviour (Falkenstein, Hohnsbein, Hoormann, & Blanke, 1991; Yeung, Botvinick, & Cohen, 2004), and has had a greater focus on how people change their mind during a decision, using movement-based response mechanisms (e.g., hand/arm movement, a joystick, or electromyographic measurement) to assess how people change from favouring one alternative to another. Interestingly, people have been found to show changes of mind both in tasks that are intentionally designed with conflicting information (Hasbroucq, Burle, Akamatsu, Vidal, & Possamai, 2001), as well as tasks that simply require sampling information from a noisy stimulus (Resulaj, Kiani, Wolpert, & Shadlen, 2009). Furthermore, research into saccade-based response paradigms, which are common in research with macaque monkeys, has found that both humans and macaque monkeys display corrective saccades on some proportion of trials after making a saccade for the incorrect response (Andersen & Buneo, 2002; Becker & Jürgens, 1979; Camalier et al., 2007; Colby & Goldberg, 1999; Murthy et al., 2007). All of these phenomena appear to show a common theoretical underpinning to double responding, where after an initial error participants either make another response, change the motor action of the current response, or alter their decision making on the subsequent trial.

Importantly, the additional information contained within double responding may provide critical constraint to distinguish between the different potential dynamics of the decision-making process. As discussed above, most EAMs can provide a good account of the choice response time distributions for most standard tasks (e.g., Brown & Heathcote, 2008; Ratcliff & Rouder, 1998; Usher & McClelland, 2001), forcing researchers to extend these models to new paradigms to distinguish between their predictions (Teodorescu & Usher, 2013; Tsetsos, Gao, McClelland, & Usher, 2012). However, another route to distinguish between models is constraining them to account for more sources of data, making it more difficult to accurately account for all sources of data. For example, accounting for the mean response time is a much easier task for models than accounting for the shape of the response time distribution, and therefore, accounting for the entire response time distribution can more clearly distinguish between models than accounting for only mean response times (Luce, 1986). Double responding provides an additional source of information about the decision-making process, serving as a behavioural manifestation of a spillover of information accumulation from the initial decision. EAMs also offer a somewhat natural extension to double responses, where evidence continues to accumulate after a response is made (see Pleskac & Busemeyer, 2010; Rabbitt & Vyas, 1981 for empirical validation of this idea, and Merkle & Van Zandt, 2006; Pleskac & Busemeyer, 2010; Van Zandt & Maldonado-Molina, 2004 for modelling frameworks that have used this "continued accumulation" assumption), and a double response is triggered if the evidence for the other alternative reaches the threshold. Therefore, double responses appear to provide a theoretically reasonable additional constraint for EAMs, which may provide a greater understanding of the dynamics of the decision-making process by distinguishing between the predictions of different models, which each represent a theoretically distinct decision process.

Our study aims to introduce double responding as a source of additional constraint for models of decision making, with double responding providing additional information about the complete decision-making process. As discussed above, double responding can be measured in standard decision making tasks, and can be easily recorded and analysed with little extra effort from the experimenter. Our study assesses implicit double responding behaviour, where participants are *not* directly instructed to make a second response, and the observed double responses can be thought of as a motor output resulting from a spillover of information from the decision-making process. Our primary reason for choosing implicit double responding behaviour over explicit double responding behaviour (i.e., where participants *are* directly instructed that they are allowed to make a second response) is due to our desire to make inferences that are generalizable to regular 2AFC tasks, and that informing participants that they are able to change their responses may alter how they respond initially. For example, if participants are made aware that they can correct their errors for a fixed amount of time after their initial response, then they may become less cautious in their initial responding compared to situations when they are not provided with this information (see also interference effects in decision-making tasks that involve initial responses; Kvam, Pleskac, Yu, & Busemeyer, 2015; Yearsley & Pothos, 2016). However, we believe that the assessment of explicit double responding behaviour is also an important avenue for future research, and we discuss the potential strengths and limitations of both implicit and explicit double responding in the discussion section.

More specifically, we aim to make two key contributions that showcase how double responding can further our understanding of the decision-making process. First, our study provides some general assessments on summary statistics relating to double responding, in order to gain a better understanding of when double responding occurs, and create potential qualitative benchmarks for models of decision-making to meet. Second, our study uses double responding to assess and compare different EAMs of decision making, in order to see which models are able to naturally account for this additional source of data within their framework, and therefore, which model provides the most accurate account of the dynamics of the decision-making process.

2. Experiment 1 (Dutilh, Vandekerckhove, Tuerlinckx, & Wagenmakers, 2009)

Our first experiment uses the data from Dutilh et al. (2009), who aimed to understand how the latent components of the decision process – as formalized in the diffusion model (Ratcliff, 1978) – changed across practice. Specifically, Dutilh et al. (2009) collected 25 blocks of trials from each participant, fitting the diffusion model to each of these blocks to assess which parameters changed over practice, and the manner in which they changed. Although double responding was recorded within this experiment, the focus of Dutilh et al. (2009) was purely on the response time and choice of each trial (and the diffusion model parameters estimated from them), and double responding was not mentioned within their article. We believe that this data set provides a particularly good starting point for assessing double responding as it contained a large number of trials per participants (10,000), allowing for a detailed assessment of double responding even if their occurrence is quite rare. Furthermore, as discussed by Smith and Little (2018), experimental designs with a large number of trials per participant provide a powerful design for making inferences about individuals, and each participant can be considered a replicate for the inference. Full details of the experimental setup can be found in Dutilh et al. (2009), though we provide a brief outline below.

2.1. Method

2.1.1. Experiment details

Four native Dutch speakers each completed 10,000 trials of a lexical decision task, where participants were presented with a letter string on each trial and had to identify, via button press, whether or not the letter string was a (Dutch) word (i.e., "word" or "nonword"). Speed emphasis was used as a between-subjects manipulation, with participants either being told to focus on responding quickly, even if it meant decreases in accuracy (2 participants), or responding accurately, even if it took a little longer (2 participants). Once participants had made an initial response any additional responses were recorded for the following 250 ms, as this was the minimum possible amount of time between trials. We defined "double responses" as being a response after the initial response for the opposite alternative, meaning that second responses for the original response were not considered. For example, in the lexical decision task, a double response would be if a participant initially responded that a letter string was a "nonword", but then less than 250 ms after this response also responded "word".

2.1.2. Theoretical models

In addition to basic assessments of double responses, we also extended a series of decision making models to incorporate this phenomenon. Specifically, all models were defined within the evidence accumulation framework, where evidence accumulates for each decision alternative until a threshold level of evidence is reached for one response, and the response is triggered. To incorporate double responding into the models, we assumed that the evidence continues to accumulate for each alternative after the initial response, which is in line with previous research (Pleskac & Busemeyer, 2010; Rabbitt & Vyas, 1981; Resulaj et al., 2009), as well as one of the potential verbal explanations proposed by Rabbitt and Vyas (1981) for the occurrence of error correcting responses. If the threshold of level of evidence is reached for the alternative that was not the initial response (i.e., the losing alternative) within 250 ms (i.e., the amount of time available for double responses in the task), we assumed that this caused a second response to be triggered, resulting in a "double response" occurring. A diagram of this process can be seen in Fig. 1.

We defined a "tree" of 9 models to test, based on the inclusion or exclusion of a series of components that are commonly included in EAMs: feed-forward inhibition (Ratcliff, 1978), lateral inhibition (Usher & McClelland, 2001), and leakage (Busemeyer & Townsend, 1993; Ratcliff & Smith, 2004; Usher & McClelland, 2001). Our most basic model was a stochastic diffusion process with only a drift rate for each alternative, a threshold level of evidence, a random uniform distribution of starting evidence, and some time dedicated to non-decision processes.[1] This is commonly known as the "racing diffusion model" (Tillman & Logan, 2017), where the continuous evidence accumulation for alternative i can be formally written as:

[1] Note that this definition means that there was only a single source of between-trial variability allowed in all models: starting point variability.

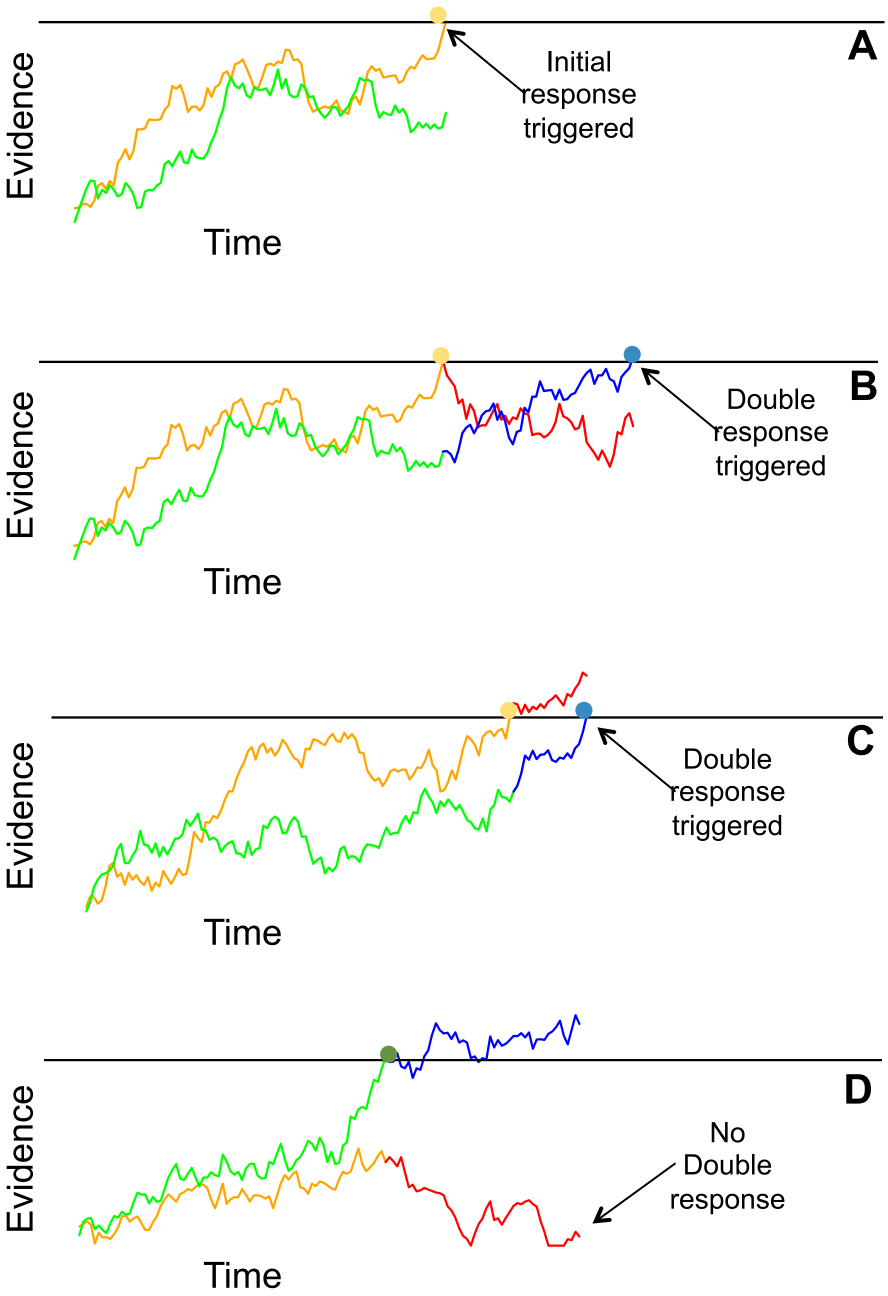

Fig. 1. A: A diagram of the typical decision process proposed within evidence accumulation models, with different decision alternatives represented by the orange and green accumulation lines. In this example, the orange accumulator reaches the threshold first, triggering a response. **B:** A diagram of our extension from the standard decision process to the double responding process, where the evidence continues to accumulate after a decision is triggered. The orange and green accumulators are replaced by red and blue accumulators, respectively, after the initial decision is made, for clarity on when the double responding process begins. In this example, the blue accumulator reaches the threshold after the initial orange response, triggering a double response for the green/blue alternative. **C:** Another example of a decision process that would result in a double response, where the orange/red alternative remains above the threshold. **D:** An example of a decision process that would **not** result in a double response. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

$$dx_i = [v_i]dt + [\sigma_i \ \epsilon]\sqrt{dt} \tag{1}$$

where x_i is the evidence for alternative i, v_i is the drift rate for alternative i, ϵ is a random variable from the univariate (i.e., independent noise between accumulators) standard normal distribution that provides the within-trial noise, σ_i is the scale of the within-trial noise for alternative i, and t is time.

From this base model, we added either feed-forward inhibition, lateral inhibition, or leakage parameter to the base model, to create an additional 3 models. Feed-forward inhibition reduces future evidence accumulation based on the accumulation rates of the other alternatives, and can be formally written as an extension of the base model:

$$dx_i = [v_i - \beta \sum_{j \neq i}^{n} v_j] dt + [\sigma_i \epsilon - \beta \sum_{j \neq i}^{n} \sigma_j \epsilon] \sqrt{dt} \tag{2}$$

where β is a parameter that controls the amount of inhibition, and n is the number of response alternatives. Note that the feed-forward inhibition is applied to both the drift rate and the within-trial noise (in both Eqs. 2 and 5), meaning that the noise processes of the different accumulators have a perfect negative correlation with one another. We also created a restricted feed-forward inhibition model with β fixed at 1, which is mathematically equivalent to the diffusion model (Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006). Lateral inhibition reduces future evidence accumulation based on the previously accumulated evidence for the other alternatives, and can be formally written as an extension of the base model:

$$dx_i = [v_i - \beta \sum_{j \neq i}^{n} x_j] dt + [\sigma_i \epsilon] \sqrt{dt} \tag{3}$$

Leakage reduces the accumulated evidence for an alternative, and can be formally written as an extension of the base model:

$$dx_i = [v_i - \lambda x_i] dt + [\sigma_i \epsilon] \sqrt{dt} \tag{4}$$

where λ is the leakage rate.

From here, we added leakage to each potential type of inhibition, creating another 3 models, giving 8 in total. The feed-forward inhibition model with leakage was defined as:

$$dx_i = [v_i - \lambda x_i - \beta \sum_{j \neq i}^{n} v_j] dt + [\sigma_i \epsilon - \beta \sum_{j \neq i}^{n} \sigma_j \epsilon] \sqrt{dt} \tag{5}$$

and again, we created a restricted feed-forward inhibition model with β fixed at 1 and leakage, which is equivalent to the Ornstein-Uhlenbeck process (Busemeyer & Townsend, 1993). The lateral inhibition model with leakage was defined as:

$$dx_i = [v_i - \lambda x_i - \beta \sum_{j \neq i}^{N} x_j] dt + [\sigma_i \epsilon] \sqrt{dt} \tag{6}$$

Lastly, we added an evidence truncation at 0 (i.e., evidence cannot be negative) to the lateral inhibition with leakage model, which gave a final model that reflected the leaky-competing accumulator (LCA; Usher & McClelland, 2001). The full tree of models can be seen in Table 1, and code for all models can be found on OSF (https://osf.io/7jmb2/).

Note that for the model fitting process we excluded trials with extremely fast response times (<250 ms; used in both Experiment 1 and Experiment 2), as these responses are likely too fast to be the result of the regular decision process. All models were fit using Bayesian parameter estimation, where we estimated the posterior distributions using Differential Evolution Markov chain Monte Carlo (DE-MCMC; Ter Braak, 2006; Turner, Sederberg, Brown, & Steyvers, 2013), as this appeared to provide a superior method for traversing the complex likelihood space than regular maximum likelihood estimation fitting. However, as our current study is mostly interested in the broad, qualitative predictions of the models, comparisons between models were largely made using 1,000,000 simulated predictions from the models using the posterior sample that had the highest probability of the data given the parameters (i.e., $P(D|\theta)$, which is the maximum likelihood estimate under the assumption that the priors are sufficiently broad to include the maximum likelihood estimate), rather than formal model selection metrics that involve the entire posterior distribution. To ensure robustness, the highest $P(D|\theta)$ (i.e., maximum likelihood) was determined by fitting each model 5 independent times (to minimize the potential influence of local maxima), obtaining the 50 highest $P(D|\theta)$ samples from each independent fit and averaging over 50 pseudo-likelihood estimates for each of these samples (to minimize the potential influence of simulation noise), and taking the sample

Table 1
The 9 models included within our assessment of double responding, and the components beyond the base model that they included. Models are presented in the same order as in the main text. Within the table and subsequent results figures, we use abbreviations for the racing diffusion model (RDM), feed-forward inhibition (FFI), the diffusion model (DM), lateral inhibition (LI), leakage (Leak), the Ornstein-Uhlenbeck process (OU), and the leaky-competing accumulator (LCA).

Model	Lateral Inhibition	Feed-forward Inhibition	Leakage	Truncation at 0
RDM (Tillman & Logan, 2017)				
FFI		✓		
DM (Ratcliff, 1978)		Fixed at 1		
LI	✓			
Leak			✓	
FFI + Leak		✓	✓	
OU (Busemeyer & Townsend, 1993)		Fixed at 1	✓	
LI + Leak	✓		✓	
LCA (Usher & McClelland, 2001)	✓		✓	✓

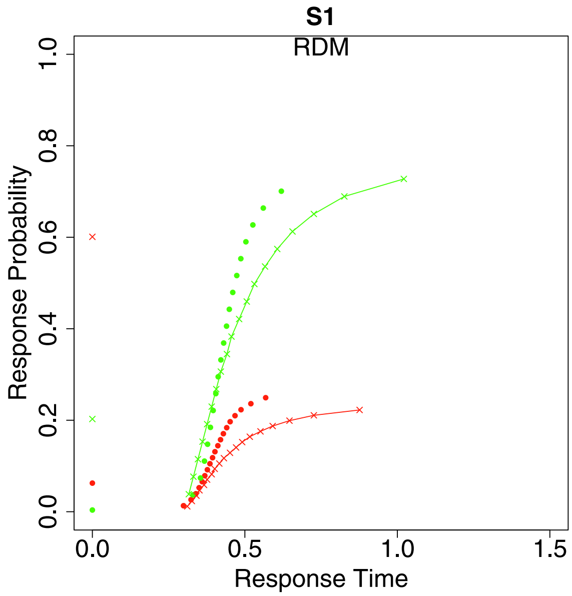

Fig. 2. Cumulative density function (CDF) plot of the empirical (dots) and racing diffusion model (RDM) predicted (lines and crosses) standard response time distributions (i.e., response times for the initial response) and double response proportions for correct (green; P(DR|C) in Table 2) and error (red; P(DR|E) in Table 2) initial responses. The x-axis displays the response times, and the y-axis displays the response proportions, with the double response proportions placed at 0 on the x-axis, to avoid confusion with the standard choice response time distributions. This format continues throughout CDF plots in subsequent figures. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

(out of all 250 samples across the 5 fits) with the highest averaged $P(D|\theta)$. However, as we discuss below, we also include a complementary quantitative comparison using the Bayesian information criterion (BIC; Schwarz, 1978), which uses the same method as above to find the maximum likelihood estimate.

Specifically, we assessed the ability of the models to account for the empirical data using defective cumulative density function (CDF) plots, plotting the empirical response time distributions for correct and incorrect responses against the model predicted distributions. An example of a CDF plot can be seen in Fig. 2, which displays the RDM predictions for participant S1 of Dutilh et al. (2009) when jointly fitting choice response time distributions and double response proportions. The dots display the defective cumulative density function for the empirical response time distributions, with different dots displaying different quantiles of the distributions (reflected by the y-axis), and the crosses with joining lines display the defective cumulative density function for the model predictions. Green dots/crosses/lines indicate correct initial responses, whereas red dots/crosses/lines indicate incorrect initial responses. The double response proportions are positioned at x = 0, with their position on the y-axis indicating the number of double responses after correct (green) and incorrect (red) initial responses, for the empirical data (dots) and model predictions (crosses). A close match between the dots and the crosses indicate an accurate account of the data, and greater discrepancies between the dots and crosses indicate greater misfit. As can be seen, there is a large amount of misfit for the predictions of the RDM, with the RDM predicting response time distributions with greater variance than the empirical data (crosses cover more of the x-axis than the dots), a higher accuracy than the empirical data (the final green cross finishes above the final green dot on the y-axis), and a higher proportion of double responses than the empirical data (crosses at x = 0 are higher on the y-axis than the dots). As our current study is mostly interested in the broad, qualitative predictions of the models, and how they compare to the empirical data, our comparisons are mostly based upon the visual fit and misfit of each model within the CDF plots (Navarro, 2019). However, to ensure that our inferences were not biased by our choices in which qualitative trends to assess (Evans, 2019c), we also provide a comparison between the models using BIC, which has been found to provide similar inferences to Bayes factors with uninformative priors in many situations (Evans, 2019a).

As the likelihood functions for most of our extended models are unknown, we used probability density approximation to obtain the likelihood function for the models (Holmes, 2015; Turner & Sederberg, 2014), which approximates the likelihood function using simulations from the model with a density kernel approximation. Specifically, for the two-choice task of Dutilh et al. (2009), the probability of an initial response occurring at time t for alternative i followed by a double response for other alternative j was given by:

$$f_i(t) \times F_j^{DR}(t_{max}^{DR}) \tag{7}$$

where $f(t)$ is the probability density function of the model for the initial response, $F^{DR}(t)$ is the cumulative density function of the model for the double response (i.e., conditioned on an initial response having already been made, with $t^{DR} = 0$ being the initial response), and t_{max}^{DR} is the maximum time available for a double response after an initial response, which in this experiment was

Table 2

Summary statistics for the participants in Experiment 1. Different rows provide different pieces of information, and different columns provide the information for different subjects. Within the table and subsequent tables, we use abbreviations for emphasis Emph, the total number of trials (N_{trials}), mean response time (MRT), probability of X (P(X)), correct responses (C), error responses (E), double responses (DR), correlation between X and Y (cor(X,Y)), trial number (t).

Sub	1	2	3	4
Emph	speed	speed	accuracy	accuracy
N_{trials}	10,000	10,000	10,000	10,000
MRT	0.4348	0.4933	0.5678	0.5296
P(C)	0.7344	0.8252	0.948	0.9286
P(DR)	0.0191	0.0255	0.0003	0.0007
P(DR|C)	0.0037	0.0039	0	0
P(DR|E)	0.0617	0.1276	0.0058	0.0098
P(C|DR)	0.1414	0.1255	0	0
cor(DR,RT)	0.0359	0.0163	−0.0134	0.0023
cor(DR,t)	0.0371	−0.0354	0.0167	0.0107
P(DR_t|DR_{t-1})	0.0209	0.0118	0	0
P(DR_t|¬DR_{t-1})	0.0191	0.0259	0.0003	0.0007

250 ms. When assessing the double response times, the probability of an initial response occurring at time t for alternative i followed by a double response at time t^{DR} for alternative j was given by:

$$f_i(t) \times f_j^{DR}(t^{DR}) \tag{8}$$

where $f^{DR}(t)$ is the probability density function of the model for the double response. The probability of an initial response occurring at time t for alternative i followed by no double response was given by:

$$f_i(t) \times S_j^{DR}(t_{max}^{DR}) \tag{9}$$

where $S^{DR}(t)$ is the survivor density function of the model for the double response. Importantly, this means that the models are not only constrained by the occurrence of a double response, but also the *lack of occurrence* of a double response, as models will be punished for both predicting too many, or too few, double responses. All models were simulated using Euler's method (e.g., Brown, Ratcliff, & Smith, 2006) with a time-step of 5 ms, through the simulation method and framework of Evans (2019b), and the density kernels for the PDA were created using 10,000 simulated trials each.

2.2. Results

2.2.1. Assessment of empirical trends

To begin, we provide some descriptive assessments of the performance of all four participants, and in particular, their double responding tendencies. These data can be seen in Table 2, where the columns give the different subjects. Note that our descriptive assessment for Experiment 1 does not attempt to make any population-level inferences about how participants under speed emphasis differ from participants under accuracy emphasis, as Experiment 1 only contains two participants per group, meaning that any generalizations to the entire population would be questionable. Rather, our aim is to provide a complete description of the trends in this sample of data for interested readers (similar to previous studies that evaluate EAMs based on a small number of participants; e.g., Ratcliff & Rouder, 1998), and then assess whether these trends are also present in Experiment 2, which uses a within-subjects manipulation of emphasis and a greater number of participants.

As can be seen, the participants who completed the experiment under speed emphasis (columns 1 and 2) have faster mean response times and lower proportions of correct responses than the participants who completed the experiment under accuracy emphasis (columns 3 and 4), which suggests that the emphasis instructions were adhered to. Interestingly, those under accuracy emphasis made very few double responses, whereas those in the speed condition made relatively more double responses. However, in all cases the number of double responses is relatively small, occurring on under 2.6% of trials at most. Double responses also very rarely occurred after correct responses, but occurred relatively more often after error responses (i.e., after up to 12.75% of errors). In addition, those in the speed condition were more likely to make a double response after an error, suggesting that the greater number of double responses in the speed condition were a result of more than just the greater number of errors made. However, this did not appear to be due to responses being faster under speed emphasis, as the trial response time and whether a double response was made showed little relation to one another (and in general, more double responses occurred after slower responses). Lastly, double responses did not appear to occur more often at certain parts of the experiment, nor did they appear to show any clear first order autocorrelation.

To gain some further insight on double responses, we also provide some descriptive assessments of the double response time (DRT) distributions for the two participants who made a large number of double responses (i.e., the two participants under speed emphasis). These distributions can be seen in Fig. 3, where both participants appear to show quite similar trends in their distributions. Interestingly, the DRT distributions appear to resemble the RT distributions of standard responses, showing a positive skew. The majority of double responses occur very rapidly after the initial response, suggesting that they serve as quick, almost identically

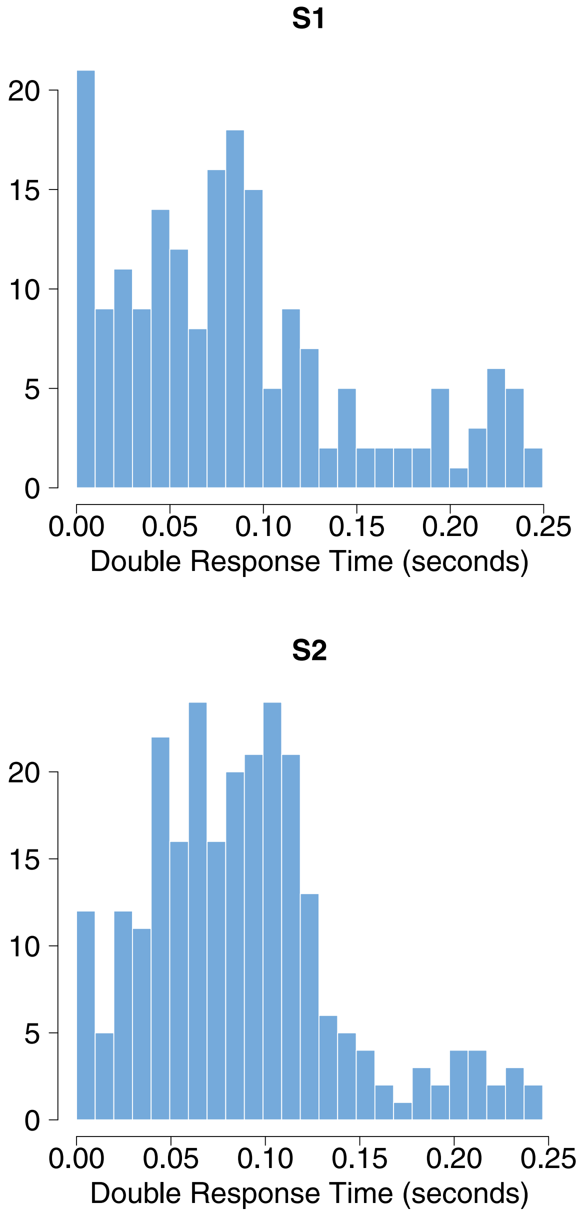

Fig. 3. Double response time distributions for the two participants (different histograms) in the speed emphasis condition of Experiment 1.

timed, “corrections” (i.e., mostly following error responses), though continue in smaller amounts out until the deadline. Another interesting question involving response time distributions is whether the initial responses on trials with double responses appear to be similar to those on regular trials, or whether the responses preceding double responses appear to be due to some other process (e.g., fast guessing; Vandekerckhove, Tuerlinckx, & Lee, 2008), which may be part of the reason for the additional response occurring afterwards. Fig. 4 plots the standard response time distributions for trials with and without double responses for each of the two participants above. As can be seen, there does not appear to be any clear qualitative difference between response times before double responses and those without double responses, and trials with double responses do not appear to have initial responses that are fast guesses.

2.2.2. Assessment of theoretical models

To begin assessing our tree of models in their ability to account for these data, we fit the models to the standard choice response time distributions of each participant. Although this is not the main question of interest, this provided an important baseline assessment, as all models should be able to provide a good account of these data. As suspected, all models provided a good account of the choice response time distributions, showing minor misfit at most (Figs. 5 and 6). We also compared the models in their ability to account for only (i.e., ignoring the response time distributions) the double response proportions (Figs. 5 and 6). All models appear to

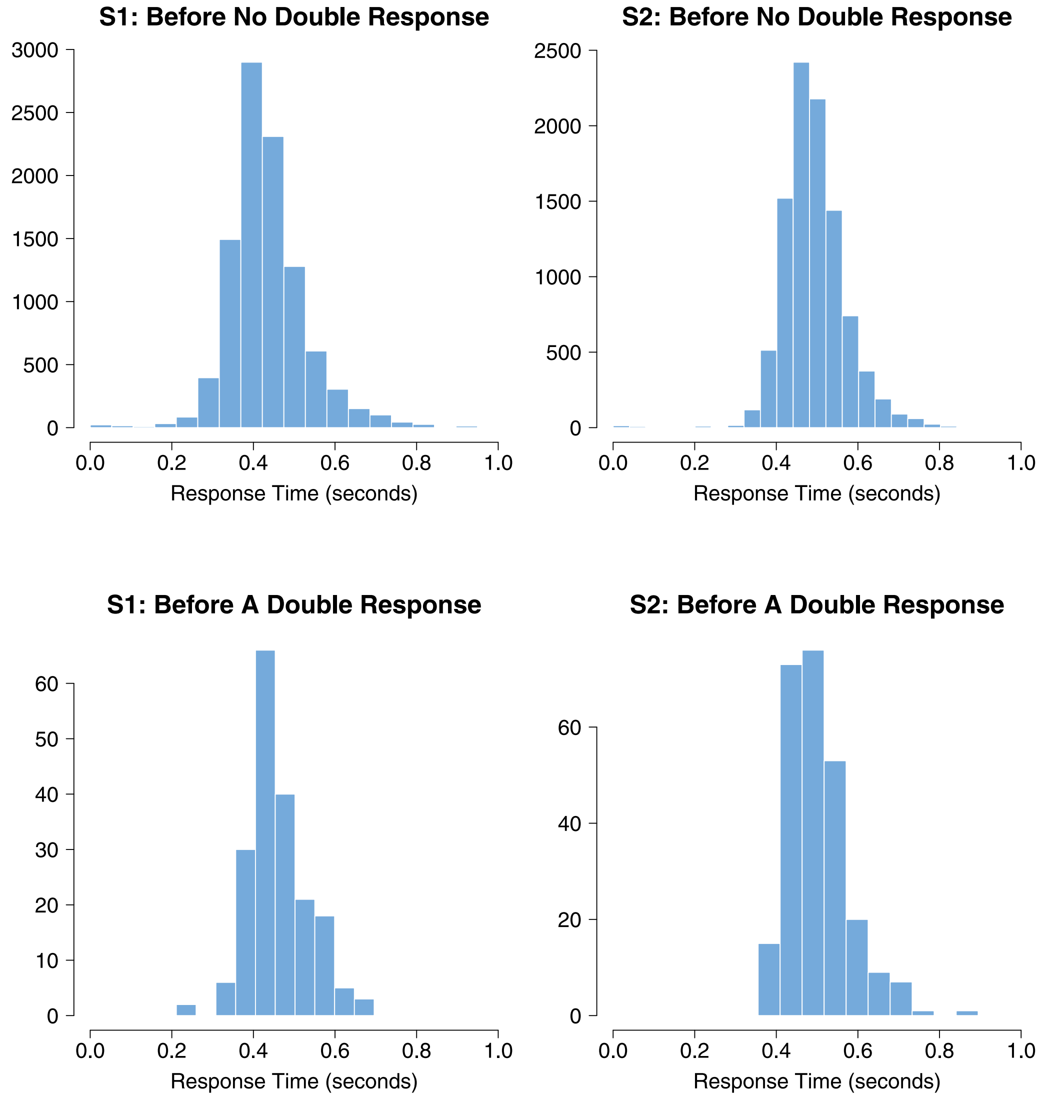

Fig. 4. Response time distributions for the two participants (columns) in the speed emphasis condition of Experiment 1. The top row displays the response time distributions for trials that did *not* have double responses following them, and the bottom row displays the response time distributions for trials that did have double responses following them. In general, these distributions seem to be fairly similar to one another, suggesting that double responses aren't clearly due to a non-standard decision process occurring during these trials.

provide a quantitatively accurate account of the double response proportions, though this is probably not surprising as each model contains at least 5 parameters and is only being fit to 2 data points (i.e., the proportions of double responses after corrects and errors), meaning that fitting to purely the double response proportions is likely not adequate to properly constrain these models.

Next, we fit the models jointly to the choice response time distributions and the double response proportions, in order to provide greater constraint on the models and hopefully better distinguish between their predictions (see Eqs. (7)–(9) in the Theoretical Models section for a complete explanation of how the joint fits were performed). These fits can be seen in Figs. 7 and 8, and seem to provide large distinctions between the models. Firstly, all models appear to provide a poorer account of both the response time distributions and double responses when fit jointly compared to when fit separately, suggesting that using double responses as an additional constraint provides a more difficult test for the models. Secondly, the models seem to be most clearly distinguished in their ability to account for the data by their type of inhibition implemented. The models with no inhibition (i.e., RDM, Leak) appear to provide the worst account of the data, greatly over-predicting the accuracy and later response time quantiles, as well as the double responses. The models with feed-forward inhibition, regardless of whether the feed-forward inhibition was included as a free parameter (i.e. FFI, FFI + Leak) or fixed at 1 (i.e., DM, OU), provide a better account of the data, with an improvement in fit to the response time distributions and a closer prediction of the double response proportions. In general, the free parameter feed-forward inhibition appears to capture the response time distributions slightly better than the fixed feed-forward inhibition models, though not by a large amount. Lastly, the models with lateral inhibition (i.e., LI, LI + Leak, LCA) appear to clearly provide the best account of the

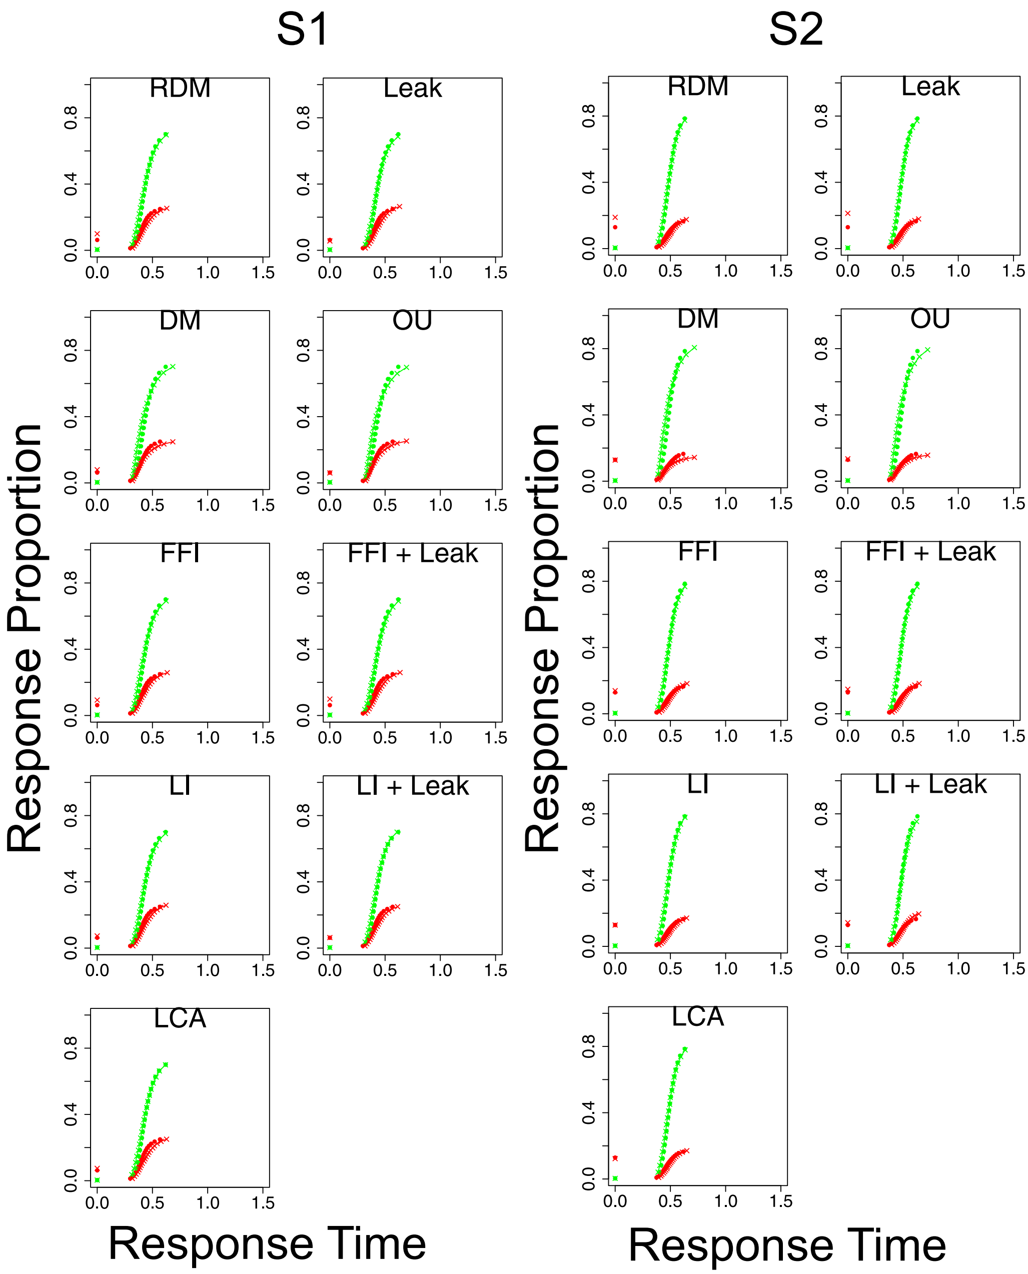

Fig. 5. CDF plots of the empirical and model predicted response time distributions and double response proportions for correct and error initial responses (see Fig. 2 for a more detailed description). Different sub-plots show the fits of different models, with the naming conventions given in Table 1. The figure contains the fits – performed separately for the response time distributions and the double response proportions (i.e., allowing different parameter values for each type of data) – for the two participants in the speed emphasis condition in Experiment 1, with the left half of the figure being one participant (S1), and the right half being the other (S2).

data, fitting the response time distributions at least as well as the feed-forward inhibition models, while also providing a better account of the double response proportions. The BIC values in Table 3 appear to provide quantitative confirmation of these qualitative differences in predictions, with the models that contain lateral inhibition strongly outperforming all other models on BIC. At an absolute level, the fit of the lateral inhibition models all appear to provide an accurate account of the data, and most importantly, provide accurate predictions for the double response proportions.

We also assessed how well the models could account for the double response time distributions, for the two participants with a large number of double responses. Fig. 9 displays the ability of the models to account for these data *without* the additional constraint of the response time distributions. All models provide an accurate account of the double response proportions, as with the fits to only the double response proportions. However, when looking at how well the models predict the double response time distributions (i.e.,

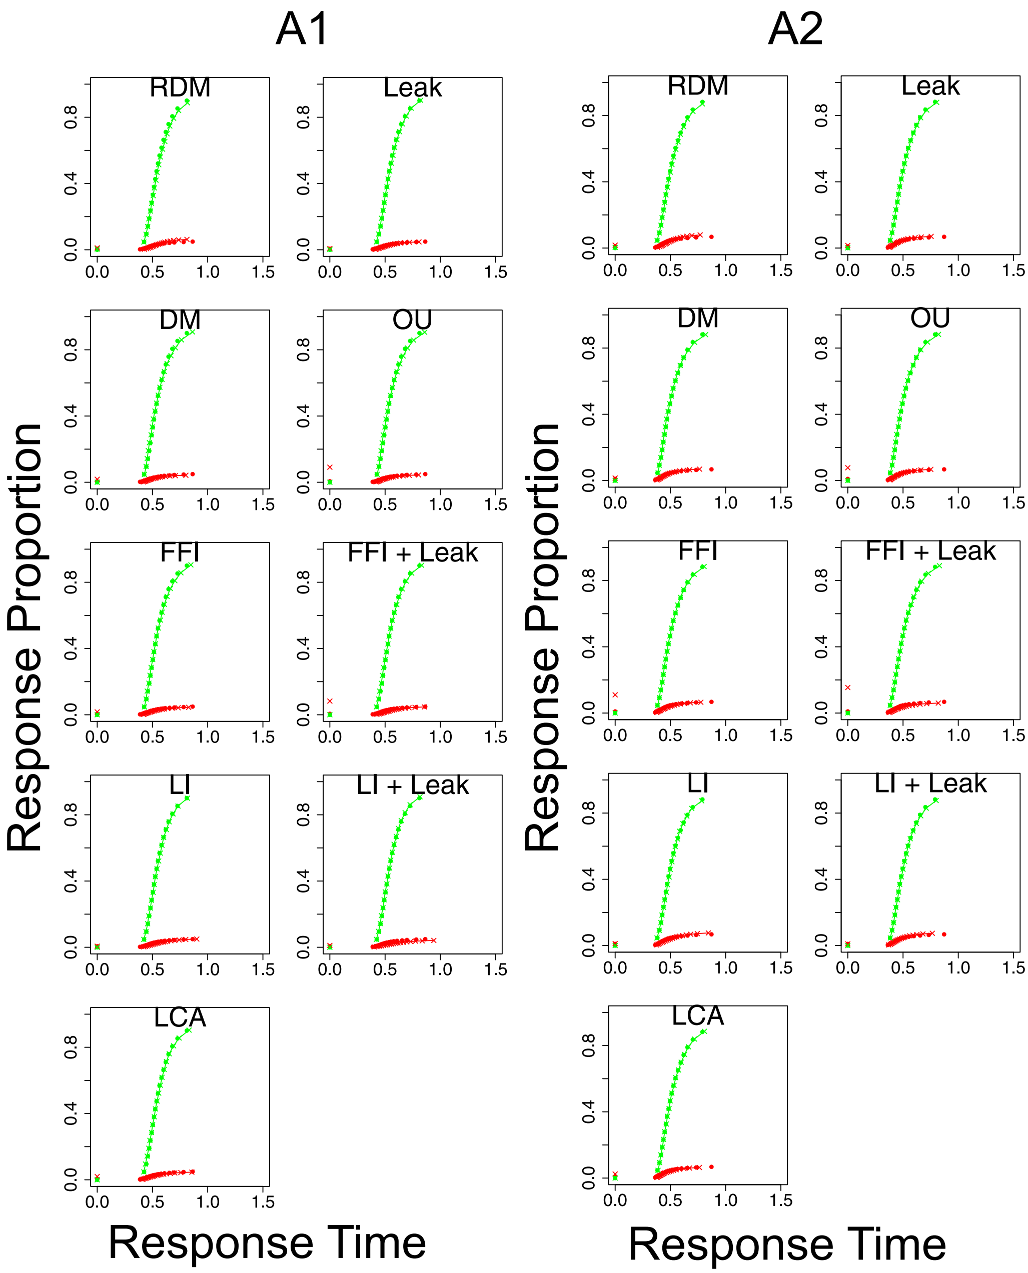

Fig. 6. CDF plots (see Fig. 2 for a more detailed description) – performed separately for the response time distributions and the double response proportions (i.e., allowing different parameter values for each type of data) – for the two participants in the accuracy emphasis condition in Experiment 1, with the left half of the figure being one participant (A1), and the right half being the other (A2).

the different quantiles), the data show a tight clustering of early quantiles, which indicates a positively skewed distribution (as shown before in Fig. 3). However, all models predict fairly evenly spaced quantiles, which indicates a distribution with equal mass throughout (i.e., a uniform distribution). Therefore, all of these models appear to have some problems with accounting for the double response time distributions. In addition, this problem only appears to worsen when fit to the choice response time and double response time distributions jointly. Fig. 10 only plots the fits to the double response time distributions from these joint fits, as the fits to the choice response time distributions remained fairly similar to those from the double response proportions. Interestingly, although the independent models still predict a fairly uniform distribution, they now greatly miss on the proportions. In addition, all other models appear to predict distributions with either a uniform distribution or with a negative skew, which is inconsistent with the empirical data (see Fig. 11 for a visual comparison between the probability density functions of the empirical data and the LCA predictions). Therefore, although the lateral inhibition models appear to provide a good account of double response proportions, even when being jointly constrained by the choice response time distributions, no model appears to be able to account for the shape of the

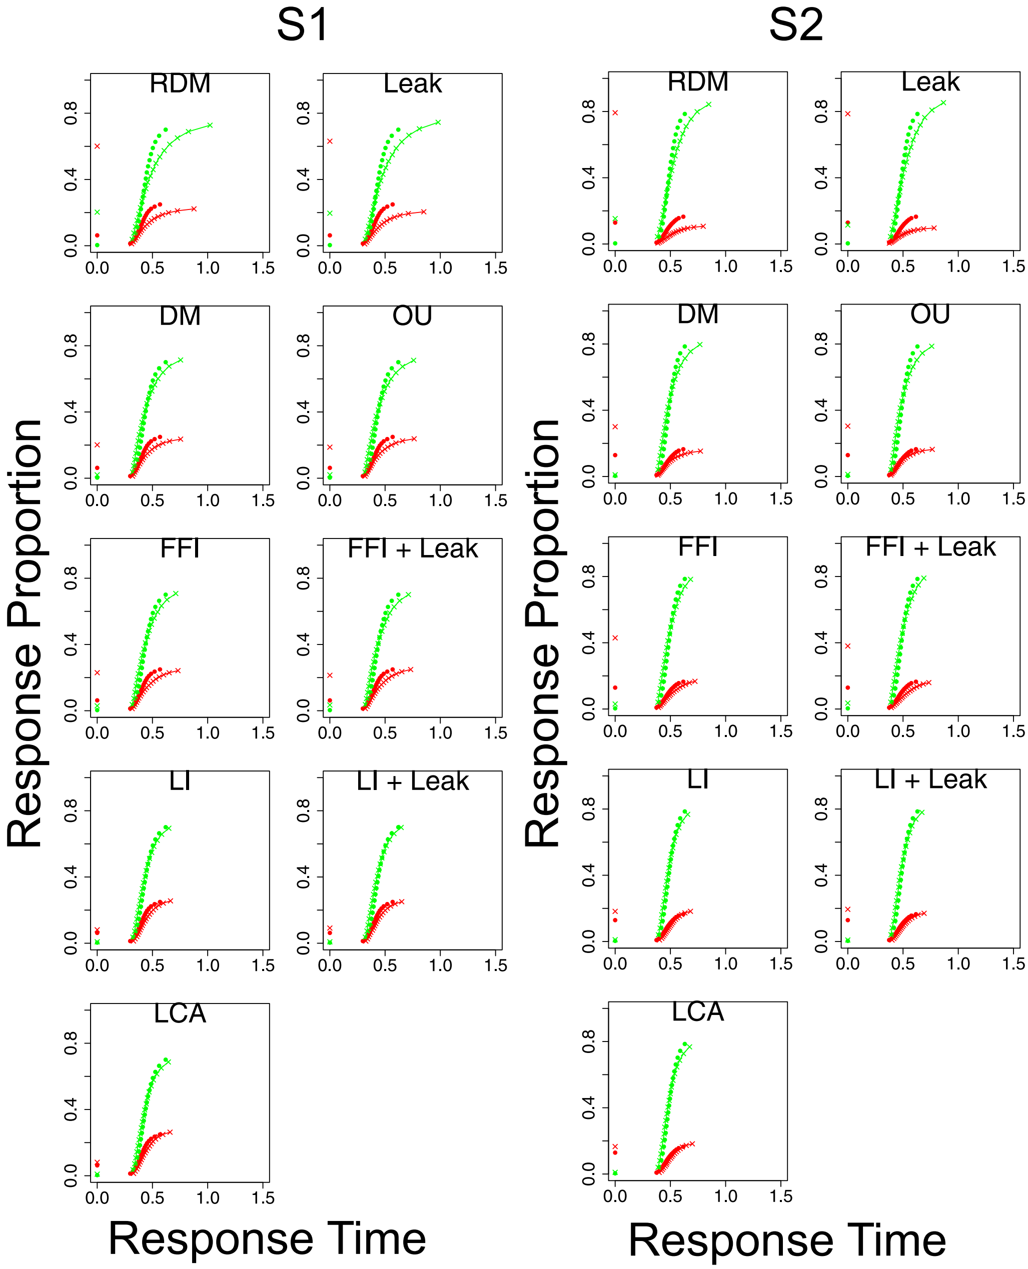

Fig. 7. CDF plots (see Fig. 2 for a more detailed description) – performed jointly for the response time distributions and the double response proportions (i.e., the same parameter values for each type of data) – for the two participants in the speed emphasis condition in Experiment 1, with the left half of the figure being one participant (S1), and the right half being the other (S2).

double response time distributions.

2.2.3. *Why does lateral inhibition help in accounting for double responding?*

In our assessment of the different theoretical models of speeded decision-making, only the models with lateral inhibition were able to provide an accurate joint account of the choice response time distributions and the double response proportions. Importantly, the choice response time distributions alone were unable to qualitatively distinguish between these models, meaning that the additional constraint provided by the double response proportions created this distinction. However, our previous analyses only provided limited insight into *why* lateral inhibition helped in accounting for the double response proportions (see Crüwell, Stefan, & Evans, 2019; Evans, 2019c; Evans & Servant, 2020 for discussions on the importance of *which* and *why* in model assessment), with the only insight being that models without lateral inhibition provided a large over-prediction for the double response proportions, particularly for trials where an error was made for the initial response.

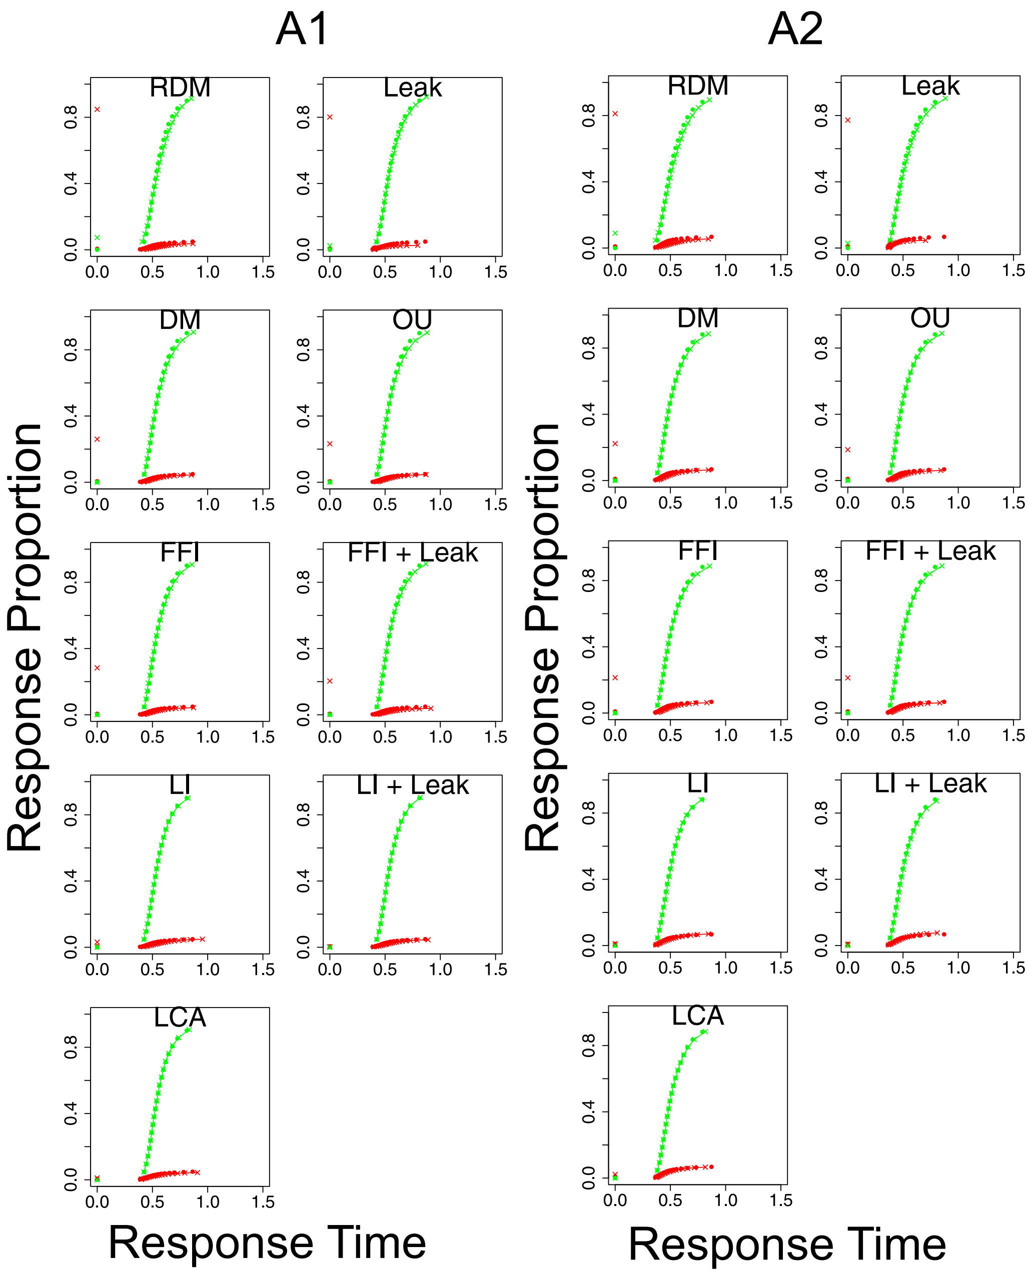

Fig. 8. CDF plots (see Fig. 2 for a more detailed description) – performed jointly for the response time distributions and the double response proportions (i.e., the same parameter values for each type of data) – for the two participants in the accuracy emphasis condition in Experiment 1, with the left half of the figure being one participant (A1), and the right half being the other (A2).

In order to provide further insight into *why* models with lateral inhibition were able to provide a superior account of double responding behaviour, we plotted the accumulation process predicted for the correct and error accumulators (green and red lines, respectively) by each model *after* an initial error, separated by whether or not a double response was made. Specifically, for each model we simulated the evidence accumulation process of 100,000 trials for 250 ms each using the best fitting set of parameters for each participant. We then separated the accumulation processes based on whether (1) the initial response was correct or an error, and (2) whether a double response was made, and averaged over accumulation processes within each category. Fig. 12 displays these averaged accumulation processes for initial error responses – using the best fitting set of parameters for participant S1, with the other participants shown in the Supplementary Materials – for three models: Leak, OU, and LCA. These models each contain a leakage component, though differ in their type of inhibition, using no inhibition, feed-forward inhibition fixed at 1, and freely estimated lateral inhibition, respectively.

As can be seen in Fig. 12, there appear to be clear differences in the accumulation trajectories after an error for each model. First

Table 3
BIC values for all 9 models (rows) for each of the 4 participants (columns) in Experiment 1. The best model for each participant is in **bold** within the table. Note that BIC is on the deviance scale, and therefore, smaller values indicate better performance.

Model	1	2	3	4
RDM	1944.67	−3747.48	−7835.96	−5530.23
Leak	2011.84	−3943.15	−9238.01	−7312.79
DM	−5221.05	−9515.55	−11286.49	−10005.30
OU	−5199.89	−9499.34	−11268.55	−10021.70
FFI	−5342.49	−10145.67	−11271.22	−10006.08
FFI + Leak	−5320.17	−10133.59	−11346.34	−10026.88
LI	−7421.07	**−12144.46**	−11816.74	−10338.38
LI + Leak	−7384.60	−12113.00	−11945.07	−10403.72
LCA	**−7445.44**	−12143.49	**−11952.44**	**−10446.06**

and foremost, OU shows a qualitatively different pattern to the other two models, which may explain its large over-prediction of double responses following errors. On both trials with and without a double response, the correct response accumulator starts a long – and similar – distance from the threshold. For the 18.3% of trials that result in a double response, the average accumulation for the correct response alternative quickly accumulates towards the threshold, reaching it at approximately 200 ms; well before the double response deadline. The 81.7% of trials without a double response show the same pattern, though with a slower average accumulation that does not reach the threshold before the double response deadline. Intuitively, one might assume that the OU (and the diffusion model) would predict a very small number of double responses due to the feed-forward inhibition during the initial response, as the correct accumulator would have to travel a long distance to reach the threshold before the double response deadline. However, the distance that the correct accumulator is required to travel to reach the threshold appears to be quite small relative to its much larger drift rate. Furthermore, feed-forward inhibition only reduces future evidence accumulation based on the accumulation rates of the other alternatives, and contains no inhibition based on previously accumulated evidence for the other alternatives. Therefore, the initially high evidence for the error alternative does little to stop the correct alternative from quickly accumulating evidence after the initial response, which also results in the error alternative quickly losing evidence due to the feed-forward inhibition.

In contrast, the pattern of accumulation between Leak and LCA appear to be much closer to one another, showing many similarities. For both models, trials with a double response have a noticeably higher amount of initial evidence for the correct response alternative than trials without a double response, suggesting that the distance between the correct and error accumulators when the initial response occurs plays a large role in whether a double response occurs. The major difference between the models in their patterns of accumulation appears to be how often the correct accumulator is close to the error accumulator after the initial response, and how quickly the correct accumulator is able to accumulate after the initial error response. For LCA, in 92.1% of trials the correct accumulator starts a reasonable distance below the error accumulator, and due to the lateral inhibition between accumulators, the accumulation for the correct accumulator is suppressed, resulting in the correct accumulator (on average) decreasing in evidence over this accumulation period. On the other 7.9% of trials the correct accumulator starts closer to the error accumulator, meaning that the lateral inhibition provides less suppression of the correct accumulator and more suppression of the error accumulator, resulting in the correct accumulator slowly accumulating evidence and moving beyond the threshold, while the error accumulator decreases in evidence over this time. In contrast, for Leak the correct accumulator is able to accumulate evidence in all cases, and does so more quickly than in the case of the LCA. This is a result of Leak not containing lateral inhibition, meaning that the initially higher accumulated evidence for the error alternative is not able to suppress accumulation for the correct alternative. In 61.7% of cases the correct accumulator starts close enough to the threshold to cross it before the double response deadline, and only starts far enough away from the threshold on 38.3% of trials to avoid crossing it. These patterns of accumulation appear to provide a fairly clear and intuitive explanation for why lateral inhibition is able to help in accurately predicting double responses, as these models are able to avoid over-predicting this trend through the suppression of evidence accumulation for the correct response alternative.

2.2.4. An alternative definition for how double responses are triggered

The primary goal of our study was to introduce double responding as a new constraint for EAMs, which required extending the standard EAM framework to incorporate double responses. Our definition in Fig. 1 attempted to provide the most natural extension of the standard EAM framework, with accumulation continuing after the initial decision, and a double response being triggered by the other accumulator hitting the threshold. However, our definition is only one of a large number of possibilities, and the choice of extension could potentially influence the conclusions regarding which model provides the best account of these data. Although we believe that it would be infeasible to implement all possible extensions within our current study, here we implement another simple definition that we believe is theoretically sensible, which can be seen in Fig. 13. This definition – which we term the "alternative definition" – is based on the potential proposal that double responding should not occur in cases where the evidence for the initial response is still higher than for the other alternative, and therefore, double responses should only occur when the initial response has dropped below the threshold. Furthermore, based on the predicted evidence accumulation plots in Fig. 12, this definition may potentially greatly help models without inhibition, as these models provide a large over-prediction of double response proportions, though in many of these cases the evidence for the initial response never drops back below the threshold. Specifically, the alternative definition is identical to the "primary definition" in every aspect except one: in the alternative definition a double response is only

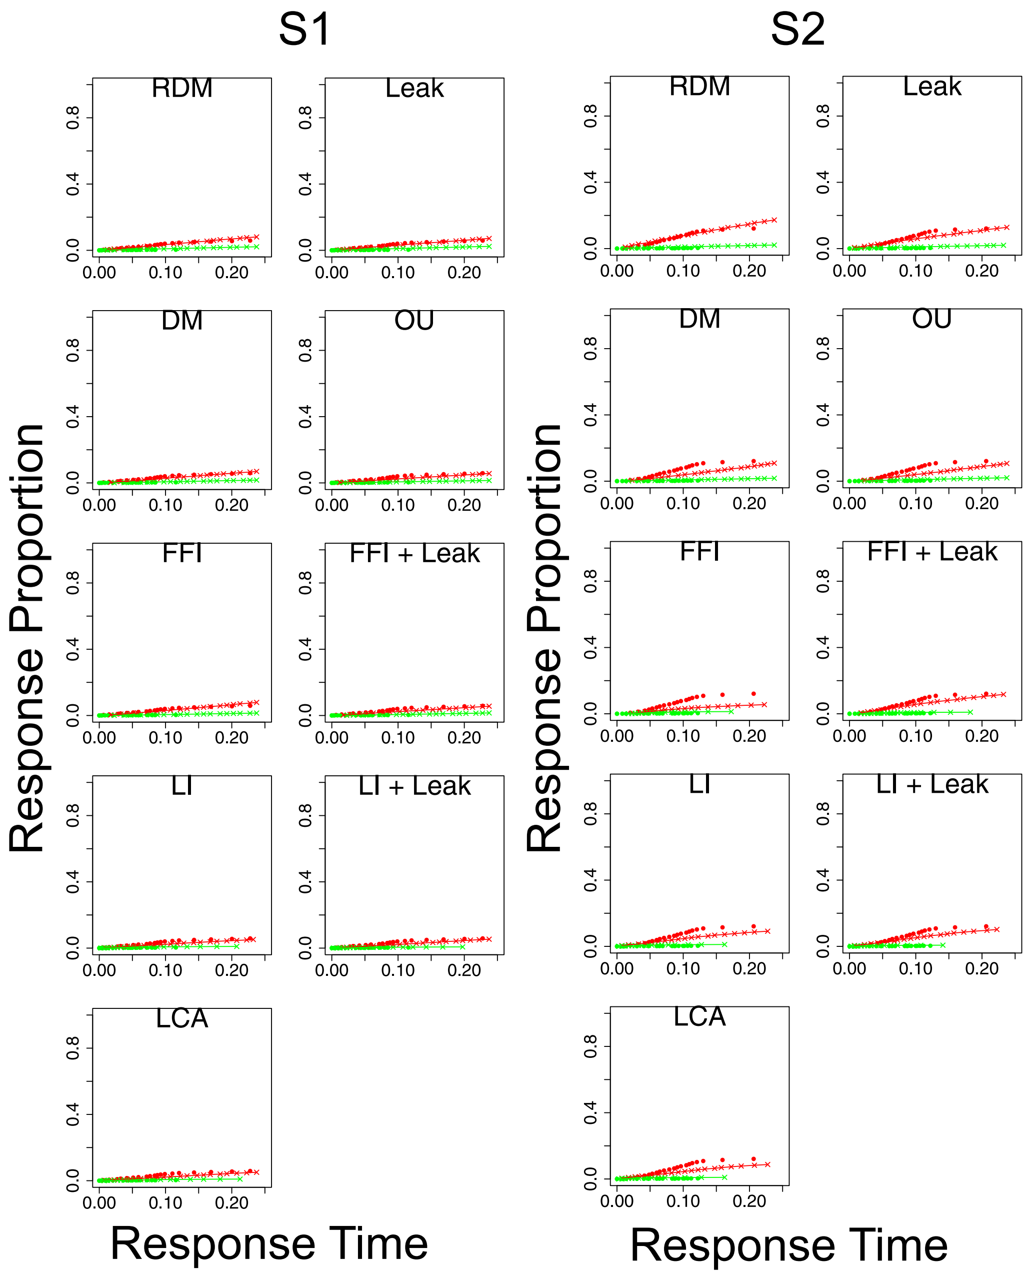

Fig. 9. CDF plots (see Fig. 2 for a more detailed description) of the empirical and model predicted double response time distributions for correct and error initial responses. The x-axis displays the double response times, and the y-axis displays the double response proportions. Different sub-plots show the fits of different models, with the naming conventions given in Table 1. The figure contains the fits – performed separately for the response time distributions and the double response time distributions (i.e., allowing different parameter values for each type of data) – for the two participants in the speed emphasis condition in Experiment 1 (left half: S1; right half: S2).

triggered if both (1) the accumulated evidence for the other alternative reaches the threshold, *and* (2) the accumulated evidence for the initial response alternative drops below the threshold. This difference can be seen in Figs. 1C and 13C, which results in a double response under the primary definition, but does not under the alternative definition.

Furthermore, our standard definitions of the different models only allowed for between-trial variability in the starting amount of evidence for each accumulator. However, two additional between-trial variability parameters are often included in EAMs, as they are able to capture certain benchmark phenomena often observed in empirical data: between-trial variability in drift rate (Brown & Heathcote, 2008; Ratcliff, 1978), and between-trial variability in non-decision time (Ratcliff & Tuerlinckx, 2002). Importantly, these between-trial variability parameters provide the models with greater flexibility, which may allow them to capture additional trends, and even potentially the pattern of double responding. Here we also implement each of the models including all three between-trial

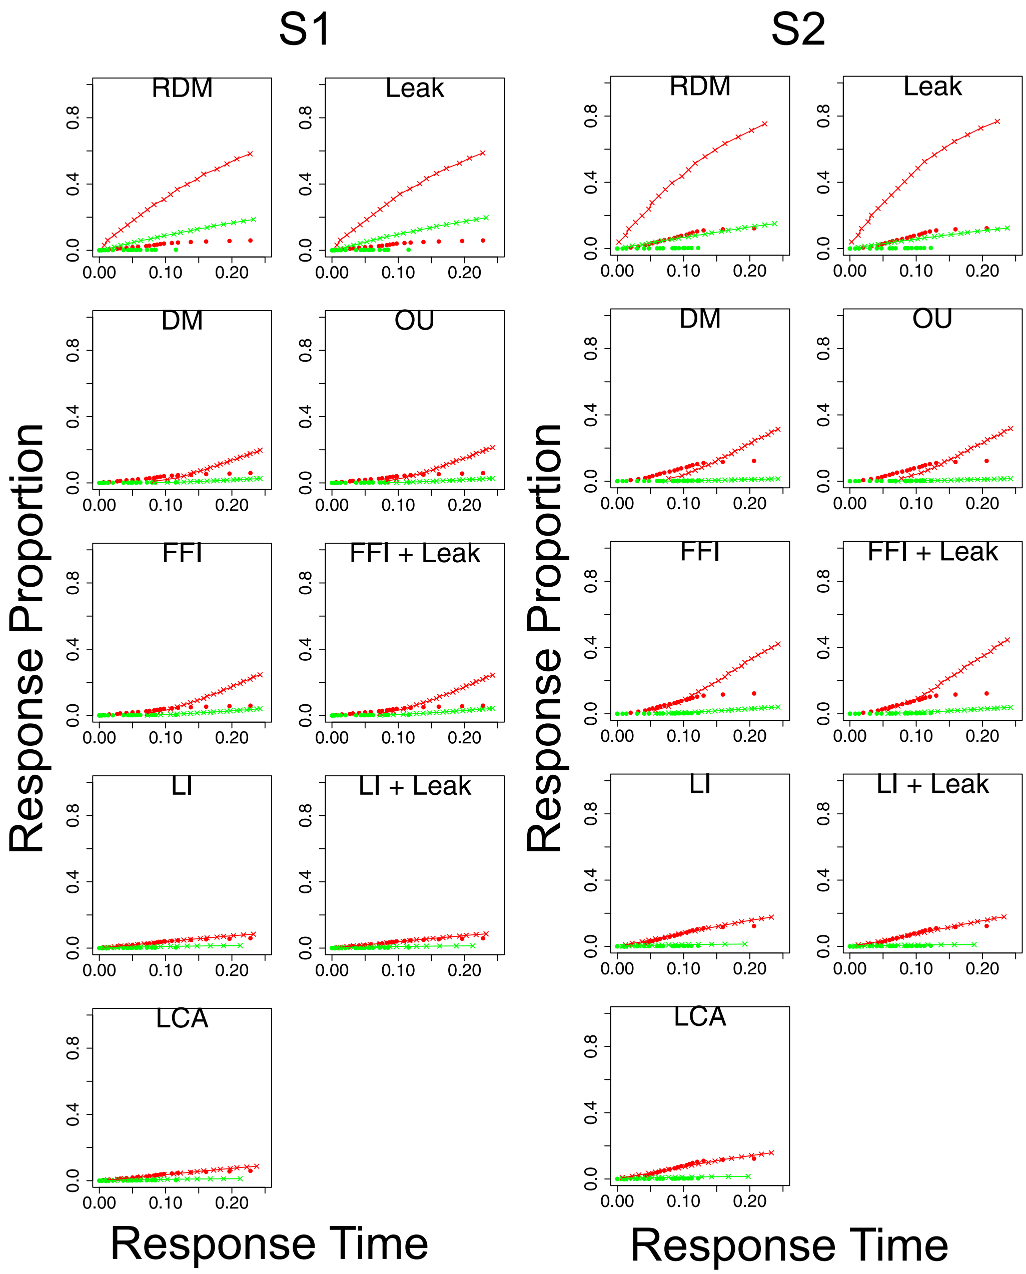

Fig. 10. CDF plots (see Fig. 2 for a more detailed description) of the empirical and model predicted double response time distributions for correct and error initial responses, based on fits that were performed jointly for the response time distributions and the double response time distributions (i.e., the same parameter values for each type of data) – for the two participants in the speed emphasis condition in Experiment 1 (left half: S1; right half: S2). Other details about the figure can be found in the caption for Fig. 9.

variability parameters.

Figs. 14 and 15 (equivalent to Figs. 7 and 8) provide the results of Experiment 1 with the alternative definition for how double responses are triggered and the inclusion of all three between-trial variability parameters. Note that we only include the fits with *both* the alternative definition *and* the between-trial variability parameters here, as changing either one of these factors separately did not result in any qualitatively different conclusions about the data, and these separate fits can be seen in the Supplementary Materials. As can be seen in Figs. 14 and 15, the combination of all three types of between trial variability parameters and the alternative definition of double responding greatly improves the ability of the models without inhibition to account for double responding, with the RDM and Leak models providing a qualitatively similar fit to the models with lateral inhibition for participants S1 and S2, which is also reflected in the RDM outperforming all other models on BIC for these participants (Table 4). The models without inhibition also provide an improved account of participants A1 and A2, though they (1) still appear to be qualitatively inferior to the models with

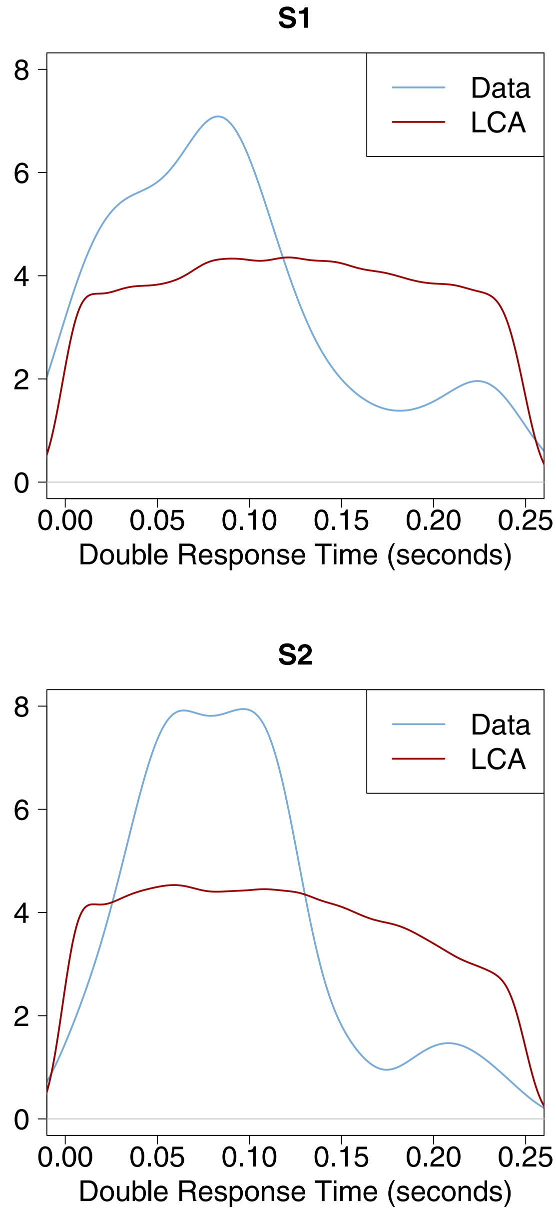

Fig. 11. Probability density function (PDF) plots of the double response time distributions after error responses for the two participants (different plots) in the speed emphasis condition of Experiment 1. Blue lines plot the empirical data, and red lines plot the predictions of the LCA when jointly fit to the response time distributions and the double response time distributions (i.e., the PDF equivalent of the CDF plots in Fig. 10 for the LCA). As can be seen, the empirical data show a positive skew in the double response time distributions, whereas the LCA predicts more uniform distributions. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

lateral inhibition, and (2) provide this improved account at the expense of being able to account for the tails of the error choice response time distributions, with both of these issues being reflected in the lateral inhibition models outperforming the models without inhibition on BIC (Table 4). Therefore, based on the participants in the accuracy emphasis condition, the models with lateral inhibition still appear to provide an overall superior account of double responding, though this is only a slight advantage when jointly considering all three between-trial variability parameters and this alternative definition of double responding. Furthermore, it should be noted that the models with forced feed-forward inhibition (i.e., the DM and OU) remain inferior to all other models, and the models with freely estimated feed-forward inhibition only show an improvement when (1) the models without inhibition also show an improvement, and (2) their estimated amount of feed-forward inhibition (i.e., FFI and FFI + Leak) is approximately zero, meaning that they essentially reduce to a model without inhibition.

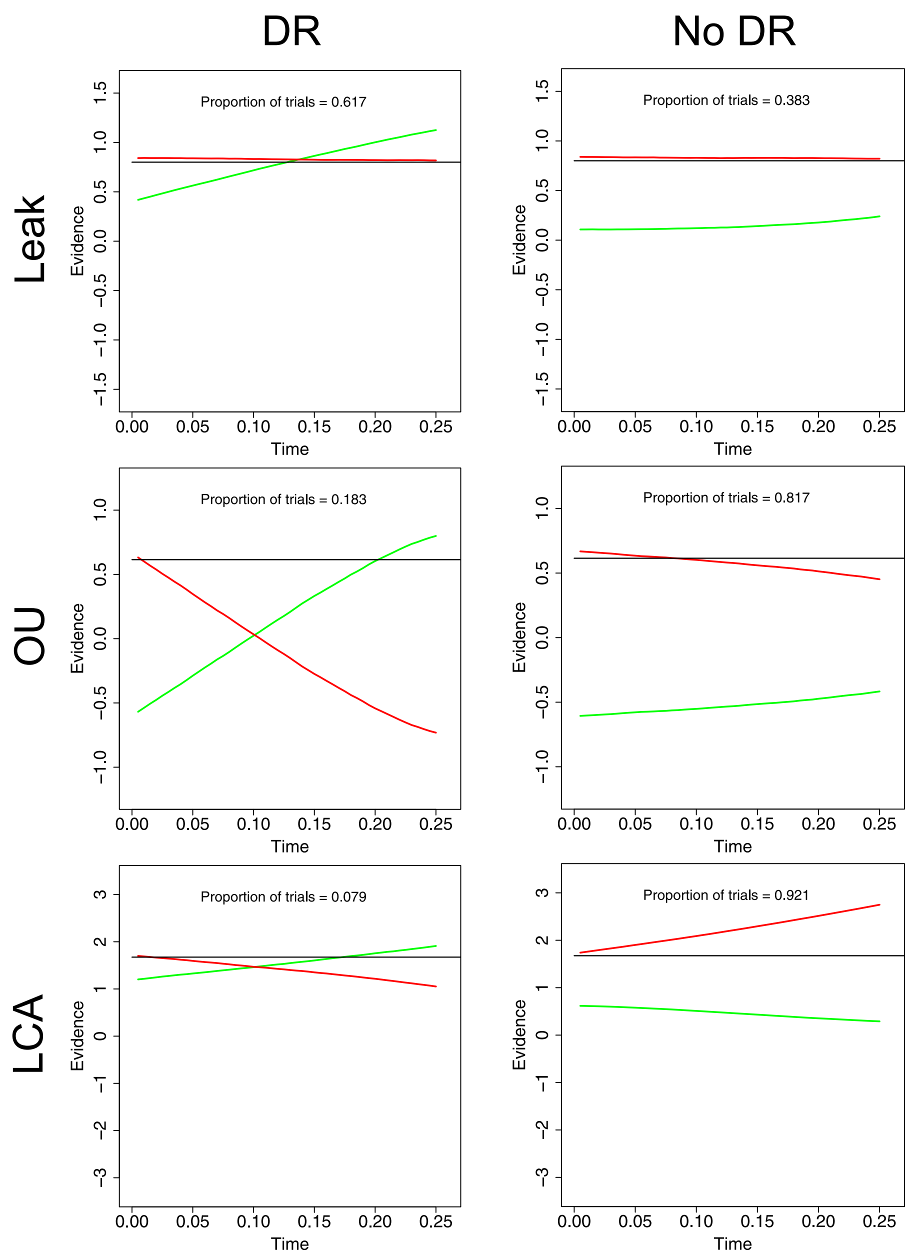

Fig. 12. Predicted evidence accumulation processes following errors for the Leak (top), OU (middle), and LCA (bottom) models based on the best fitting parameters for participant S1. The x-axis displays the time from the initial response, up until the double response time limit of 0.25 s, and the y-axis displays the accumulated evidence, starting at the finishing evidence at the initial response. The green line displays the evidence for the correct response alternative, and the red line displays the evidence for the error response alternative, with the latter being the initial response. The straight black line displays the response threshold. The left panel displays accumulation processes that result in double responses (though the accumulation continues beyond the threshold for this plot), and the right panel displays accumulation processes that do not result in double responses. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

3. Experiment 2 (Dutilh, Forstmann, Vandekerckhove, & Wagenmakers, 2013)

Our second experiment uses the data from Dutilh et al. (2013), who aimed to understand how the latent components of the decision process – via the EZ diffusion model (Wagenmakers, van der Maas, & Grasman, 2007) – were effected by errors, and how this differed between younger and older adults. Specifically, Dutilh et al. (2013) had each participant make decisions under speed and accuracy emphasis (separate sessions), using the EZ diffusion model to estimate the parameter values for decisions made after correct responses, and decisions made after error responses. As with Experiment 1, although double responding was recorded within this experiment, the focus of Dutilh et al. (2013) was purely on the response time and choice of each trial (and the diffusion model

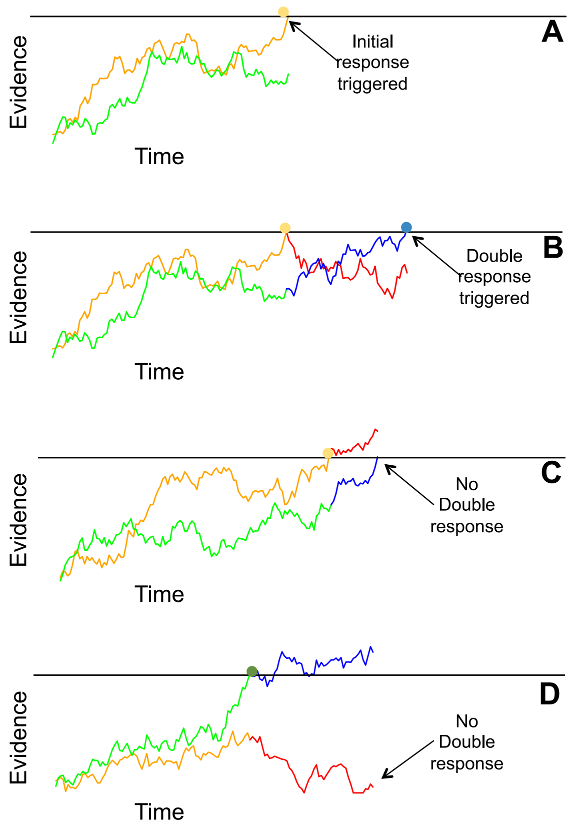

Fig. 13. Our alternative definition for how the standard EAM framework may extend to double responses. **A:** As in Fig. 1. **B:** As in Fig. 1. **C:** An example of a decision process that would **not** result in a double response, as the orange/red alternative remains above the threshold. **D:** As in Fig. 1. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

parameters estimated from them), and double responding was not mentioned within their article. We believe that this data set provides a suitable follow-up to Experiment 1 for our assessment of double responding, as it contains a within-subjects manipulation of task emphasis, which was suggested to influence double responding in Experiment 1. In addition, Dutilh et al. (2013) contains a larger number of participants (15 young adults) than Experiment 1, though at the expense of a smaller number of trials per person. Full details of the experimental setup can be found in Dutilh et al. (2013), though we provide a brief outline below. Also note that we only assess the data from the young participants within our study, as our interest is not in how double responding is influenced by aging.

3.1. Method

Fifteen participants each completed two sessions of a random dot motion task (Pilly & Seitz, 2009), where participants were presented with a cloud of moving dots during each trial and had to identify, via button press, whether the cloud of dots was generally moving towards the left or the right of screen. Each session went for 1 h and was performed under either speed or accuracy emphasis, meaning that the number of trials completed differed between participants. The range of the number of trials completed in the speed emphasis condition was 102–923, and the range in the accuracy emphasis was 22–440. Double responses were defined in the same manner as the previous experiment.

All models were extended and fit in the same manner as the previous experiment. To model the different emphasis conditions for each person, all parameters were constrained to have the same value for both emphases, apart from threshold (Dutilh et al., 2018; Lerche & Voss, 2019; Rae, Heathcote, Donkin, Averell, & Brown, 2014; Voss, Rothermund, & Voss, 2004), drift rate (Rae et al., 2014), and non-decision time (Lerche & Voss, 2019; Rae et al., 2014; Voss et al., 2004).

3.2. Results

To begin, we provide the same descriptive assessments as Experiment 1 for each participant, for each emphasis condition (Tables 5 and 6). As with Experiment 1, the emphasis manipulation appeared to be effective, with participants being faster and less accurate when under speed emphasis. Again, very few double responses were made (the most being 2.73% of trials in the speed condition), though the number of double responses was much greater in the speed condition than the accuracy condition. In addition, in the

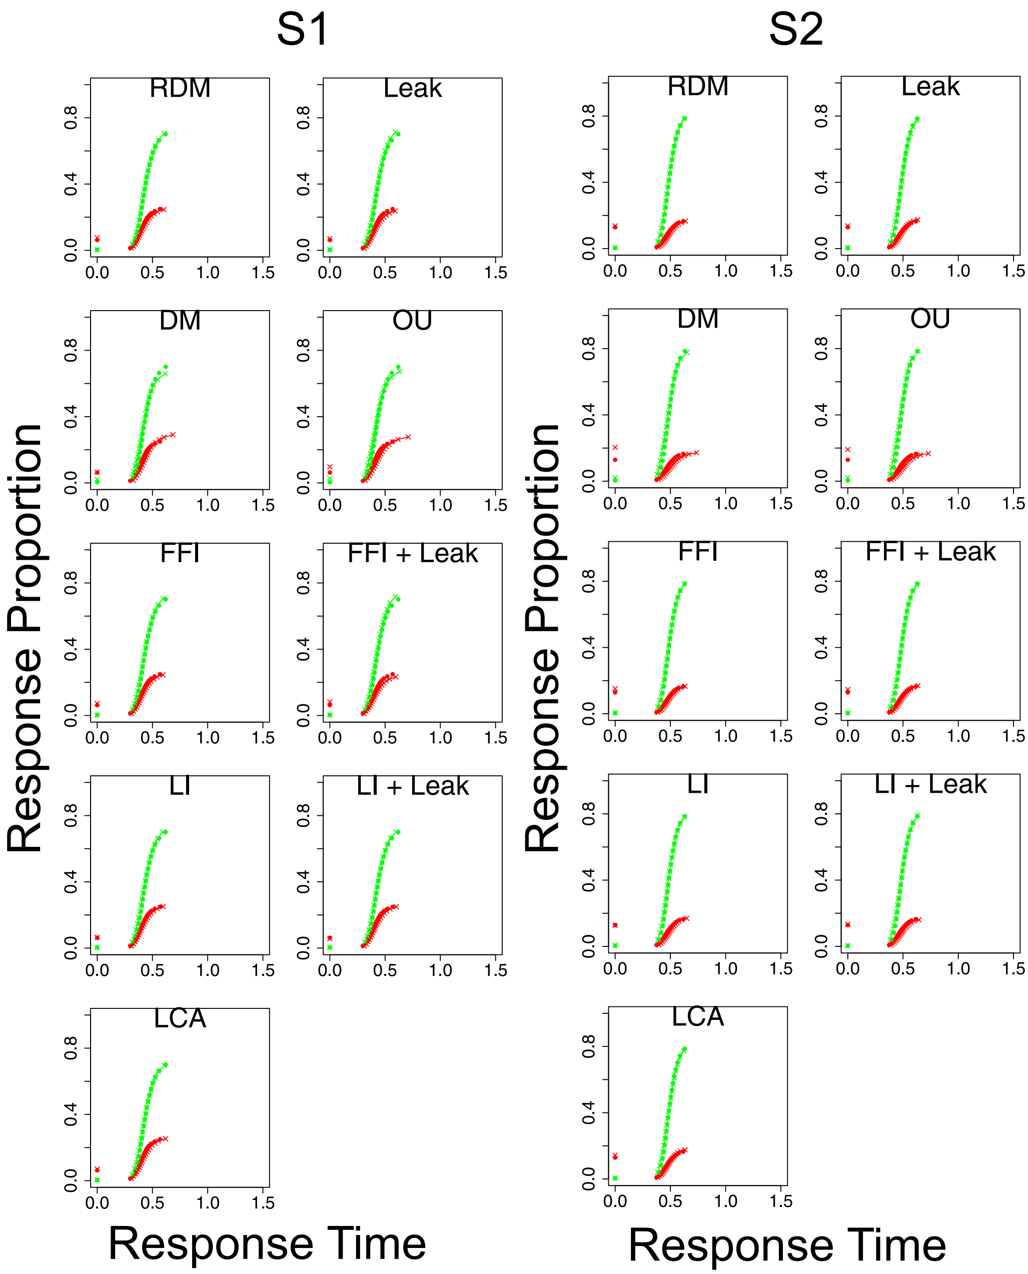

Fig. 14. CDF plots (see Fig. 2 for a more detailed description) – performed jointly for the response time distributions and the double response proportions (i.e., the same parameter values for each type of data) – for the two participants in the speed emphasis condition in Experiment 1, with the left half of the figure being one participant (S1), and the right half being the other (S2). These fits were performed using the alternative definition of double responding and all three between-trial variability parameters.

speed condition double responses occurred much more frequently after error responses than correct responses, again suggesting that these double responses are mostly "corrective" (i.e., following error responses).

We fit the models jointly to the choice response time distributions and the double response proportions, which can be seen in Fig. 16. As with Experiment 1, the joint fits appear to clearly distinguish between the models, and provide similar results to Experiment 1. Once again, the type of inhibition in the model appears to be the decisive factor, with independent models vastly over-predicting the number of double responses, feed-forward inhibition models improving upon the independent models, and lateral inhibition models providing the best account of double responses. These findings are also supported by the BIC analysis (Table 7), with lateral inhibition models outperforming all other models for all participants. There also appears to be a tradeoff within the feed-forward inhibition models in accounting for the double responses and the response time distributions; the more-constrained diffusion model provides a closer account of double responding than the less constrained feed-forward inhibition model, though this is at the

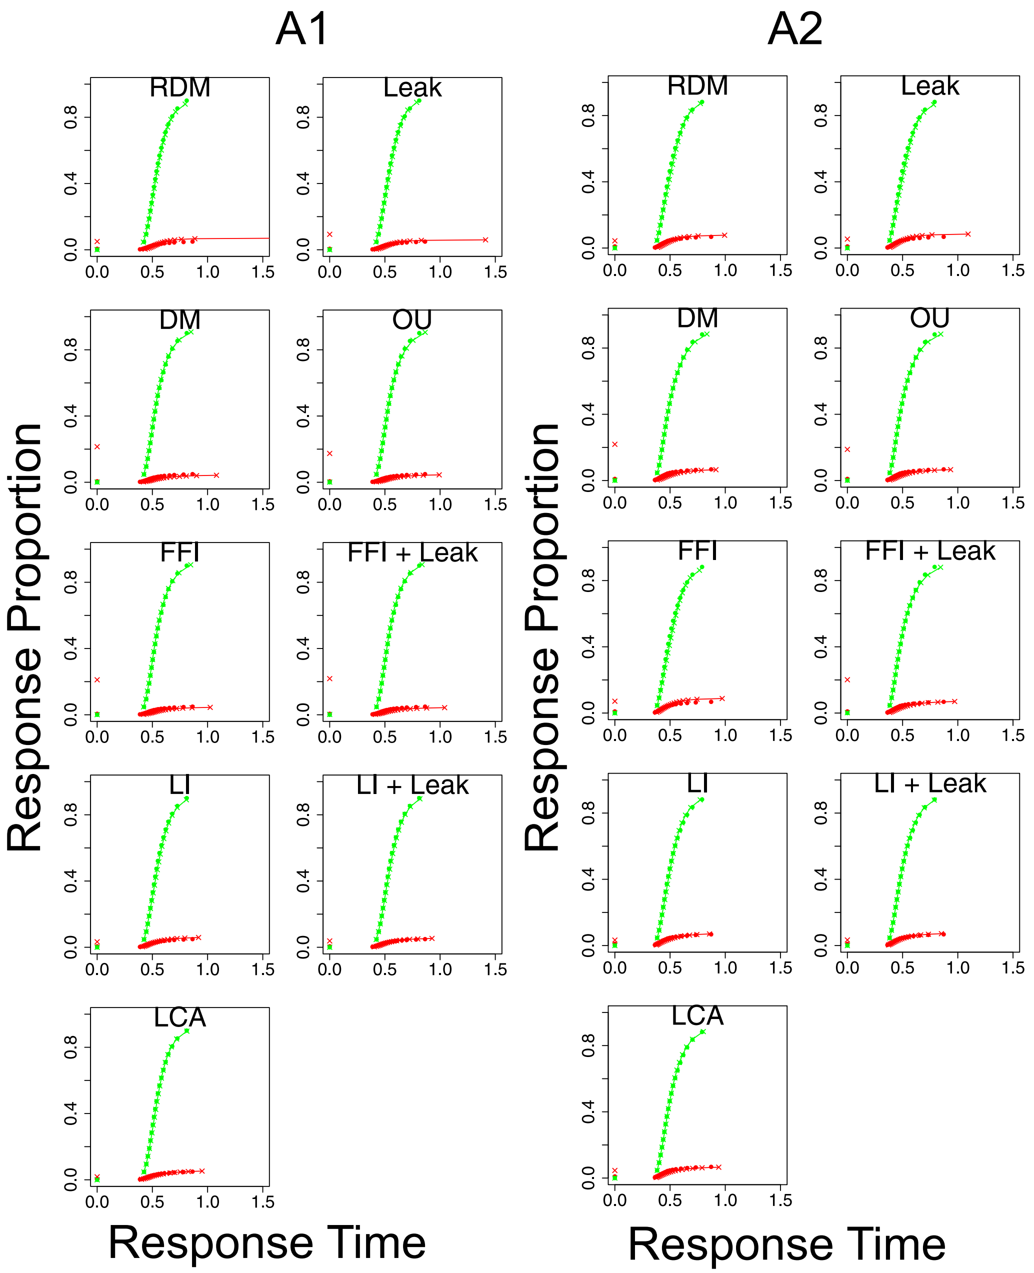

Fig. 15. CDF plots (see Fig. 2 for a more detailed description) – performed jointly for the response time distributions and the double response proportions (i.e., the same parameter values for each type of data) – for the two participants in the speed emphasis condition in Experiment 1, with the left half of the figure being one participant (A1), and the right half being the other (A2). These fits were performed using the alternative definition of double responding and all three between-trial variability parameters.

expense of providing a poorer fit to the response time distributions. The independent models also provide a poor fit to both the double responses and the response time distributions, suggesting that to capture the response time distributions properly the independent models would have had to provide an even poorer account of double responding. These findings place a further emphasis on the importance of fitting the response time distributions and double responses jointly, as models can provide a poor account of the standard response time and choice data to more closely account for the double responses.

However, as in Experiment 1, when considering *both* the alternative definition for how double responses are triggered and all three between-trial variability parameters, the models without inhibition provide much more accurate predictions about these data. These results can be seen in Fig. 17, with the corresponding BIC analysis in Table 8. As can be seen, both the models without inhibition and the models with lateral inhibition provide an accurate account of the choice response time distributions, with only a slight – and similar – over-prediction for the proportion of double responses in each condition. However, unlike in Experiment 1

Table 4
BIC values for all 9 models (rows) for each of the 4 participants (columns) in Experiment 1, using the alternative definition of double responding and all three between-trial variability parameters. The best model for each participant is in **bold** within the table. Note that BIC is on the deviance scale, and therefore, smaller values indicate better performance.

Model	1	2	3	4
RDM	**−8409.98**	**−13259.74**	−11610.48	−10099.26
Leak	−8397.24	−13257.20	−11627.22	−10107.05
DM	−7304.49	−12162.50	−11685.12	−10071.10
OU	−7303.49	−12134.47	−11695.08	−10095.90
FFI	−8408.66	−13255.61	−11678.24	−10110.17
FFI + Leak	−8368.34	−13222.14	−11715.04	−10058.60
LI	−8407.78	−13212.73	−11885.96	−10370.22
LI + Leak	−8387.80	−13240.42	−11955.85	**−10459.55**
LCA	−8365.80	−13237.99	**−11983.08**	−10440.10

Table 5
Summary statistics for the participants in Experiment 2 for the speed emphasis condition. Different columns provide different pieces of information, and different rows provide the information for different subjects. Abbreviations can be found in Table 2.

Sub	MRT	P(C)	P(DR)	P(DR|C)	P(DR|E)	P(C|DR)	N_{trials}
1	0.4297	0.9295	0.005	0	0.0714	0	397
2	0.4609	0.8326	0.0088	0	0.0526	0	227
3	0.4328	0.8421	0.0115	0	0.0729	0	608
4	0.392	0.6308	0.0063	0	0.0171	0	474
5	0.4558	0.6111	0.0033	0.0018	0.0056	0.3333	923
6	0.4648	0.9496	0.0168	0	0.3333	0	119
7	0.3693	0.8255	0.0169	0	0.0968	0	533
8	0.4286	0.8786	0.0022	0	0.0182	0	453
9	0.3957	0.7244	0.0079	0	0.0286	0	762
10	0.4828	1	0	0	NA	NA	102
11	0.4051	0.891	0.0062	0	0.0571	0	321
12	0.4094	0.9057	0.0101	0	0.1071	0	297
13	0.4665	0.9447	0.0158	0	0.2857	0	253
14	0.4703	0.8371	0.0181	0	0.1111	0	221
15	0.4155	0.8484	0.0273	0	0.1798	0	587
Average	0.4319	0.8434	0.0104	0.0001	0.1027	0.0238	418

Table 6
Summary statistics for the participants in Experiment 2 for the accuracy emphasis condition. Different columns provide different pieces of information, and different rows provide the information for different subjects. Abbreviations can be found in Table 2.

Sub	MRT	P(C)	P(DR)	P(DR|C)	P(DR|E)	P(C|DR)	N_{trials}
1	0.4437	0.9786	0	0	0	NA	187
2	0.4718	0.9433	0	0	0	NA	141
3	0.4793	0.9653	0	0	0	NA	259
4	0.5221	0.9423	0	0	0	NA	104
5	0.6147	0.8068	0.0023	0	0.0118	0	440
6	0.5092	1	0	0	NA	NA	22
7	0.4259	0.9371	0.0063	0	0.1	0	159
8	0.4764	0.9837	0	0	0	NA	123
9	0.4821	0.95	0	0	0	NA	180
10	0.5104	0.9583	0	0	0	NA	48
11	0.4171	0.9639	0	0	0	NA	194
12	0.4234	0.9327	0.0048	0	0.0714	0	208
13	0.4922	0.9427	0.0088	0	0.1538	0	227
14	0.4834	0.9914	0	0	0	NA	116
15	0.4411	0.9839	0.004	0	0.25	0	249
Average	0.4795	0.952	0.0017	0	0.0419	0	177

where the models with lateral inhibition were still superior for the participants in the accuracy emphasis condition, both models appear to have nearly identical accuracy in their predictions. This is also supported by the BIC analysis, where models without inhibition outperform other models for 8/15 participants, models with lateral inhibition outperform other models for 7/15 participants, and for all participants the BICs between these classes of models are quite similar. Similar to Experiment 1, models with forced feed-forward inhibition (i.e., the DM and OU) remain inferior to all other models, providing (1) a poorer account of the choice response time distributions, (2) a greater over-prediction of the double response proportions, and (3) never outperform the other

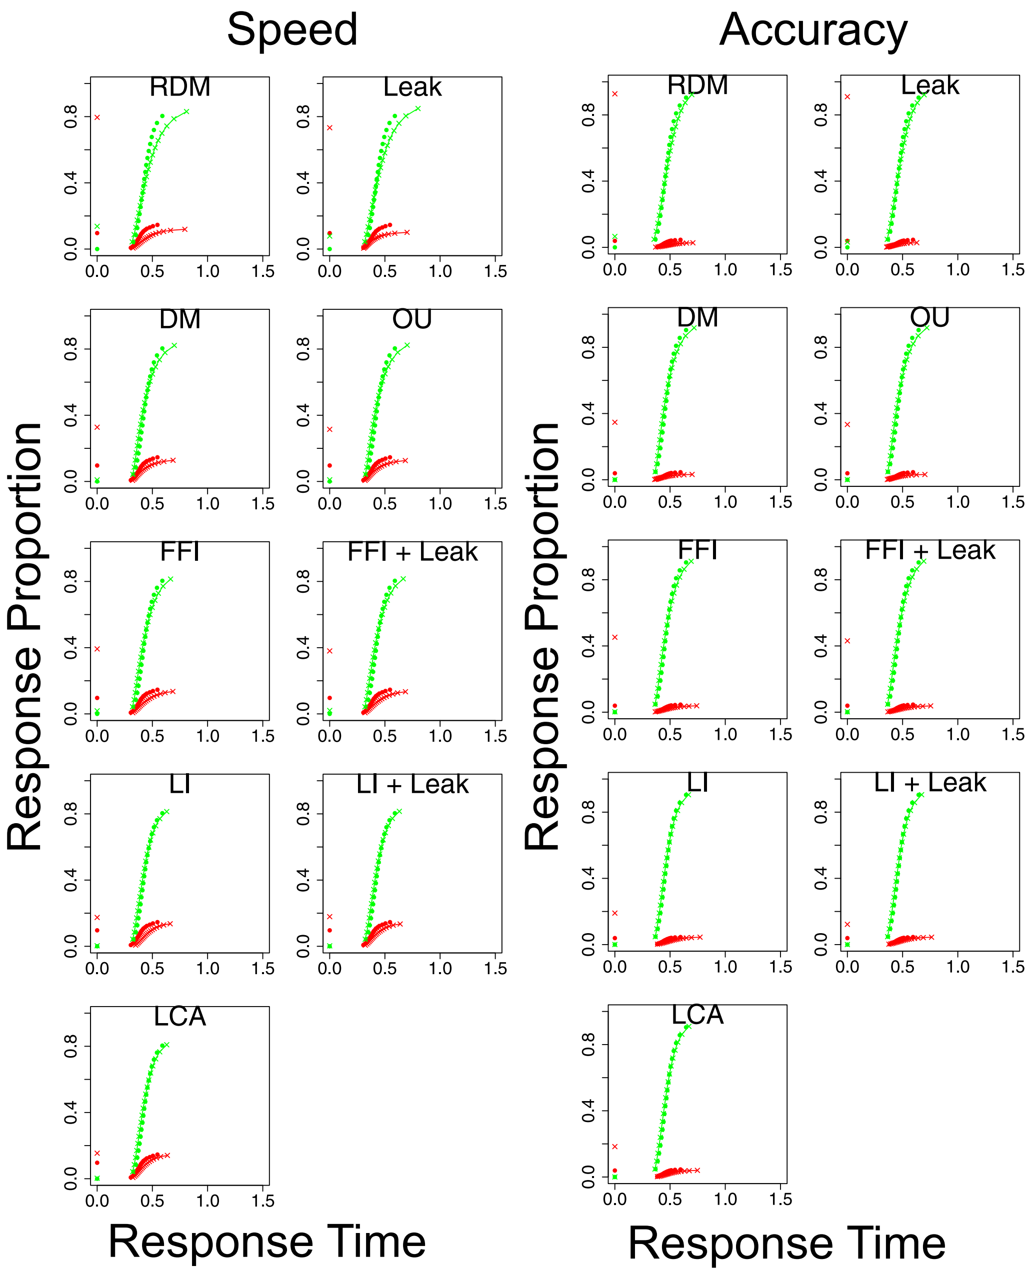

Fig. 16. CDF plots (see Fig. 2 for a more detailed description) – performed jointly for the response time distributions and the double response proportions (i.e., the same parameter values for each type of data), for the speed (left) and accuracy (right) conditions of Experiment 2. The data of each participant was fit separately to obtain the observed and predicted response time quantiles and response proportions, with these quantiles and proportions then averaged over participants to create these group-averaged CDF plots.

models on BIC for any participant.

4. Discussion

Our study proposed double responding as a new constraint for models of decision-making, which allows for a better understanding of the dynamics of the complete decision-making process by separating the predictions of different theoretical models of decision-making. Previous studies within the decision making literature have mostly focused on response time and choice, as these two variables are able to provide insights into the decision making process using computational models (Brown & Heathcote, 2008; Ratcliff, 1978; Ratcliff & Rouder, 1998; Usher & McClelland, 2001). Although other phenomena have also been explored in some studies (e.g., choice confidence; Pleskac & Busemeyer, 2010; Ratcliff & Starns, 2009, 2013), measuring these phenomena often

Table 7
BIC values for all 9 models (columns) for each of the 15 participants (rows) in Experiment 2. The best model for each participant is in **bold** within the table. Note that BIC is on the deviance scale, and therefore, smaller values indicate better performance.

Sub	RDM	Leak	DM	OU	FFI	FFI + Leak	LI	LI + Leak	LCA
1	−748	−826	−953	−946	−967	−964	**−1030**	−1024	−1025
2	−127	−165	−297	−291	−297	−291	**−348**	−343	−343
3	−505	−567	−915	−902	−933	−925	**−1061**	−1058	−1055
4	124	132	−291	−285	−290	−285	**−430**	−422	−430
5	1158	1171	215	223	139	152	**−71**	−60	−70
6	−207	−204	−223	−218	−223	−216	**−236**	−229	−229
7	−599	−703	−1120	−1113	−1111	−1107	**−1274**	−1268	−1270
8	−503	−546	−788	−781	−802	−797	−902	−897	**−903**
9	−44	−88	−653	−642	−649	−642	**−855**	−852	−852
10	−227	−225	−230	−225	−231	−225	**−233**	−230	−230
11	−575	−701	−886	−882	−882	−878	**−959**	−957	−951
12	−647	−797	−981	−978	−974	−973	**−1051**	−1044	−1041
13	−578	−661	−763	−757	−760	−756	**−790**	−787	−784
14	−172	−195	−313	−311	−314	−310	**−348**	−341	−334
15	−829	−915	−1205	−1197	−1233	−1219	**−1349**	−1339	−1343

require an alteration of the experimental paradigm, and prevents their assessment in "standard" decision making paradigms. Here, we suggest that double responding may be an interesting phenomenon for researchers to attempt to explain in addition to response times and choice, which can be collected as part of any standard experimental paradigm. We found that double responses are quite rare, mostly appear to be corrective (i.e., correcting previous errors; see Rabbitt & Rodgers, 1977 for more details about "error correcting responses"), and that they occur proportionally more after errors when participants were performing an experiment under speed emphasis than accuracy emphasis. When assessing the shape of the double response time distributions in the two participants who provided a sizable number of double responses, we found the distributions to display the positive skew found in standard response time distributions. These three observations could serve as benchmarks for theories that wish to provide a more comprehensive explanation of decision making by integrating double responses into their framework.

Importantly, we showed that double responding can be used to further constrain different models of decision making, and that lateral inhibition is required to explain the complete decision making process. As most studies within the decision making literature only focus on response time and choice, the prominent models of decision making – evidence accumulation models (EAMs; Laming, 1968; Ratcliff, 1978; Stone, 1960) – have been designed to account for these response time and choice phenomena, resulting in the models being difficult to distinguish in standard paradigms (though see Teodorescu & Usher, 2013 for an example of an experimental test between EAMs). We provided a natural extension of 9 variants of EAMs to double responding, where evidence continued to accumulate for the alternatives after the initial decision had been made, and a double response was triggered if the opposing alternative reached the threshold. Under our primary definition for how double responses are triggered within the EAM framework, where any instance of the accumulated evidence for the other response alternative reaching the decision threshold triggered a double response, models with lateral inhibition – where the future evidence accumulation for an alternative is reduced based on the previously accumulated evidence for the other alternatives – provided the best account of double responding, suggesting that they occur quite rarely, whereas all model variants predicted that they occurred in large quantities. However, when simultaneously considering an alternative definition – where a double response can only be triggered if the evidence for the initial response alternative has also dropped below the threshold – and between trial variability in drift rate, threshold, and non-decision time, models without inhibition were also able to correctly predict the small quantities of double responses. Therefore, although models with lateral inhibition provided accurate predictions under all definitions, models without inhibition showed that they can also capture these trends under certain circumstances. Furthermore, the models with feed-forward inhibition fixed at 1 – such as the popular diffusion model (Ratcliff, 1978) – were unable to provide an accurate account of these data under any definition. Importantly, these distinctions could only be made when the models were forced to account for response time, response choice, and double responses simultaneously; when fitting only the double responses, all models provided a good account of the data. This places further emphasis on the importance of constraining models of additional sources of data, rather than just different sources of data, as important distinctions can often be made based on their joint predictions.

However, we wish to note that the aim of our study was only to distinguish between and compare different theories of speeded decision-making, and that our study does not have implications for the measurement properties of EAMs. Importantly, one of the most common uses of EAMs is as "measurement tools" of the decision-making process, where researchers use EAMs to estimate the latent parameters of the decision-making process – such as drift rate and threshold – and make inferences about how they vary across experimental conditions and/or groups (Crüwell et al., 2019). Having measurement tools that are theoretically accurate is important, as inaccuracies in the theoretical underpinning of the measurement tools may make the parameters estimated from them meaningless. However, we believe that the importance of having reliable measurement tools cannot be overstated, as having unreliable measurement tools (e.g., models with identifiable parameters) will guarantee inaccurate estimates and conclusions. Therefore, although our study suggested that models with lateral inhibition may be theoretically superior to models with feed-forward inhibition for 2AFC tasks, models with lateral inhibition may be a questionable choice for use as measurement tools, as previous research has

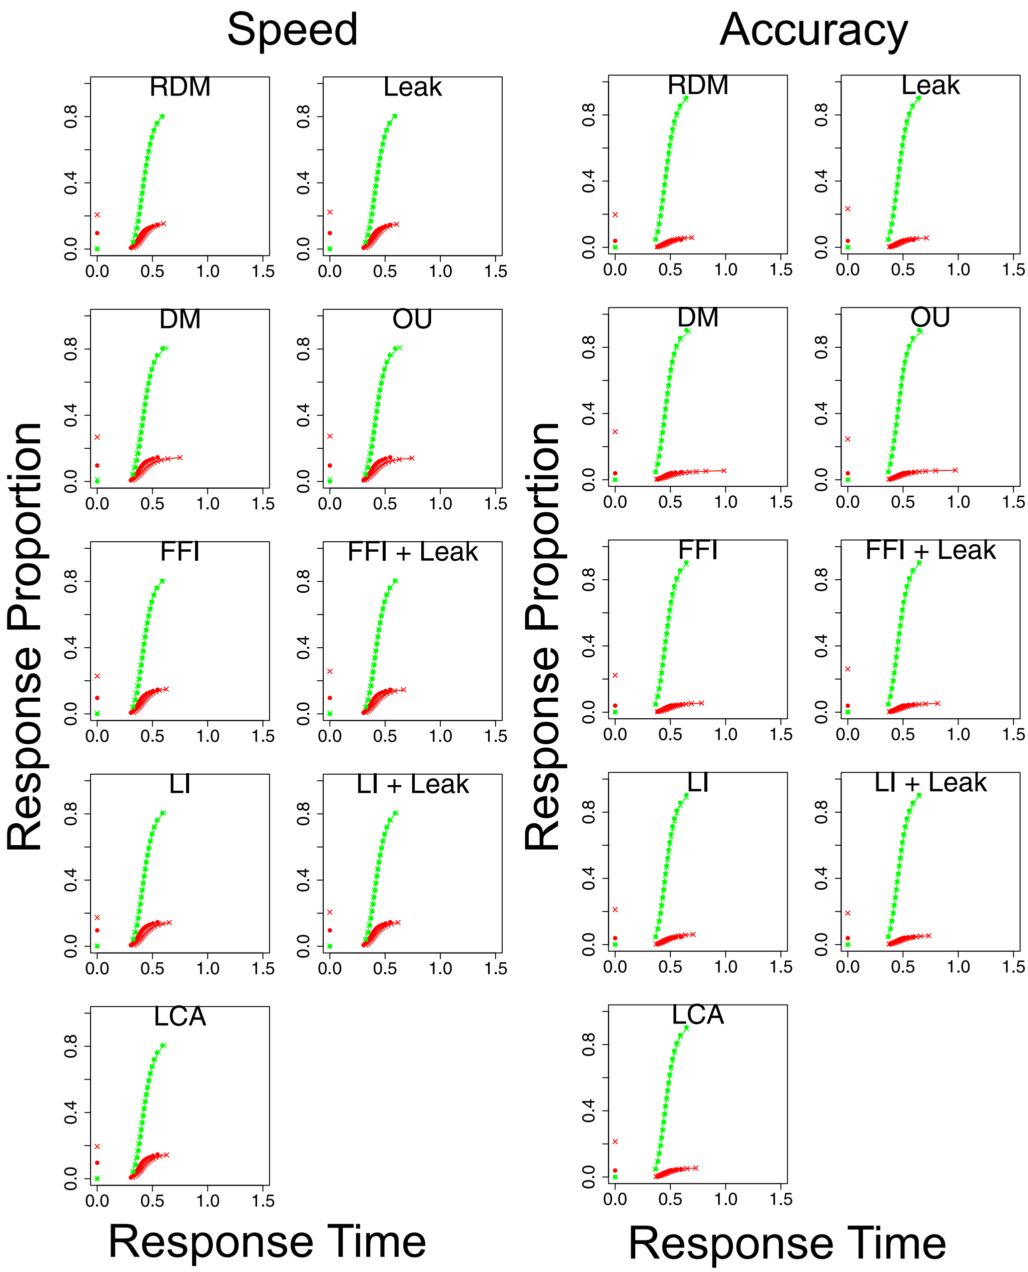

Fig. 17. CDF plots (see Fig. 2 for a more detailed description) – performed jointly for the response time distributions and the double response proportions (i.e., the same parameter values for each type of data), for the speed (left) and accuracy (right) conditions of Experiment 2. The data of each participant was fit separately to obtain the observed and predicted response time quantiles and response proportions, with these quantiles and proportions then averaged over participants to create these group-averaged CDF plots. These fits were performed using the alternative definition of double responding and all three between-trial variability parameters.

suggested that these models contain identifiability problems (Miletić, Turner, Forstmann, & van Maanen, 2017). Therefore, we again wish to reiterate that researchers should not use the results of our study to guide which EAMs they choose to use as measurement tools. However, we believe that an interesting avenue for future research would be to assess whether double responses – or other additional sources of data – are able to provide adequate constraint to allow models with lateral inhibition to have identifiable parameters, and whether including double responses as a constraint in models that are already identifiable leads to changes in parameter estimates.

Table 8
BIC values for all 9 models (columns) for each of the 15 participants (rows) in Experiment 2, using the alternative definition of double responding and all three between-trial variability parameters. The best model for each participant is in **bold** within the table. Note that BIC is on the deviance scale, and therefore, smaller values indicate better performance.

	RDM	Leak	DM	OU	FFI	FFI + Leak	LI	LI + Leak	LCA
1	−1043	−1035	−1026	−1022	−1037	−1023	**−1050**	−1048	−1041
2	−353	−347	−322	−318	−344	−335	**−355**	−349	−346
3	−1124	−1120	−1038	−1035	−1115	−1106	**−1126**	−1123	−1122
4	**−480**	−469	−412	−404	−473	−452	−471	−464	−460
5	−159	−156	−95	−88	−153	−139	−155	−157	**−159**
6	**−253**	−247	−230	−225	−247	−241	−247	−242	−241
7	**−1326**	−1320	−1228	−1227	−1316	−1307	−1319	−1309	−1310
8	**−943**	−940	−884	−880	−936	−923	−937	−929	−929
9	−925	**−927**	−777	−772	−905	−903	−917	−911	−910
10	**−235**	−229	−232	−227	−229	−223	−229	−224	−224
11	−963	−956	−959	−955	−951	−951	**−966**	−964	−948
12	**−1047**	−1040	−1010	−1007	−1040	−1025	−1038	−1029	−1031
13	−771	−767	−781	−779	−774	−773	**−785**	−778	−779
14	−354	−351	−359	−355	−354	−349	**−361**	−356	−351
15	**−1427**	−1421	−1341	−1341	−1416	−1405	−1420	−1414	−1413

4.1. Implicit or explicit double responding?

As briefly discussed at the end of our introduction, our study assessed implicit double responding behaviour, where participants are *not* directly instructed to make a second response. The observed double responses from an implicit paradigm can be thought of as a spillover of information from the decision-making process, which results in a second motor output for the other alternative. However, an alternate way to assess double responding would be through explicit double responding behaviour, where participants *are* directly instructed that they are allowed to make a second response. The observed double responses from an explicit paradigm can be thought of as a combination of the previously outlined spillover of information from the decision-making process – as these implicit double responses would still occur under direct instruction – and intentional corrections of errors, where participants make a second response knowing that the task allows for them to correct their errors. Importantly, when studying double responding, researchers must choose whether to assess implicit or explicit double responses, which may provide very different measurements of double responding behaviour. Here, we provide a discussion of the advantages and disadvantages of assessing each type of double responding behaviour, as well as how our choice of implicit double responding over explicit double responding may have impacted our results.

Implicit double responding provides an easy-to-implement measurement of double responding behaviour. Researchers do not need to change their existing speeded decision-making tasks, as the only specific requirement is that researchers record any responses that participants make between the initial response and the following trial. Importantly, implicit double responding ensures that participants do not alter their behaviour from a standard speeded decision-making tasks, as they are unaware that their second responses are being recorded, and therefore, cannot adjust their decision strategy accordingly. This equivalence provides a clear generalization for the findings of implicit double responding assessments to standard speeded decision-making tasks, as participants should be performing both tasks in the same manner. However, the key limitation of implicit double responding is that participants may intentionally avoid making double responses. For example, as participants are not aware that their second responses are being recorded, they may directly attempt to suppress their natural instinct to make a second response in some situations, as making a second response may impair their ability to prepare for the next trial. This would lead to an underestimation of implicit double responding behaviour, as only information spillovers that participants fail to inhibit would result in a double response being observed.

Explicit double responding provides a clear solution to the key limitation of implicit double responding. As participants are instead directly instructed that they are allowed to make a second response whenever they believe that their initial response was incorrect, participants would be unlikely to attempt to suppress second responses, as the suppression of these responses would result in a poorer performance in the task. Explicit double responding behaviour would consist of both implicit double responses, as these would not be suppressed, and purposeful corrections of perceived errors, in accordance with the instructions of the task. However, the key limitation of explicit double responding is that informing participants that they are able to correct their initial errors may alter how they complete the task. For example, as participants are aware that their initial response does not need to be their final response, they may adopt a less cautious strategy in their initial responses than they would normally, as they believe that they will be able to correct many of the errors that they make from being overly urgent in their initial responses. Any change in strategy between standard decision-making tasks and the explicit double responding paradigm would decrease the generalizability for the findings of explicit double responding assessments to standard speeded decision-making tasks, making the assessment of double responding less useful in understanding the decision process more generally.

As discussed above, implicit and explicit double responding behaviour each have their own unique advantages and limitations, and therefore, we do not wish to make any claims about which type of double responding behaviour researchers should assess. Furthermore, our choice of implicit double responding behaviour is not intended as an endorsement of implicit double responding over explicit double responding, and was based upon (1) our desire to keep our results as generalizable as possible to standard speeded decision-making tasks, and (2) the availability of existing implicit double responding data. However, we believe that future research may be able to provide a solution to this problem through careful experimentation. Specifically, as the primary concern regarding explicit double responding is that participants' initial responses may be influenced by their knowledge that they can correct their initial response with a second response, experimentally establishing that introducing this information does not alter participants' initial responding behaviour would help to mitigate these concerns. However, finding a set of instructions that are able to convey this information, while also restricting participants' freedom to change how they perform the task, may prove challenging. Beyond informing participants of their ability to correct their errors through secondary responses, these instructions could also attempt to restrict participants' ability to rely on second responses by informing participants that they (1) should avoid relying on the secondary responses, (2) will only be partially correct when correcting an initial error, (3) are only allowed a certain number of corrections and must use them wisely, or (4) will only be able to correct their errors on a certain proportion of randomly selected trials. However, given that previous research has suggested that instructing participants to make multiple responses can alter their decision-making behaviour (Yearsley & Pothos, 2016), even in the case of perceptual decision-making (Kvam et al., 2015), it may not be possible to instruct participants to make secondary responses in a way that does not also influence their initial responding behaviour.

One important point that we wish to address is how our choice of implicit double responding over explicit double responding may have influenced our results and conclusions. As discussed above, implicit double responding paradigms will only consist of involuntary second responses that involve some spillover of information from the decision-making process, where explicit double responding paradigms will consist of both implicit double responses and purposeful corrections of initial errors. Importantly, this means that the proportion of double responses should always be *higher* for explicit double responding paradigms than implicit double responding paradigms, as explicit double responding paradigms will consist of both implicit and explicit double responding behaviour. In our study, models with lateral inhibition provided the best account of empirical data from two experiments, primarily due to all other models providing a large over-estimation of the double response proportion, particularly when participants were instructed to perform the tasks with an emphasis on accurate responding. However, as the use of an explicit double responding paradigm should theoretically result in a higher proportion of double responses than an implicit double responding paradigm, our choice of implicit double responding may have led to the small proportion of double responses that most models were unable to capture, and an explicit double responding paradigm may have led to a larger proportion of double responses that would have provided less distinction between the models. However, it should be noted that although explicit double responding should increase the proportion of double responses, this increase may not be enough to change our results. Furthermore, as our study focused on the joint predictions of the models for both the initial choice response time distributions and double responding proportions, and as the explicit double responding paradigm may change how participants behave for the initial response, changing to explicit double responding may also alter the observed initial choice response time distributions, meaning that the models may still differ in their joint predictions. Although we believe that our choice of implicit double responding behaviour was sensible, and that our results and conclusions have potentially meaningful implications for our understanding of the decision-making process, we believe that the potential differences between implicit and explicit double responding behaviour further motivates both (1) the assessment of explicit double responding behaviour, and (2) an investigation into how to inform participants of their ability to make a second, corrective response without altering how they behave for initial responses. It should also be noted that our 250 ms double response deadline may not be adequate time to allow participants to make purposeful corrections of perceived errors, and that researchers who use explicit double responding paradigms may wish to implement longer deadlines.

4.2. Future directions related to double responding

We believe that the introduction of double responding as an additional constraint for models of speeded decision-making provides many potential avenues for future research, and we attempt to discuss some of these possibilities here. Firstly, although models with lateral inhibition were able to provide an accurate account of the joint response time/choice/double response data under both potential definitions of how double responses are triggered, their account of double responding was not perfect. Specifically, the lateral inhibition models still over-predicted the number of double responses that occurred – especially in Experiment 2 – and this was also true of the models without inhibition under the alternative definition for how double responses are triggered. In addition, no model – including the lateral inhibition models – predicted the positive skew present in the double response time distributions, with all models predicting either uniform distributions or distributions with a slight negative skew. Therefore, future research could aim to create new variants of EAMs that may be able to meet these benchmarks. Furthermore, we would again like to note that our study is the first to use double responding as an additional constraint for models of speeded decision-making, and therefore, we do not believe that it would be appropriate for our study to provide broad conclusions about which models of decision-making are superior to others. However, we believe that the general success of lateral inhibition within our study, in combination with previous experimental tests favouring lateral inhibition (Teodorescu & Usher, 2013), shows a growing body of evidence for the importance of lateral inhibition in speeded decision-making.

Another important point of consideration is that our extensions of EAMs to double responding are only two of a large number of possibilities. Our primary choice of extension was based on trying to provide the most natural extension of the standard EAM framework, with accumulation continuing after the initial decision, and a double response being triggered by the other accumulator hitting the threshold. Our alternative extension was based on the potential proposal that double responding should not occur in cases where the evidence for the initial response is still higher than for the other alternative, requiring the additional condition that the initial response alternative has dropped below the threshold. However, there are many other sensible and simple extensions that could be implemented, which could lead to different, and potentially better, accounts of double responding. For example, one could suggest that double responses might purely be due to a change in which accumulator has the greatest share in evidence, meaning that a double response would be triggered if the evidence for the losing alternative were to go above the evidence for the winning alternative after the initial decision. Furthermore, previous research has suggested that the level of lateral inhibition may increase after a decision is made (Bronfman et al., 2015), meaning that allowing greater levels – or in the case of the independent models with the alternative definition, some lateral inhibition – during the double responding process may further improve the ability of these models to account for double response times and proportions. Alternatively, our assumption that double responses are the result of changes in the decision process may be incorrect, and instead double responding may result from motor errors where participants give the wrong response (i.e., press the wrong button), which they attempt to quickly correct following their initial response. We believe that our study has provided an important initial comparison between different EAMs in how well they can account for double responding, showing that double responding can be used to help distinguish between different EAMs, and future research should aim to compare other possible extensions of EAMs to double responding.

More generally, we believe that our study displays the importance of tightly constraining cognitive models, and how this can be done using additional sources of data. However, it should be noted that these additional sources of data should be theoretically sensible, as the use of theoretically unrelated data (i.e., data produced by a difference psychological process) may result in misleading conclusions. We believe that double responding provides a theoretically sensible constraint for models of speeded decision-making, as EAMs can be naturally extended to capture these data, but that future research should attempt to find further, and tougher, constraints for these models. For example, future research could perform a similar comparison to our study, though using partial errors measured through EMG recordings (Coles et al., 1985; Servant et al., 2015) as the additional constraint. Moreover, future studies could potentially combine multiple constraints, providing an even tougher test for models of speeded decision-making. For example, the combination of choice response time distributions, double responses, *and* partial errors provide information about the decision process from before, during, and after the initial response, and therefore, would likely provide an extremely tough constraint for EAMs.

Aside from the potential to further constrain models of decision making, double responding also provides an interesting avenue of theoretical research in better understanding why people change their mind. Previous studies have assessed change of mind in a variety of different ways, both assessing how people's behaviour changes after making an error (e.g., error correcting responses, post-error slowing), and how people's preferences change of mind during a decision effects their motor output (e.g., joystick movement, EMG). Although some models currently exist that involve change of mind, such as the diffusion-based models of conflict tasks (White, Servant, & Logan, 2018), these models are still only constrained by the standard data of response time and choice, and are only intended for specific types of tasks (i.e., conflict tasks). Therefore, we believe that double responding provides an interesting avenue to better understand why people change their mind, which can be used to extend decision making models to explain these post-decision changes in mind. There are also several other common experimental manipulations that may be able to provide more double responding benchmarks for decision making models to explain, such as difficulty manipulations (e.g., do double responses occur proportionally more often after errors on easy trials, or errors on hard trials?), response bias manipulation (e.g., do double responses happen more often after people respond in favour of the initial bias alternative?), or manipulations that place a stronger emphasis on fast responding (e.g., do double responses happen more often as more emphasis is placed on fast responding) such as deadline paradigms (e.g., Evans, Hawkins, & Brown, 2020). Moreover, future research may benefits from exploring individual differences in double responding, in order to understand why some people may be more likely to double respond than others. This may involve attempting to relate double responding behaviour to other constructs (e.g., personality traits), or to the parameters of EAMs (e.g., threshold) estimated from the standard choice response time distributions. Finally, double responding could be used as an additional source of constraint in the application of EAMs to theoretical and applied problems, such as assessing the robustness of selective influence manipulations (e.g., Dutilh et al., 2018). We believe that each of these potential research questions provide informative directions for future research involving double responding, which may help to further our understanding of speeded decision-making and changes of mind.

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at https://doi.org/10.1016/j.cogpsych.2020.101292.

References

Andersen, R. A., & Buneo, C. A. (2002). Intentional maps in posterior parietal cortex. *Annual Review of Neuroscience, 25*(1), 189–220.

Becker, W., & Jürgens, R. (1979). An analysis of the saccadic system by means of double step stimuli. *Vision Research, 19*(9), 967–983.

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & 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*, 700–765.

Bronfman, Z. Z., Brezis, N., Moran, R., Tsetsos, K., Donner, T., & Usher, M. (2015). Decisions reduce sensitivity to subsequent information. *Proceedings of the Royal Society B: Biological Sciences, 282*(1810), 20150228.

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

Brown, S. D., Ratcliff, R., & Smith, P. L. (2006). Evaluating methods for approximating stochastic differential equations. *Journal of Mathematical Psychology, 50*, 402–410.

Burle, B., Spieser, L., Servant, M., & Hasbroucq, T. (2014). Distributional reaction time properties in the Eriksen task: Marked differences or hidden similarities with the

Simon task? *Psychonomic Bulletin & Review, 21*(4), 1003–1010.
Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic–cognitive approach to decision making in an uncertain environment. *Psychological Review, 100*, 432–459.
Camalier, C., Gotler, A., Murthy, A., Thompson, K., Logan, G., Palmeri, T., & Schall, J. (2007). Dynamics of saccade target selection: race model analysis of double step and search step saccade production in human and macaque. *Vision Research, 47*(16), 2187–2211.
Colby, C. L., & Goldberg, M. E. (1999). Space and attention in parietal cortex. *Annual Review of Neuroscience, 22*(1), 319–349.
Coles, M. G., Gratton, G., Bashore, T. R., Eriksen, C. W., & Donchin, E. (1985). A psychophysiological investigation of the continuous flow model of human information processing. *Journal of Experimental Psychology: Human Perception and Performance, 11*(5), 529.
Crüwell, S., Stefan, A. M., & Evans, N. J. (2019). Robust standards in cognitive science. *Computational Brain & Behavior,* 1–11.
Ditterich, J. (2010). A comparison between mechanisms of multi-alternative perceptual decision making: ability to explain human behavior, predictions for neurophysiology, and relationship with decision theory. *Frontiers in Neuroscience, 4*, 184.
Donkin, C., Brown, S., Heathcote, A. J., & Wagenmakers, E. J. (2011). Diffusion versus linear ballistic accumulation: Different models for response time, same conclusions about psychological mechanisms? *Psychonomic Bulletin & Review, 55*, 140–151.
Dutilh, G., Annis, J., Brown, S. D., Cassey, P., Evans, N. J., Grasman, R. P., & Donkin, C. (2018). The quality of response time data inference: A blinded, collaborative assessment of the validity of cognitive models. *Psychonomic Bulletin & Review*.
Dutilh, G., Forstmann, B. U., Vandekerckhove, J., & Wagenmakers, E. J. (2013). A diffusion model account of age differences in posterror slowing. *Psychology and Aging, 28*, 64–76.
Dutilh, G., Vandekerckhove, J., Tuerlinckx, F., & Wagenmakers, E. J. (2009). A diffusion model decomposition of the practice effect. *Psychonomic Bulletin & Review, 16*, 1026–1036.
Evans, N. J. (2019a). Assessing the practical differences between model selection methods in inferences about choice response time tasks. *Psychonomic Bulletin & Review,* 1–29.
Evans, N. J. (2019b). A method, framework, and tutorial for efficiently simulating models of decision-making. *Behavior Research Methods,* 1–15.
Evans, N.J. (2019c). What factors are most important in finding the best model of a psychological process? Comment on Navarro (2018). *PsyArXiv*.
Evans, N. J., Bennett, A. J., & Brown, S. D. (2018). Optimal or not; depends on the task. *Psychonomic Bulletin & Review,* 1–8.
Evans, N. J., & Brown, S. D. (2017). People adopt optimal policies in simple decision-making, after practice and guidance. *Psychonomic Bulletin & Review, 24*, 597–606.
Evans, N. J., & Hawkins, G. E. (2019). When humans behave like monkeys: Feedback delays and extensive practice increase the efficiency of speeded decisions. *Cognition, 184*, 11–18.
Evans, N. J., Hawkins, G. E., & Brown, S. D. (2020). The role of passing time in decision-making. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 46*(2), 316–326.
Evans, N. J., Rae, B., Bushmakin, M., Rubin, M., & Brown, S. D. (2017). Need for closure is associated with urgency in perceptual decision-making. *Memory & Cognition, 45*, 1193–1205.
Evans, N. J., & Servant, M. (2020). A comparison of conflict diffusion models in the flanker task through pseudolikelihood bayes factors. *Psychological Review, 127*(1), 114.
Falkenstein, M., Hohnsbein, J., Hoormann, J., & Blanke, L. (1991). Effects of crossmodal divided attention on late ERP components. II. Error processing in choice reaction tasks. *Electroencephalography and Clinical Neurophysiology, 78*, 447–455.
Gigerenzer, G., Hoffrage, U., & Kleinbölting, H. (1991). Probabilistic mental models: A Brunswikian theory of confidence. *Psychological Review, 98*, 506–528.
Hasbroucq, T., Burle, B., Akamatsu, M., Vidal, F., & Possamai, C. A. (2001). An electromyographic investigation of the effect of stimulus–response mapping on choice reaction time. *Psychophysiology, 38*, 157–162.
Holmes, W. R. (2015). A practical guide to the probability density approximation (PDA) with improved implementation and error characterization. *Journal of Mathematical Psychology, 68*, 13–24.
Kvam, P. D., Pleskac, T. J., Yu, S., & Busemeyer, J. R. (2015). Interference effects of choice on confidence: Quantum characteristics of evidence accumulation. *Proceedings of the National Academy of Sciences, 112*(34), 10645–10650.
Laming, D. R. J. (1968). *Information theory of choice–reaction times.* London: Academic Press.
Leite, F. P., & Ratcliff, R. (2010). Modeling reaction time and accuracy of multiple–alternative decisions. *Attention, Perception & Psychophysics, 72*, 246–273.
Lerche, V., & Voss, A. (2019). Experimental validation of the diffusion model based on a slow response time paradigm. *Psychological Research, 83*, 1194–1209.
Luce, R. D. (1986). *Response times.* New York: Oxford University Press.
Maylor, E. A., & Rabbitt, P. (1987). Effects of alcohol and practice on choice reaction time. *Perception & Psychophysics, 42*, 465–475.
Merkle, E. C., & Van Zandt, T. (2006). An application of the poisson race model to confidence calibration. *Journal of Experimental Psychology: General, 135*(3), 391.
Miletić, S., Turner, B. M., Forstmann, B. U., & van Maanen, L. (2017). Parameter recovery for the leaky competing accumulator model. *Journal of Mathematical Psychology, 76*, 25–50.
Moran, R., Teodorescu, A. R., & Usher, M. (2015). Post choice information integration as a causal determinant of confidence: Novel data and a computational account. *Cognitive Psychology, 78*, 99–147.
Murthy, A., Ray, S., Shorter, S. M., Priddy, E. G., Schall, J. D., & Thompson, K. G. (2007). Frontal eye field contributions to rapid corrective saccades. *Journal of Neurophysiology, 97*(2), 1457–1469.
Navarro, D. J. (2019). Between the devil and the deep blue sea: Tensions between scientific judgement and statistical model selection. *Computational Brain & Behavior, 2*(1), 28–34.
Pilly, P. K., & Seitz, A. R. (2009). What a difference a parameter makes: A psychophysical comparison of random dot motion algorithms. *Vision Research, 49*, 1599–1612.
Pleskac, T. J., & Busemeyer, J. R. (2010). Two-stage dynamic signal detection: A theory of choice, decision time, and confidence. *Psychological Review, 117*, 864–901.
Rabbitt, P. (1966). Errors and error correction in choice–response tasks. *Journal of Experimental Psychology, 71*, 264–272.
Rabbitt, P. (1967). Time to detect errors as a function of factors affecting choice-response time. *Acta Psychologica, 27*, 131–142.
Rabbitt, P. (1968). Three kinds of error-signalling responses in a serial choice task. *Quarterly Journal of Experimental Psychology, 20*, 179–188.
Rabbitt, P. (1969). Psychological refractory delay and response-stimulus interval duration in serial, choice-response tasks. *Acta Psychologica, 30*, 195–219.
Rabbitt, P. (2002). Consciousness is slower than you think. *The Quarterly Journal of Experimental Psychology Section A, 55*, 1081–1092.
Rabbitt, P., Cumming, G., & Vyas, S. (1978). Some errors of perceptual analysis in visual search can be detected and corrected. *The Quarterly Journal of Experimental Psychology, 30*, 319–332.
Rabbitt, P., & Rodgers, B. (1977). What does a man do after he makes an error? An analysis of response programming. *Quarterly Journal of Experimental Psychology, 29*, 727–743.
Rabbitt, P., & Vyas, S. (1981). Processing a display even after you make a response to it. How perceptual errors can be corrected. *The Quarterly Journal of Experimental Psychology Section A, 33*, 223–239.
Rae, B., Heathcote, A., Donkin, C., Averell, L., & Brown, S. (2014). The hare and the tortoise: Emphasizing speed can change the evidence used to make decisions. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 40*, 1226–1243.
Rahnev, D., & Denison, R. N. (2018). Suboptimality in perceptual decision making. *Behavioral and Brain Sciences,* 1–107.
Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review, 85*, 59–108.
Ratcliff, R., & Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science, 9*, 347–356.
Ratcliff, R., & Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Review, 111*, 333–367.
Ratcliff, R., & Starns, J. J. (2009). Modeling confidence and response time in recognition memory. *Psychological Review, 116*, 59–83.
Ratcliff, R., & Starns, J. J. (2013). Modeling response times, choices, and confidence judgments in decision making. *Psychological Review, 120*, 697–719.
Ratcliff, R., & Tuerlinckx, F. (2002). Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability.

Psychonomic Bulletin & Review, 9, 438–481.
Resulaj, A., Kiani, R., Wolpert, D. M., & Shadlen, M. N. (2009). Changes of mind in decision–making. *Nature, 461*, 263–266.
Schwarz, G. (1978). Estimating the dimension of a model. *Annals of Statistics, 6*, 461–464.
Servant, M., White, C., Montagnini, A., & Burle, B. (2015). Using covert response activation to test latent assumptions of formal decision-making models in humans. *Journal of Neuroscience, 35*(28), 10371–10385.
Smith, P. L., & Little, D. R. (2018). Small is beautiful: In defense of the small-n design. *Psychonomic Bulletin & Review, 25*(6), 2083–2101.
Stone, M. (1960). Models for choice–reaction time. *Psychometrika, 25*, 251–260.
Teodorescu, A. R., & Usher, M. (2013). Disentangling decision models: From independence to competition. *Psychological Review, 120*, 1–38.
Ter Braak, C. J. (2006). A markov chain monte carlo version of the genetic algorithm differential evolution: Easy bayesian computing for real parameter spaces. *Statistics and Computing, 16*, 239–249.
Tillman, G., & Logan, G. D. (2017). The racing diffusion model of speeded decision making. doi: https://doi.org/10.31234/osf.io/xuwbk.
Tsetsos, K., Gao, J., McClelland, J. L., & Usher, M. (2012). Using time-varying evidence to test models of decision dynamics: Bounded diffusion vs. the leaky competing accumulator model. *Frontiers in Neuroscience, 6*, 79.
Turner, B. M., & Sederberg, P. B. (2014). A generalized, likelihood-free method for posterior estimation. *Psychonomic Bulletin & Review, 21*, 227–250.
Turner, B. M., Sederberg, P. B., Brown, S. D., & Steyvers, M. (2013). A method for efficiently sampling from distributions with correlated dimensions. *Psychological Methods, 18*, 368–384.
Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science, 211*, 453–458.
Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review, 108*, 550–592.
Van Ravenzwaaij, D., Dutilh, G., & Wagenmakers, E. J. (2012). A diffusion model decomposition of the effects of alcohol on perceptual decision making. *Psychopharmacology, 219*, 1017–1025.
Vandekerckhove, J., Tuerlinckx, F., & Lee, M. D. (2008). A Bayesian approach to diffusion models of decision–making. In V. M. Sloutsky, B. C. Love, & K. McRae (Eds.). *Proceedings of the 30th annual conference of the cognitive science society* (pp. 1429–1434). Cognitive Science Society.
Van Zandt, T., & Maldonado-Molina, M. M. (2004). Response reversals in recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 30*(6), 1147.
Vickers, D., & Lee, M. D. (1998). Dynamic models of simple judgments: I. Properties of a self–regulating accumulator module. *Nonlinear Dynamics, Psychology, and Life Sciences, 2*, 169–194.
Vickers, D., & Lee, M. D. (2000). Dynamic models of simple judgments: II. Properties of a self–organizing PAGAN (parallel, adaptive, generalized accumulator network) model for multi–choice tasks. *Nonlinear Dynamics, Psychology, and Life Sciences, 4*, 1–31.
Voss, A., Rothermund, K., & Voss, J. (2004). Interpreting the parameters of the diffusion model: An empirical validation. *Memory & Cognition, 32*, 1206–1220.
Wagenmakers, E. J., van der Maas, H. J. L., & Grasman, R. P. P. P. (2007). An EZ–diffusion model for response time and accuracy. *Psychonomic Bulletin & Review, 14*, 3–22.
White, C. N., Servant, M., & Logan, G. D. (2018). Testing the validity of conflict drift-diffusion models for use in estimating cognitive processes: A parameter-recovery study. *Psychonomic Bulletin & Review, 25*, 286–301.
Yearsley, J. M., & Pothos, E. M. (2016). Zeno's paradox in decision-making. *Proceedings of the Royal Society B: Biological Sciences, 283*(1828), 20160291.
Yeung, N., Botvinick, M. M., & Cohen, J. D. (2004). The neural basis of error detection: conflict monitoring and the error-related negativity. *Psychological Review, 111*, 931–959.