



UvA-DARE (Digital Academic Repository)

Cognitive Control of Choices and Actions

Heathcote, Andrew; Verbruggen, Frederick; Boehler, C. Nico; Matzke, Dora

DOI

[10.1007/978-3-031-45271-0_14](https://doi.org/10.1007/978-3-031-45271-0_14)

Publication date

2024

Document Version

Final published version

Published in

An Introduction to Model-Based Cognitive Neuroscience

License

Article 25fa Dutch Copyright Act (<https://www.openaccess.nl/en/policies/open-access-in-dutch-copyright-law-taverne-amendment>)

[Link to publication](#)

Citation for published version (APA):

Heathcote, A., Verbruggen, F., Boehler, C. N., & Matzke, D. (2024). Cognitive Control of Choices and Actions. In B. U. Forstmann, & B. M. Turner (Eds.), *An Introduction to Model-Based Cognitive Neuroscience* (2nd ed., pp. 361-385). Springer. https://doi.org/10.1007/978-3-031-45271-0_14

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, P.O. Box 19185, 1000 GD Amsterdam, The Netherlands. You will be contacted as soon as possible.

Cognitive Control of Choices and Actions



Andrew Heathcote, Frederick Verbruggen, C. Nico Boehler, and Dora Matzke

Abstract We review model-based neuroscience work on cognitive control of choices and actions. We consider both strategically deployed executive processes and more automatic influences, first in binary choice tasks and then in more complex tasks. These include “conflict” tasks, where automatic and executive control processes sometimes act in opposition; delay discounting tasks, which require self-control to obtain larger rewards; and tasks where routine actions are occasionally interrupted by cues requiring different action or the inhibition of action. For all of these tasks, dynamic cognitive models have been developed based on the idea of accumulating evidence. They have also been studied by traditional neuroscience methods, but direct links to the cognitive models have not always been made. We detail the way in which progress has been made with model-based neuroscience methods in some cases and in others highlight how this points the way towards opportunities for progress. We emphasise generative Bayesian estimation methods that are well suited to the complexities of model-based neuroscience and provide exercises with open-source code that allow readers to develop skills with models relevant to cognitive control.

Keywords Cognitive control · Conflict tasks · Evidence accumulation models

1 Introduction

Choice and action control impacts on the entire gamut of psychological processes from signal detection to action selection and action execution (Verbruggen et al., 2014), and so it needs cognitive models that can encompass this full range.

A. Heathcote (✉)
University of Amsterdam, Amsterdam, The Netherlands
University of Newcastle, Callaghan, NSW, Australia

F. Verbruggen · C. N. Boehler
Ghent University, Ghent, Belgium

D. Matzke
University of Amsterdam, Amsterdam, The Netherlands

© Springer Nature Switzerland AG 2024

B. U. Forstmann, B. M. Turner (eds.), *An Introduction to Model-Based Cognitive Neuroscience*, https://doi.org/10.1007/978-3-031-45271-0_14

Evidence-accumulation models (EAMs), which choose actions based on gradually accruing required amounts of particular types of information (Donkin & Brown, 2018), fulfil this desideratum, and have provided the framework of choice for models of cognitive control. Action control can also require complicated computations that resolve conflicting information and conflicting demands among two or more stimuli, stimulus attributes, and response options. For this reason, cognitive control models have increasingly used a flexible type of evidence-accumulation architecture constituted of accumulators racing, either independently or interactively, in order to select an internal state or an external action (for a review, see Heathcote and Matzke (2022)). In the exercises for this chapter, we enable the reader to fit models based on various such architectures, for instance, using an independent race between single barrier diffusion processes (Tillman et al., 2020; see Box 1).

Box 1: The Racing Diffusion (RD) Model of Binary Choice

Each accumulator adds up noisy evidence (red and green lines) accruing at different rates (greater for the correct than error choice). The chosen response corresponds to the first accumulator to reach its threshold, with RT equal to the time to do so plus the sum of times to encode the stimulus and produce a response (non-decision time). Rates are a locus for controlling signal detection through selective attention (evidence quality) and effort (overall speed), as is the encoding portion of non-decision time. Thresholds and the response production portion of non-decision time provide loci for controlling action selection and action execution. The distribution of finishing times for a single accumulator follows a shifted Wald distribution (Heathcote, 2004; Matzke & Wagenmakers, 2009), with Fig. 1c showing the RT distributions for the 78% of trials where the correct (i.e. higher mean rate) accumulator wins and for the remaining 22% of trials when the error (i.e. lower mean rate) accumulator wins.




Fig. 1 Ten simulated trajectories for correct (left panel; rate = 2) and error (middle panel; rate = 0.5) accumulators and corresponding RT distributions (right panel; threshold crossing time for the winner plus 0.25 s non-decision time). <https://tinyurl.com/4n9h89ch> under CC-BY 2.0 license (<https://creativecommons.org/licenses/by/2.0/>)

EAMs specify what can be controlled but not how control is achieved, although identifying where control is exerted can often be informative on this question. Importantly, EAMs support rigorous theory testing by linking these details to a fine-grained characterisation of behaviour in terms of the responses that are executed and corresponding response times (RT)—the period that elapses between the appearance of a stimulus and the execution of a motor response—as well as how often responses are withheld. Crucially, in the current context, EAM parameters provide a natural medium to link neuroscience and behavioural data. Unlike “manifest” (i.e. directly observable) behavioural quantities, model parameters are typically “latent” quantities that can only be estimated indirectly based on finding values that provide an accurate and parsimonious account of the data. Correspondences between latent parameters and manifest physiological measures create a synergy between cognitive modelling and neuroscience: the difficult (and in some cases impossible) task of inferring parameters is made easier by linking them to neuroscience measures, and the model aids in the interpretation of neuroscientific data that can otherwise be ambiguous.

In the following, we first briefly review early successes of model-based neuroscience combining standard EAMs and functional magnetic resonance imaging (fMRI) to understand the control of *when to act*. We then review two further areas of cognitive control, control over *which actions to make* and *which actions to withhold*, that make use of fMRI; electroencephalographic (EEG) signals, both in the amplitude and frequency domains; transcranial magnetic stimulation (TMS); and electromyography (EMG). In each of these three areas, we highlight recent advances in cognitive modelling that could benefit from the model-based neuroscience approach. Note that our focus here is on the cognitive and neural processes associated with cognitive control rather than the representations on which it is based (Freund et al., 2021). In terms of Nigg’s (2017) classification, these processes represent low-level executive functions exerting control over actions at short time scales. We do address control over impulsive choice in the delay-discounting paradigm, which is usually studied separately from action control and also associated to some degree with higher-level executive functions, longer time scales, and different manifestations in clinical groups (see also Dalley and Robbins (2017)). Our review shows that EAMs can provide a framework to integrate these two aspects of cognitive control at a process level. However, we do not address more complex cognitive control (e.g. Alexander & Brown, 2015; Botvinick et al., 2009, Dayan, 2008; Rougier et al., 2005; see O’Reilly et al., 2010, for a review).

2 Controlling When to Act

As illustrated in Box 1, EAM threshold parameters provide a means of controlling when to act. It takes longer to accrue enough evidence to reach a higher threshold, and so decisions are slowed. However, the decisions also become more accurate because longer accumulation diminishes the potential for noise in the accumulation

process (evident in the irregular trajectories in Fig. 1) to allow the wrong accumulator to win. Hence, the evidence threshold controls a trade-off between speed and accuracy that participants have to manage when making decisions under time pressure.

In an early example of model-based neuroscience, Forstmann et al. (2008) studied the speed-accuracy trade-off through fMRI and fitting an EAM, the linear ballistic accumulator (LBA, Brown & Heathcote, 2008). The LBA is a racing accumulator model, but unlike the racing diffusion (RD) model shown in Box 1, accumulation is deterministic with a constant rate (i.e. the trajectories in Fig. 1 would be straight lines). Instead of decision noise being in the moment-to-moment evidence total, in the LBA, it manifests in normally distributed trial-to-trial variations in the rates (an assumption shared with the LATER model; Carpenter & Williams, 1995) and uniformly distributed noise in the starting point of evidence accumulation. Forstmann et al. cued participants on a trial-by-trial basis to make either quicker or more accurate decisions about the predominant direction in a cloud of moving dots and accounted for the associated changes in speed and accuracy with a change in the LBA's threshold parameters. The fMRI data revealed that the right anterior striatum and the right pre-supplementary motor area (pre-SMA) were modulated by the speed vs. accuracy instructions. Across participants, the magnitude of these changes correlated with the magnitude of estimated LBA threshold changes. The results of the "two-stage" procedure (i.e. first fitting a process model to behavioural data and a descriptive model to brain data, and then correlating model parameter estimates) were consistent with the proposal that the propensity to act is controlled by a cortico-striatal loop additionally involving the sub-thalamic nucleus (StN; Frank, 2006; Wiecki & Frank, 2013).

Turner et al. (2015) criticised the two-stage procedure on two grounds: it neglects trial-to-trial changes and it is not *statistically reciprocal*, as the neural data do not influence estimation of the parameters of the behavioural model and vice versa. It also fails to properly account for uncertainty, which attenuates correlation estimates and makes inference overly confident (Ly et al., 2018; Matzke et al., 2017c). Turner et al. asserted that instead a "joint-modelling" framework is the way forward for model-based neuroscience. Joint modelling requires two statistical models for the data on each trial, in their case an independent components analysis (ICA) model for the fMRI data (Eichele et al., 2008), and a Wiener diffusion model (Stone, 1960; often referred to as the drift diffusion model in the neuroscience literature) of the behavioural data. As depicted in Fig. 2, their behavioural model is similar to the RD model, but it consists of a single accumulator with two thresholds, one for each response in a binary choice task. The variant of this model most widely used with behavioural data alone, the diffusion decision model (Ratcliff & McKoon, 2008), makes the same assumptions about trial-to-trial rate and starting-point variability as the LBA. However, in both cases, these variations are mathematically integrated out, so that only estimates of the parameters of the overall distributions are obtained. In contrast, Turner et al. used a Bayesian hierarchical model with both trial and participants levels (Vandekerckhove et al., 2011), enabled by the numerical methods of Navarro and Fuss (2009), to estimate rates and starting points for individual

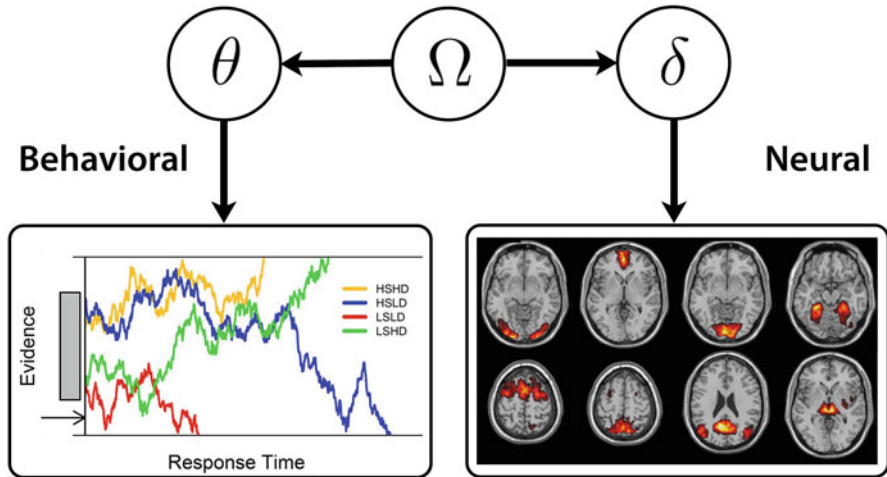


Fig. 2 The neural drift diffusion model. Parameters of the neural ICA model (δ) are linked to parameters of the behavioural model (θ) through parameters (Ω) quantifying the variance that they share. Four evidence trajectories are depicted for the behavioural (diffusion decision) model in the lower left panel exemplifying combinations of high and low starting points (HS and LS) and high and low drift rates (HD and LD). (Reproduced with permission from Turner et al. (2015))

trials and link them to trial-by-trial BOLD estimates through a multi-variate normal distribution quantifying the correlations they share, which they called the neural drift diffusion model (NDDM).

We will not describe Turner et al.'s (2015) results further here, as their analysis did not focus on cognitive control. However, a later analysis by Turner et al. (2017b) used a refined version of joint modelling to investigate the conventional wisdom that EAMs produce speed-accuracy trade-offs purely by threshold changes. More recent work with behavioural data alone has found that non-decision time and accumulation rates also change (e.g. Lerche & Voss, 2018; Rae et al., 2014). Turner et al.'s (2017b) refined method replaced the NDDM's multi-variate normal link—which rapidly becomes impractical when investigating large brain networks because the required number of parameters is a quadratic function of the number of brain areas modelled—with a factor analysis model, where the increase is only linear. They used this model to examine a large number of brain areas and found evidence for changes in the brain networks associated with non-decision time and rates between speed and accuracy emphasis conditions, consistent with thresholds not being the only mechanism involved.

We close this section with an opportunity related to increasing concerns with the role of time in decision-making in circumstances where more reward might be gained from correct responses per unit time by requiring less evidence as decision time progresses (e.g. when decision difficulty varies unpredictably from trial to trial). This has led to the proposal of models with “collapsing bounds” (thresholds that decrease over time) or adding an “urgency” input that drives responding (e.g.

Churchland et al., 2008). Although it is now clear that neither mechanism provides much reward gain (Boehm et al., 2020), there are some circumstances both in humans and primates that are better fit by collapsing than constant bounds (Palestro et al., 2018b; Evans et al., 2020). Unfortunately, collapsing bounds and urgency models are usually computationally expensive (Hawkins et al., 2015) and hence difficult to work with (see, however, Trueblood et al. (2021)). The first opportunity box describes a new more tractable approach that is ripe for investigation with model-based neuroscience methods.

Opportunity 1: Further Exploring Control of Action Timing

Hawkins and Heathcote (2021) proposed an elaboration of the model depicted in [Box 1](#), the timed racing diffusion model (TRDM) that, consistent with theories of human time estimation (Simen et al., 2016), measures the passage of time with an extra diffusion process. The TRDM, with a choice being guessed if the timer wins, provides as good or better fits to data than the leading alternative models in paradigms favouring both fixed and collapsing bounds. It can also simultaneously account for decision-making and time estimation tasks by the same individuals. As TRDM is analytically highly tractable, it is ideal for joint modelling, which could explore the overlap between the cortico-striatal decision-making circuit (Frank, 2006) and the cortico-thalamic-basal ganglia timing circuit (Merchant et al., 2013), as well as joint modelling of more than one type of task (Kvam et al., 2021). As discussed further in the Sect. 4, it also provides a good model of the sustained attention response task used to study failures of control due to mind wandering.

3 Controlling Which Actions to Take

We begin by discussing a series of papers focusing on choice selection in the delay-discounting task (see also Dai & Busemeyer, 2014). Rodriguez et al. (2014) set the stage with an LBA model of behaviour alone, which was first extended to two-stage fMRI modelling (Rodriguez et al., 2015) and joint modelling of EEG and fMRI (Turner et al., 2016). This work mainly addressed where and how value is represented in the brain, whereas the final study (Turner et al., 2019) focused more on control processes, using the Leaky Competing Accumulator (LCA; Usher & McClelland, 2001) model to address behaviour and returning to a two-stage modelling approach. We then turn to a prototypical approach to studying cognitive control through how people resolve decision conflicts in tasks such as the Stroop (MacLeod, 1991), Hommel (2011), and Flanker (Eriksen, 1995). There is an associated voluminous experimental literature, numerous applications to the assessment of individual differences in selective attention, and a range of cognitive

models that have been applied to one (e.g. White et al., 2011) or two (e.g. Ulrich et al., 2015) of these tasks. Here we focus on the Flanker task, where participants indicate the direction of a central arrow and the conflicting information comes from surrounding arrows (e.g. <<><<), and the work of Weichart et al. (2020), which incorporates a type of two-stage modelling of EEG data that Turner et al. (2017a) called a “latent-input” approach as it is based on the dynamics of the model rather than its parameters.

Before addressing these tasks, we note that this section also illustrates a tension in model-based neuroscience work between two-stage procedures and the more statistically optimal joint modelling approach (Palestro et al., 2018a). To date, success with the latter approach has required mathematically tractable EAMs, with likelihood functions that can be quickly and accurately computed, such as the RD, LBA, and diffusion models. More recent work discussed in this section uses the LCA model, which (at least in its full form) is intractable and impractical to estimate in a fully Bayesian manner (Miletić et al., 2017). However, its likelihood can be approximated using computationally intensive probability density approximation (PDA) methods (Turner & Sederberg, 2014), allowing a partially Bayesian analysis.

3.1 Delay Discounting

The delay-discounting task is a type of multi-attribute value-based choice, where participants must choose between immediate “smaller-sooner” and delayed “larger-later” options (\$1 now or \$10 in one month). Research on this task has traditionally focused on the function that computes the subjective value (V) of immediate (i) and delayed (d) options by combining information about the delay (t) and the magnitude of the reward (r), most prominently using a hyperbolic model (e.g. $V = r/(1 + kt)$, where k is an estimated *discount rate* parameter). The dynamics of the process by which an option is selected has received less attention and has typically focused only on choice probability assuming a probabilistic “softmax” decision rule, where the probability of choosing the delayed option is given by $p_D = 1/(1 + e^{-m(V_d - V_i)})$, where m is estimated. Rodriguez et al. (2014) report behavioural modelling using the LBA to provide a simultaneous account of choice and RT. Their experiment randomly selected delays and paired them with rewards that resulted in five conditions with p_D in the set (0.1, 0.3, 0.5, 0.7, 0.9) according to the hyperbolic softmax model of subjective value. For each condition, they freely estimated a rate parameter for an LBA accumulator corresponding to the delay option, v_d , which races with an accumulator for the immediate option with rate $1 - v_d$. The LBA model was fit with hierarchical Bayesian methods (Gelman & Hill, 2007; Lee, 2011) and provided an accurate account of choice probabilities and RT distributions. Consistent with rates encoding value, maximum a posteriori (MAP, that is, the value of the posterior mode) estimates of the difference in its rate parameters (i.e. $2v_d - 1$) for each participant were directly related to the difference in subjective value of the options according to, among others, the hyperbolic model of subjective value.

Rodriguez et al.'s (2014) experiment also collected fMRI and EEG data for different groups of participants. Rodriguez et al. (2015) built on this work, using the fMRI group to examine the general hypothesis that prefrontal cortex (PFC) maintains representations that provide bias signals to other areas that help achieve goals (Duncan, 2006; Miller & Cohen, 2001). Specifically, they examined the hypothesis that ventro-medial PFC (vmPFC) represents value and dorso-lateral PFC (dlPFC) implements the self-control required to obtain greater rewards by choosing delayed options (Figner et al., 2010; Hare et al., 2009). Trial-by-trial estimates of a subset of LBA parameters (van Maanen et al., 2011) were obtained by an empirical Bayes approach, first using MAP estimates from Rodriguez et al. (2014) to sample posterior distributions for each trial, followed by an optimisation procedure to choose the parameter values that best represented that trial. A function of these single trial parameter values was then used as a regressor in a whole-brain general-linear model (GLM) that identified dorsal medial frontal cortex (dmFC), posterior parietal cortex (pPC), and dlPFC. These same areas were identified with relative value (i.e. $|V_d - V_i|$) and RT as regressors, providing converging evidence about the validity of the model-based approach. Functional connectivity analyses indicated that these regions received inputs from relative-value representations in vmPFC and that they provided output to motor regions. Based on the sum of this evidence, they concluded that these areas are the locus of value accumulation.

Turner et al. (2016) moved from the earlier two-stage approach to joint modelling. They made use of data from both Rodriguez et al.'s (2014) fMRI and EEG groups, which included a small subset of participants with both types of data. The neural data consisted of single trial EEG values from four electrodes and the time window (i.e. 300–850 ms) that maximally differentiated value conditions, along with single trial beta values from an fMRI GLM analysis. They showed that fully Bayesian estimation of a multi-variate normal joint model with the LBA (i.e. without a two-stage analysis or empirical Bayesian methods to infer individual-trial parameter estimates) provided an accurate unified account of all three (behavioural, fMRI, and EEG) types of data. Further, this analysis improved as more data entered in the model in two senses. First, relative to fits to behaviour alone, the LBA parameters were more precisely estimated in a joint model with either type of neural data. Precision was even greater in a joint model with both types. Second, as more types of data were used, predictions for new behavioural data (i.e. data to which the model had not been fit) improved. These results highlight the validity, and fundamental advantages of a model-based neuroscience approach, in understanding the links between cognitive and neural processes.

Turner et al. (2019) performed a new experiment to test whether dlPFC indirectly implements self-control by modulating vmPFC value representations (Hare et al., 2009) or does so by acting directly on the decision process without affecting the representation of value (Figner et al., 2010). Building on the work of Turner et al. (2018), who reported results favouring LCA-based over LBA-based models of multi-attribute value choice, they developed a new model of delay discounting based on the LCA. They used the LCA to instantiate a mechanistic account of temporal discounting (rather than the earlier descriptive approach) using a stochastic

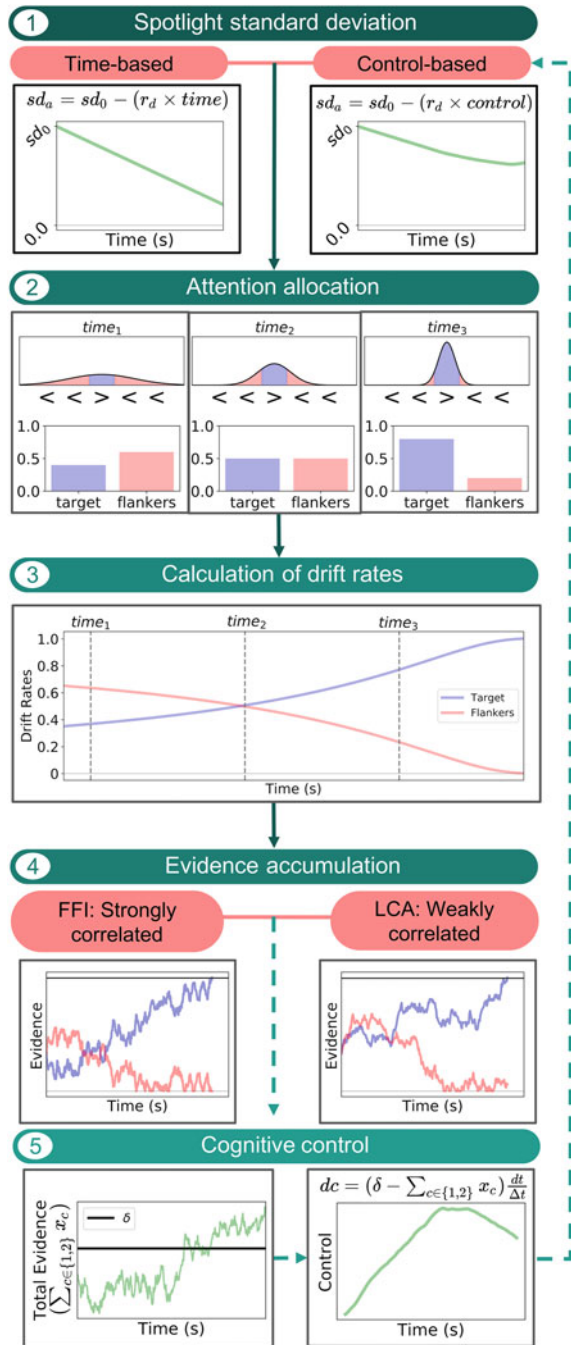
feature sampling process to dynamically construct value (cf., Roe et al., 2001). The complexity of this model makes it mathematically intractable, so that likelihoods had to be obtained by PDA; furthermore, its parameters are not identifiable when all mechanisms are in play, so model selection was used to compare subsets of mechanisms. Full Bayesian model fitting was not possible; however, informative priors enabled point (MAP) estimates to be obtained, although this precluded the use of model selection methods that take account of posterior uncertainty. As in the initial LBA work, empirical Bayesian methods provided individual-trial parameter estimates used in a two-step procedure to link to neural data. Model selection favoured the operation of both direct and indirect mechanisms, with parameters for the direct mechanism linked to all three previously identified value accumulation areas, but no areas linked to the parameter for the indirect value modulation mechanism. Turner et al. (2019) closed by noting the desirability of developing a joint modelling approach with this model.

3.2 *Controlling Spatial Attention*

The indirect mechanism discussed by Turner et al. (2019) is an example of cognitive control exerted through Verbruggen et al.'s (2014) signal detection stage, by altering the representations that provide input to the decision process. Weichart et al. (2020) proposed and compared four alternative models of the Flanker task that exclusively use this sort of control, crossing two types of decision processes with two types of control (see Fig. 3). All of the models are based on White et al.'s (2011) shrinking spotlight model, where flanker information flows to the decision process through an attentional "spotlight" with a Gaussian sensitivity profile. Initially, the spotlight is wide, but as time progresses, its standard deviation (sd) shrinks, increasing the relative influence of the central (target) information. White et al. used a diffusion process, whereas Weichart et al. compared a closely related feedforward inhibition (FFI) process, where accumulators for each response (left vs. right) are perfectly anti-correlated (i.e. a change in one accumulator is mirrored in the opposite change in the other accumulator). They compared the FFI to the LCA, where lateral inhibition causes a weaker negative correlation between accumulators (see Fig. 3). They also compared White et al.'s mechanism, where the spotlight shrinks linearly with time, with one based on Braver's (2012) idea of reactive control exercised through a within-trial reaction to conflict. Their reactive mechanism was based on a control signal (c) that changes at a rate proportional to the difference between the total activation of both accumulators and a target level (δ), so that in contrast to the time-based mechanism, the spotlight eventually stops shrinking.

Weichart et al. (2020) fit these four models to data from a standard Flanker task using PDA-based likelihoods and fully Bayesian estimation. Model selection clearly favoured the control over the time-based mechanism and weakly favoured the LCA over the FFI decision process, although the LCA was qualitatively better at describing the relationship between accuracy and RT distributions. A second

Fig. 3 Flowchart of alternative models fit by Weichart et al. (2020)



experiment that manipulated the similarity of target and flanker stimuli (Servant et al., 2014) more clearly favoured the LCA over FFI due to diverging predictions about RT distributions. A third experiment returned to the standard Flanker task and collected EEG data in order to identify a neural signature of the within-trial control process. MAP parameter estimates were used to simulate behavioural data with associated vectors of control signals, and subsets that sufficiently matched condition, choice, and RT were used to characterise the average evolution of the control signal. These control signals for the ~60% of trials with sufficient matches were then correlated with stimulus-locked EEG in a window from $\tau/2$ after the stimulus to $\tau/2$ before the response, where τ is the estimated non-decision time. Correlations were greatest in right posterior regions and greatest for the LCA activation-based control model, although an LCA model with noise added to the time-based control signal better accounted for differences in EEG between short, medium, and long RT trials.

Opportunity 2: Further Exploring Control of Action Choice

Braver's (2012) dual modes of control framework, which contrasts reactive control with proactive control (e.g. control exerted in anticipation based on knowledge about the upcoming trial), has also been influential in a key component of multi-tasking, event-based prospective memory (PM, that is, carrying out an intended action when cued after a delay period filled with unrelated ongoing activity). Initial EAM analysis indicated a role for proactive control of when to act through a higher threshold delaying responding in the ongoing task in contexts where PM cues occur (Heathcote et al., 2015b). A subsequent LBA-based model, PM decision control (PMDC; Strickland et al., 2018), revealed a major role for reactive control over which action to take, whereby the PM cue reduces the rate of accumulation for the routine choice task through a feedforward inhibitory connection. PM is supported by a broad brain network (Burgess et al., 2011), with both frontal control-related processes in PFC (particularly Brodmann Area 10) related to the dual modes of control framework and ventral parietal network related to memory retrieval (e.g. Lamichhane et al. (2018)). PMDC is mathematically tractable and so provides a suitable basis for a model-based neuroscience investigation using joint modelling.

4 Controlling Which Actions to Withhold

Control of Verbruggen et al.'s (2014) final "action-execution" stage has primarily been studied through either a go/no-go or stop-signal task. In the go/no-go task, participants make a binary choice but are required to execute an action for one

choice and to withhold it for another. Typically, there is a consistent mapping between the choices and acting vs. withholding, which transfers the required control to an automatic process (Shiffrin & Schneider, 1977). When both choices are correct equally often, standard EAMs with a boundary for each response but modified rule for what happens when a threshold is reached (i.e. responding as usual for one threshold and not making any response for the other) work well (e.g. Gómez et al., 2007). However, this is not the case in a variant, the sustained attention response task, where no-go trials are rare and commission errors (i.e. responding on no-go trials) are used to index failures of inhibitory control. In this case, unusually fast commission errors that have been ascribed to rhythmic responding cannot be described by standard EAMs (Hawkins et al., 2019) but can be accommodated by the TRDM model described above, presenting opportunities for further study through model-based neuroscience that have not presently been taken up.

Here, we will focus on the stop-signal task (Matzke et al., 2018), where participants make “go” responses—typically easy binary choices (e.g. does an arrow point left or right)—that they are occasionally (e.g. on 25% of trials) required to withhold when a stop signal (e.g. a tone) occurs (see Fig. 4). The delay between the appearance of the go and stop signals, with longer stop-signal delays (SSD) making it harder to stop, is manipulated, typically using a staircase algorithm, so that half of the stop trials are successfully inhibited (“signal-inhibit” trials) and half not (“signal-respond” trials). Because the mapping between go stimuli and

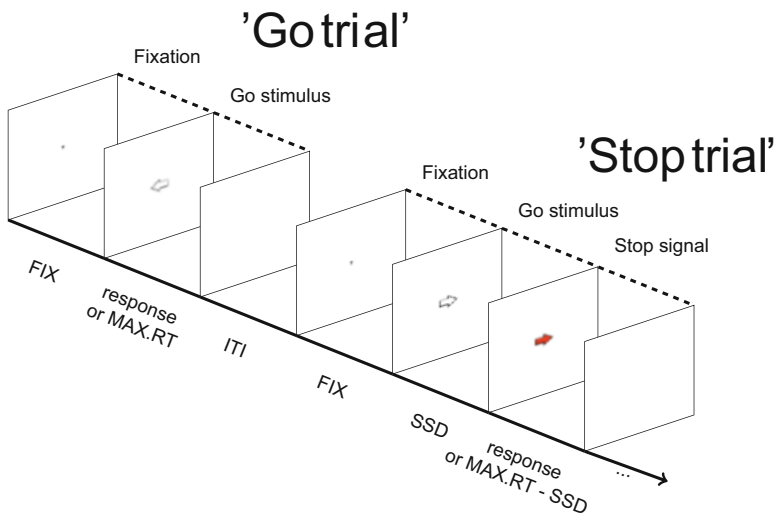


Fig. 4 Trial types in the stop-signal paradigm. The choice task is to press a button corresponding to the direction of the arrow. The stop signal is the arrow turning red. ITI inter-stimulus interval, SSD stop-signal delay, FIX fixation interval, MAX.RT is the maximum period allowed before a new trial commences. (Adapted and reproduced from Verbruggen et al. (2019) under CC-BY license)

responding vs. withholding is inconsistent, automaticity cannot develop, and so control is instead attributed to an “executive” process (Logan & Cowan, 1984; Verbruggen & Logan, 2008).

The key measure of inhibitory ability, the time from the onset of the stop signal to the go process being halted, the stop-signal RT (SSRT), cannot be directly observed based on behaviour alone, so cognitive modelling is viewed as indispensable. For this purpose, Logan and Cowan (1984) proposed a “horse-race” model in which processes (“runners”) triggered by go and stop stimuli race with each other, with the resulting action, or lack of it, determined by the winner. The horse-race model has typically been used to estimate mean SSRT in a non-parametric manner that makes no assumptions about the processes which determine the racer’s speeds and only very general assumptions about the associated distribution of finishing times (e.g. that they do not vary at all). However, this approach requires other strong assumptions, including reliable stop triggering (i.e. the stop racer always enters the race), stochastic independence, go context independence (i.e. the presence of the stop racer does not affect the speed of the go racers), and stop context independence (i.e. the speed of the stop racer is not affected by the SSD).

More recently, Matzke et al. (2013) used the ex-Gaussian distribution (a right skewed distribution providing a good description of RT distributions) to model finishing times in a way that is flexible, although only descriptive (i.e. the ex-Gaussian does not have a clear process interpretation; Matzke & Wagenmakers, 2009). This so-called BEESTS model was later augmented with methods to deal with failures of assumptions made by the non-parametric approach, widening the scope of application to cases with trigger failures (Matzke et al., 2017b) and failures of go context independence (Matzke et al., submitted), as well as allowing modelling of error-prone choices (Matzke et al., 2019). Given the horse-race model has much in common with racing EAMs, it is perhaps surprising that model-based neuroscience using EAMs has not been more widely applied (Sebastian et al., 2018). We review the EAM-based work that has been done, in both non-human primates (Boucher et al., 2007; Logan et al., 2015) and humans (Mittner et al., 2014), as well as more recent work that used the BEESTS model in humans (Skippen et al., 2020; Jana et al., 2020).

4.1 The Stop-Signal Paradigm in Non-human Primates

Boucher et al. (2007) studied the “saccadic-countermanding” stop-signal task where responding was through either a left or right eye movement. They identified aspects of the go and stop processes in the race model with movement and fixation neurons, respectively, in a gaze control network in frontal eye field (FEF) and superior colliculus (SC). A two-stage approach was used to investigate what they claimed is the paradox of mutual inhibition between these neurons and behaviour being very well described by the independent horse-race model. Behavioural data was fit using a point estimation approach with an LCA model corresponding to movement and

fixation neurons in FEF and SC. Additional delays external to the LCA accounted for the remaining time course of go and stop processes in the race model: go delays with assumed and constant values (for the two monkeys studied 80 ms and 35 ms, respectively) for encoding prior to the LCA go unit becoming active, an assumed and constant 10 ms for go response production, and a constant encoding delay before the stop unit becomes active, which was estimated. In the interactive version of this model, lateral inhibition parameters were also estimated, whereas in the independent version, they were set to zero.

Both interactive and independent models accounted for the behavioural data well, but only the interactive model could account for the drop in movement neuron firing rates just prior to SSRT on signal-inhibit trials. Closer investigation of the interactive model showed that it could be mimicked in behaviour by an independent race because there was little interaction during most of the period of go accumulation—due to a substantial stop delay (~60 ms), much faster stop than go accumulation rate (~20 ms to reach threshold), and much stronger inhibition from stop to go units than vice versa—followed by late but potent inhibition of the go unit by the stop unit. This implies that SSRT measured by the independent race model is valid but that most of SSRT is due to an encoding outside the LCA model, which describes (at least for the stop component) only a very short “interruption” stage. Lo et al. (2009) implemented and Wong-Lin et al. (2010) studied versions of the model that were modified to account for the high firing rates that are seen in fixation neurons prior to target onset.

Logan et al. (2015) further pursued this two-stage strategy but rejected Boucher et al.’s (2007) fundamental assumptions mapping stop and go processes to movement and fixation units in the gaze control network. They first considered again a model originally rejected by Boucher et al., where fixation and movement units race independently and an external stop process blocks input to the movement unit when a stop signal occurs. Boucher et al. reported that this “blocked-input” model fit worse than the interactive race model. However, using the same data and an altered fitting procedure, Logan et al. reported the models fit equally well. Logan et al. then extended both the interactive and blocked-input models to fix another problem that had been missed by Boucher et al.: an inability to accommodate steady-state firing rates in the gaze control network before the go stimulus is presented. They compared the fits of these extended “interactive 2.0” and “blocked-input 2.0” models and a “boosted fixation 1.0” model, where the external stop process increases the activation of the fixation unit instead of blocking the input to the movement unit. The interactive 2.0 model was clearly worse, with the other two equal except that the boosted fixation model over-predicted post-inhibition steady-state fixation unit firing rates, so overall the blocked-input 2.0 model was preferred.

These results illustrate the way in which extra constraint from neural data can help to adjudicate between models that fit behaviour equally well but also underline the critical role of mapping assumptions being correct if that constraint is not to be misleading. For example, Boucher et al.’s (2007) conclusions of strong asymmetries between fixation and movement unit inhibition, speed of accumulation, and encoding times were all overturned by the corresponding estimates from the

blocked-input 2.0 model, where they were all similar. However, lateral inhibition was estimated as being weak and the time consumed to complete the external stop-control process substantial (60–80 ms), so the blocked-input 2.0 model does confirm that the race is effectively independent. These results also underline the important difference between the site of inhibition (e.g. the gaze-control network) and its source, something that Logan et al.'s (2015) research sheds little light on, except that in the case of saccadic countermanding, the source of inhibition likely does not coincide with the site at which it acts.

4.2 *The Stop-Signal Paradigm in Humans*

Logan et al. (2014) proposed an independent RD model of the stop-signal paradigm with one accumulator for each choice and one accumulator for the stop signal that inhibits responding if it wins. Mittner et al. (2014) used this “RD-SS” model to investigate mind-wandering that causes a failure to exert the control necessary to perform a task (see Hawkins et al. (2017), on the role of EAMs in research). Mittner et al. examined a paradigm using an auditory stop signal requiring easy visual (arrow direction) choices. After 10% of the trials, participants answered “thought probes”, rating the degree to which their thoughts had been on vs. off task during the previous trial. Mind-wandering has been associated with increased activity in the default mode network measured by fMRI (Christoff et al., 2009) and changes in pupil diameter (Smallwood et al., 2011). Consistent with these findings, Mittner et al. showed that a nonlinear support vector machine (Scholkopf & Smola, 2002) could reliably predict the probability of off-task ratings based on 28 features derived from fMRI and pupil diameter measurements. Mittner et al. then used the classifier they developed to predict the probability of on-task vs. off-task states in trials without thought probes and applied them to construct mixtures of two RD-SS models. The two RD-SS models had different parameters to represent on-task and off-task states. By comparing the complexity-adjusted fit of mixture models that allowed different parameters to vary between states, they were able to identify mechanisms through which failures of control due to mind-wandering affected behaviour in the stop-signal task. Specifically, the effects manifested in both the go and stop processes, with the rate parameter for the correct choice accumulator and the stop accumulator as well as the threshold parameter (which was assumed to be the same for all accumulators) being decreased during mind wandering.

The remaining two studies that we review linked the BEESTS model to neural measures of the time course of processing in the stop-signal task using a two-stage procedure. Skippen et al. (2020) used the BEESTS model developed by Matzke et al. (2019) to account for choice errors and provide estimates of both SSRT and trigger failures in a sample of more than 150 participants performing a stop-signal task based on number parity judgements. They found that the field-standard non-parametric integration method (Verbruggen et al., 2013, 2019) produced an average SSRT estimate (~200 ms) almost 50% larger than the BEESTS-based estimate

(~135 ms). While the two SSRT estimates were only moderately correlated, the non-parametric estimate was strongly correlated with trigger failures, suggesting that it confounds the speed of inhibition with attention-related factors (e.g. maintenance of a stop goal). Matzke et al. (2017a) studying schizophrenia, Weigard et al. (2019) studying attention deficit/hyperactivity disorder, Skippen et al. (2019) studying impulsivity, and Doekemeijer et al. (2021) studying reward effects also found reductions in SSRT when trigger failures are accounted for and similar confounding of inhibitory and attention effects.

Skippen et al. (2020) examined the correlation between individual differences in stop-signal performance and the amplitude and peak latency of fronto-central N1 and P3 event-related potentials. Both components have traditionally been related to stopping as they peak earlier and have a larger amplitude, for signal-inhibit than signal-respond trials. P3 has received more attention as an index of stopping because its timing aligns better with non-parametric SSRT, although it peaks still ~100 ms later (Wessel & Aron, 2015). However, when trigger failures are taken into account, N1 is much better temporally aligned. For signal-respond trials, N1 peak latency correlated with all behavioural stopping measures, and for participants with better stop-signal performance (i.e. fewer trigger failures/faster SSRT), the N1 peaked earlier for signal-inhibit than signal-respond trials. Skippen et al. interpreted their results in terms of Keneman's (2015) proposal of potentiation by right inferior frontal gyrus (rIFG) of a link between stop-signal detection and a reactive inhibition that is greater (as reflected in a faster and stronger N1) on successful stop trials (see Aron et al., 2014, for further discussion). P3 amplitude was also significantly correlated with all three measures of stopping behaviour, but given its time course, they suggested that this, and the larger and earlier P3 for successful stopping, reflected the relationship between P3 and the detection of rare events (Boehler et al., 2011; Waller et al., 2021).

The final study, Jana et al. (2020), identified the fine-grained time course of processing in the stop-signal task over five studies by combining information from EMG, TMS, and frequency-domain EEG. At ~120 ms following a stop signal, they found an increase in right frontal beta (13–30 Hz) EEG, consistent with rIFG being the source of inhibition. Shortly thereafter (~140 ms), they measured broad skeleton-motor suppression with single pulse TMS applied to the motor cortex, which they attributed to the impact of the StN on basal ganglia output. They then used EMG to define “cancel time” as the peak of partial myographic bursts that occurred on signal-inhibit trials (i.e. trials on which muscle contraction occurred but was not sufficient to cause a button press). Average cancel time across several studies ranged from 146 ms to 166 ms, substantially shorter than the average non-parametric SSRT estimate of ~220 ms. Given that Jana et al. used an auditory stop signal, which is registered ~30 ms earlier (Ramautar et al., 2006) than the visual stop signals used by Skippen et al. (2019, 2020), cancel time aligns quite well with their SSRT estimates that take account of trigger failures (~135 ms). Cancel time also aligns well with trigger-failure-corrected SSRT estimates obtained by Matzke et al. (2017a) for control participants in two experiments using an auditory stop signal (141 ms and 160 ms). However, when Jana et al. applied Matzke et al.'s (2017b) BEESTS model to their first two studies, they found a very low trigger-failure rate of 4% compared to

Skippen et al.'s (2019) 22% and Matzke et al.'s (2017a) ~9% and a longer BEESTS-based SSRT estimate than cancel time (~205 ms). Still, trigger failure rate and SSRT were both significantly positively correlated with cancel time. They concluded that their results are compatible with a race model with an interaction between the stop and the go processes (as measured by TMS applied to motor cortex) happening very late, in the last 20 ms of cancel time.

4.3 Problems with Modelling Unobserved “Responses”

Estimating the time course of action withholding based on purely behavioural data is difficult; when inhibition succeeds, the time course cannot be directly observed, and when it fails, observed signal-respond RT provides a biased estimate as they omit the faster times associated with successful inhibition. Because of this, Logan (1994) observed that “there is no avoiding the race model” (p. 207), which enables the bias to be corrected. Although it seems natural to characterise the racers with evidence-accumulation processes, recent work by Matzke et al. (2020) has thrown doubt on the identifiability of parameters in such models. In an extensive simulation study, they showed that for the RD-SS model and for a simpler model based on a lognormal-race evidence-accumulation process (Heathcote & Love, 2012), estimation of the non-decision time for the stop accumulator is problematic. Even with the very large number of trials per participant collected in Logan et al.'s (2014) original study (over 3000), non-decision time was greatly overestimated and the variability in the finishing time of the stop runner grossly underestimated. Further, because of the highly correlated nature of parameters in EAMs, these problems caused distortions in other model parameters.

Fixing non-decision time to be the same for all accumulators can address the problem. However, this assumption is unlikely to be very accurate, because go non-decision time likely includes a greater response production component depending on what proportion of that period is ballistic (i.e. cannot be halted by the stop process). This is particularly so in paradigms with an auditory stop signal and visual go stimulus, as the former is encoded more quickly than the latter (Ramautar et al., 2006). Matzke et al.'s (2020) results cast doubt on the parameter values reported by Logan et al. (2014) and Mittner et al. (2014), who both estimated separate non-decision times for stop and go racers, but they do not necessarily mean that their primary analyses using model selection are invalid. Matzke et al.'s results also potentially raise concerns about Boucher et al.'s (2007) and Logan et al.'s (2015) parameter estimates, as they did not use neural measures directly in their fitting, potentially explaining the inconsistency between their results. They both found that any period of interaction in their models is brief, and so their models may well act similarly to the RD-SS model studied by Matzke et al. Estimation may even have a more severe problem as the greater complexity of the Boucher et al. and Logan et al. models is likely to already make identifiability challenging.

Jana et al. (2020) suggested that their cancel time measure could be used instead of a model-based SSRT estimate. Raud and Huster (2017) made the same suggestion

in the context of a task where the go response requires a button press with both hands with the stop signal indicating that one or the other button press should be selectively withheld. There are two problems with this suggestion. First, it suffers from the same overestimation bias as signal-respond RT, although likely less so. This is because the fact that partial EMG bursts occur only on some stop-signal trials (e.g. 47% in Raud and Huster) indicates that true SSRT is variable, and so cancel time ignores the cases where it is fast enough to suppress any EMG activity. Second, forsaking modelling also foregoes the benefits associated with accounting for attentional goal-neglect influences on the control of action withholding by estimating trigger failures. Hence, we believe that there is an opportunity for model-based neuroscience based on joint-modelling, as we discuss in the final opportunity box.

Opportunity 3: Further Exploiting Control of Action Withholding

Although the BEESTS model enables measurement of many important quantities for action withholding research, EAM-based models enable more theoretically constrained modelling. A partial solution is provided by a hybrid architecture, with racing diffusion processes to describe performance on the go task and the ex-Gaussian distribution to describe the finishing times of the stop racer. Tanis et al. (2024) show that this “hybrid” model is identifiable, even in complex designs. Weigard et al. (2023) applied a similar hybrid model to a large neuro-imaging study (Casey et al., 2018) with an unfortunate design choice where the go choice stimulus was replaced by the stop signal. This caused the amount of information on which the go choice was based to vary with SSD, resulting in severe violations of go context independence. This feature was easily handled with the hybrid model by allowing the rate of go accumulation to change as a function of SSD. Joint modelling with a temporally precise neural measure could support a more complete solution through providing constraints on the estimate of stop non-decision time. For example, Nunez et al. (2019) used the N2 component of EEG to measure the completion of visual encoding and the onset of evidence accumulation, an approach that may provide the first step towards constraining the non-decision time of the stop runner.

5 Discussion

We have reviewed previous research on the cognitive control of different aspects of actions, either the time at which an action is made or which action is made or withheld. Clearly, the difficulty of modelling increases in the latter case because

some temporal information is not directly available (i.e. there is no RT for signal-inhibit trials). Even when modelling only observed actions, more complex models can be difficult to estimate. A key methodological lesson from the foregoing review is that any new model should be accompanied by simulation studies to determine estimation performance in the design to which it is to be applied. Such model and parameter recovery studies determine, respectively, if the model can be reliably selected from among competing models and how well the model's parameters can be estimated (Heathcote et al., 2015a).

The work we reviewed took a variety of different cognitive models, ranging from tractable models with easily computed likelihoods to intractable models that require simulation methods to approximate their likelihood and different approaches to linking these models to neural data, ranging from two-stage analyses to full joint modelling. Weichert et al.'s (2020) results, and those of Turner et al. (2019), illustrate the freedom afforded by approximate likelihood methods to explore complex models but also the challenges this brings. We agree with Turner et al. (2018) that more tractable models are not necessarily less flexible and that models should not be selected based on their tractability but rather on the way they navigate the trade-off between goodness-of-fit and flexibility. However, at present, the intractable models of cognitive control that have been developed have not been able to access the advantages of joint modelling.

Weichert et al. (2020) were at least able to use fully Bayesian estimation and so base model selection on a criterion that takes account of posterior uncertainty, but for both them and Turner et al. (2019), less than ideal methods of generating individual trial model estimates were needed to link to neural data, indicating an important direction for future methodological developments. Weichert et al.'s results, and those of Turner et al. (2018), also illustrate another key methodological point: constraining models with a comprehensive characterisation of behaviour by simultaneously accounting for both choices and RT distributions. In the case of Weichert et al., it was RT distributions that best differentiated the LCA and FFI models, although the variation of input strength in the experimental design played a key role. For Turner et al., their conclusion in favour of LCA-based models, which was based on only choice data, was later reversed by Evans et al. (2019) when RT distributions were also taken into account. Similarly, the work of Logan et al. (2015) demonstrates that extra neural data can help to discriminate between models that otherwise give an equivalently good account of behaviour. Combined with Turner et al.'s (2016) demonstration of the estimation advantages of accounting for more than one type of neural data in a joint model, and the demonstration by Jana et al. (2020) of the utility of multi-modal electrophysiological data in identifying the precise time-course of processing, it seems certain that our understanding of cognitive control will benefit from the increasing use of joint modelling to integrate diverse sources of data.

A.1 Exercises

The annotated R code available on <https://osf.io/zs245/> enables the reader to collect data from a Flanker task and fit standard EAMs (lognormal race, Wiener diffusion, diffusion decision, LBA, and RD models) to these data. The accompanying exercises illustrate how to evaluate the models' ability to account for the pattern of data from a conflict task.

B.1 Recommended Reading

- The stages of control: Verbruggen, McLaren, and Chambers (2014).
- An overview and comparison of different EAMs: Donkin & Brown (2018).
- Racing diffusion models: Tillman, Van Zandt, and Logan (2020).
- Bayesian parameter recovery for EAMs and BEESTS: Heathcote et al. (2018).

References

- Alexander, W. H., & Brown, J. W. (2015). Hierarchical error representation - a computational model of anterior cingulate and dorsolateral prefrontal cortex. *Neural Computation*, 27(11), 2354–2410.
- Aron, A. R., Robbins, T. W., & Poldrack, R. A. (2014). Inhibition and the right inferior frontal cortex: One decade on. *Trends in Cognitive Sciences*, 18(4), 177–185.
- Boehler, C. N., Appelbaum, L. G., Krebs, R. M., Chen, L.-C., & Woldorff, M. G. (2011). The role of stimulus salience and attentional capture across the neural hierarchy in a stop-signal task. *PLoS One*, 6, e26386.
- Boehm, U., van Maanen, L., Evans, N. J., Brown, S. D., & Wagenmakers, E.-J. (2020). A theoretical analysis of the reward rate optimality of collapsing decision criteria. *Attention, Perception, & Psychophysics*, 82, 1520–1534.
- Botvinick, M. M., Niv, Y., & Barto, A. G. (2009). Hierarchically organized behavior and its neural foundations: A reinforcement learning perspective. *Cognition*, 113(3), 262–280.
- Boucher, L., Palmeri, T. J., Logan, G. D., & Schall, J. D. (2007). Inhibitory control in mind and brain: An interactive race model of countermanding saccades. *Psychological Review*, 114, 376–397.
- Braver, T. S. (2012). The variable nature of cognitive control: A dual mechanisms framework. *Trends in Cognitive Sciences*, 16, 106–113.
- Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice response time: Linear Ballistic Accumulation. *Cognitive Psychology*, 57, 153–178.
- Burgess, P. W., Gonen-Yaacovi, G., & Volle, E. (2011). Functional neuroimaging studies of prospective memory: What have we learnt so far? *Neuropsychologia*, 49, 2246–2257.
- Carpenter, R., & Williams, M. (1995). Neural computation of log likelihood in control of saccadic eye movements. *Nature*, 377, 59–62.
- Casey, B. J., Cannonier, T., Conley, M. I., Cohen, A. O., Barch, D. M., Heitzeg, M. M., et al. (2018). The Adolescent Brain Cognitive Development (ABCD) study: Imaging acquisition across 21 sites. *Developmental Cognitive Neuroscience*, 32, 43–54.

- Christoff, K., Gordon, A. M., Smallwood, J., Smith, R., & Schooler, J. W. (2009). Experience sampling during fMRI reveals default network and executive system contributions to mind wandering. *Proceedings of the National Academy of Sciences*, *106*, 8719–8724.
- Churchland, A. K., Kiani, R., & Shadlen, M. N. (2008). Decision-making with multiple alternatives. *Nature Neuroscience*, *11*, 693–702.
- Dai, J., & Busemeyer, J. R. (2014). A probabilistic, dynamic, and attribute-wise model of intertemporal choice. *Journal of Experimental Psychology: General*, *143*(4), 1489–1514.
- Dalley, J. W., & Robbins, T. W. (2017). Fractionating impulsivity: Neuropsychiatric implications. *Nature Reviews Neuroscience*, *18*(3), 158–171.
- Dayan, P. (2008). Simple substrates for complex cognition. *Frontiers in Neuroscience*, *2*(2), 255–263.
- Doekemeijer, R. A., Verbruggen, F., & Boehler, C. N. (2021). Face the (trigger) failure: Trigger failures strongly drive the effect of reward on response inhibition. *Cortex*, *139*, 166–177.
- Donkin, C., & Brown, S. D. (2018). Response times and decision-making. In E.-J. Wagenmakers & J. T. Wixted (Eds.), *Stevens' handbook of experimental psychology and cognitive neuroscience, Volume five: Methodology* (4th ed.). Wiley.
- Duncan, J. (2006). EPS Mid-Career Award 2004: Brain mechanisms of attention. *Quarterly Journal of Experimental Psychology*, *59*(1), 2–27.
- Eichele, T., Debener, S., Calhoun, V. D., Specht, K., Engel, A. K., Hugdahl, K., von Cramon, D. Y., & Ullsperger, M. (2008). Prediction of human errors by maladaptive changes in event-related brain network. *Proceedings of the National Academy of Sciences, National Academy of Sciences*, *105*, 6173–6178.
- Eriksen, C. W. (1995). The flankers task and response competition: A useful tool for investigating a variety of cognitive problems. *Visual Cognition*, *2*, 101–118.
- Evans, N. J., Holmes, W. R., & Trueblood, J. S. (2019). Response-time data provide critical constraints on dynamic models of multi-alternative, multi-attribute choice. *Psychonomic Bulletin & Review*, *26*, 901–933.
- 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*, 316–326.
- Figner, B., Knoch, D., Johnson, E. J., Krosch, A. R., Lisanby, S. H., Fehr, E., & Weber, E. U. (2010). Lateral prefrontal cortex and self-control in intertemporal choice. *Nature Neuroscience*, *13*, 538–539.
- Forstmann, B. U., Dutilh, G., Brown, S., Neumann, J., Cramon Von, D. Y., Ridderinkhof, K. R., & Wagenmakers, E.-J. (2008). Striatum and pre-SMA facilitate decision-making under time pressure. *Proceedings of the National Academy of Sciences of the United States of America*, *105*, 17538–17542.
- Frank, M. J. (2006). Hold your horses: A dynamic computational role for the subthalamic nucleus in decision making. *Neural Networks*, *19*, 1120–1136.
- Freund, M. C., Etzel, J. A., & Braver, T. S. (2021). Neural coding of cognitive control: The representational similarity analysis approach. *Trends in Cognitive Sciences*, *25*, 622–638.
- Gelman, A., & Hill, J. (2007). *Data analysis using regression and multilevel/ hierarchical models*. Cambridge University Press.
- Gómez, P., Ratcliff, R., & Perea, M. (2007). A model of the go/no-go task. *Journal of Experimental Psychology: General*, *136*, 389–413.
- Hare, T. A., Camerer, C. F., & Rangel, A. (2009). Self-control in decision-making involves modulation of the vmPFC valuation system. *Science*, *324*, 646–648.
- Hawkins, G. E., & Heathcote, A. (2021). Racing against the clock: Evidence-based vs. time-based decisions. *Psychological Review*, *128*, 222–263.
- Hawkins, G. E., Forstmann, B. U., Wagenmakers, E.-J., Ratcliff, R., & Brown, S. D. (2015). Revisiting the evidence for collapsing boundaries and urgency signals in perceptual decision-making. *The Journal of Neuroscience*, *35*, 2476–2484.
- Hawkins, G. E., Mittner, M., Forstmann, B. U., & Heathcote, A. (2017). On the efficiency of neurally-informed cognitive models to identify latent cognitive states. *Journal of Mathematical Psychology*, *76*, 142–155.

- Hawkins, G. E., Mittner, M., Forstmann, B. U., & Heathcote, A. (2019). Modeling distracted performance. *Cognitive Psychology*, *112*, 48–80.
- Heathcote, A. (2004). Fitting Wald and ex-Wald distributions to response time data: An example using functions for the S-PLUS package. *Behavior Research Methods*, *36*, 678–694.
- Heathcote, A., & Love, J. (2012). Linear deterministic accumulator models of simple choice. *Frontiers in Psychology: Cognitive Science*, *3*, 292.
- Heathcote, A., & Matzke, D. (2022). Winner takes all! What are race models, and why and how should psychologists use them? *Current Directions in Psychological Science*, *31*, 383–394.
- Heathcote, A., Brown, S. D., & Wagenmakers, E.-J. (2015a). An introduction to good practices in cognitive modeling. In B. U. Forstmann & E.-J. Wagenmakers (Eds.), *An introduction to model-based cognitive neuroscience* (pp. 25–48). Springer.
- Heathcote, A., Loft, S., & Remington, R. W. (2015b). Slow down and remember to remember! A delay theory of prospective memory costs. *Psychological Review*, *122*, 367–410.
- Heathcote, A., Lin, Y.-S., Reynolds, A., Strickland, L., Gretton, M., & Matzke, D. (2018). Dynamic models of choice. *Behavior Research Methods*, *51*, 961–985.
- Hommel, B. (2011). The Simon effect as tool and heuristic. *Acta Psychologica*, *136*, 189–202.
- Jana, S., Hannah, R., Muralidharan, V., & Aron, A. R. (2020). Temporal cascade of frontal, motor and muscle processes underlying human action-stopping. *eLife*, *9*, e50371.
- Kenemans, J. L. (2015). Specific proactive and generic reactive inhibition. *Neuroscience Biobehavioral Review*, *56*, 115–126.
- Kvam, P. D., Romeu, R. J., Turner, B. M., Vassileva, J., & Busemeyer, J. R. (2021). Testing the factor structure underlying behavior using joint cognitive models: Impulsivity in delay discounting and Cambridge gambling tasks. *Psychological Methods*, *26*, 18–37.
- Lamichhane, B., McDaniel, M. A., Waldum, E. R., & Braver, T. S. (2018). Age-related changes in neural mechanisms of prospective memory. *Cognitive, Affective, & Behavioral Neuroscience*, *18*, 982–999.
- Lee, M. D. (2011). How cognitive modeling can benefit from hierarchical Bayesian models. *Journal of Mathematical Psychology*, *55*, 1–7.
- Lerche, V., & Voss, A. (2018). Speed–accuracy manipulations and diffusion modeling: Lack of discriminant validity of the manipulation or of the parameter estimates? *Behavior Research Methods*, *50*(6), 2568–2585.
- Lo, C.-C., Boucher, L., Paré, M., Schall, J. D., & Wang, X.-J. (2009). Proactive inhibitory control and attractor dynamics in countermanding action: A spiking neural circuit model. *The Journal of Neuroscience*, *29*, 9059–9071.
- Logan, G. D. (1994). On the ability to inhibit thought and action: A users' guide to the stop signal paradigm. In D. Dagenbach & T. H. Carr (Eds.), *Inhibitory processes in attention, memory, and language* (pp. 189–239). Academic Press.
- Logan, G. D., & Cowan, W. B. (1984). On the ability to inhibit thought and action: A theory of an act of control. *Psychological Review*, *91*, 295–327.
- Logan, G. D., Van Zandt, T., Verbruggen, F., & Wagenmakers, E.-J. (2014). On the ability to inhibit thought and action: General and special theories of an act of control. *Psychological Review*, *121*, 66–95.
- Logan, G. D., Yamaguchi, M., Schall, J. D., & Palmeri, T. J. (2015). Inhibitory control in mind and brain 2.0: Blocked-input models of saccadic countermanding. *Psychological Review*, *122*, 115–147.
- Ly, A., Boehm, U., Heathcote, A., Turner, B. M., Forstmann, B., Marsman, M., & Matzke, D. (2018). A flexible and efficient hierarchical Bayesian approach to the exploration of individual differences in cognitive-model-based neuroscience. In A. A. Moustafa (Ed.), *Computational models of brain and behavior* (pp. 467–480). Wiley Blackwell.
- MacLeod, C. M. (1991). Half a century of research on the Stroop effect: An integrative review. *Psychological Bulletin*, *109*, 163–203.
- Matzke, D., & Wagenmakers, E.-J. (2009). Psychological interpretation of the ex-Gaussian and shifted Wald parameters: A diffusion model analysis. *Psychonomic Bulletin & Review*, *16*, 798–817.

- Matzke, D., Dolan, C. V., Logan, G. D., Brown, S. D., & Wagenmakers, E.-J. (2013). Bayesian parametric estimation of stop-signal reaction time distributions. *Journal of Experimental Psychology: General*, *142*, 1047–1073.
- Matzke, D., Hughes, M., Badcock, J. C., Michie, P., & Heathcote, A. (2017a). Failures of cognitive control or attention? The case of stop-signal deficits in schizophrenia. *Attention, Perception, & Psychophysics*, *79*, 1078–1086.
- Matzke, D., Love, J., & Heathcote, A. (2017b). A Bayesian approach for estimating the probability of trigger failures in the stop-signal paradigm. *Behavior Research Methods*, *49*, 267–281.
- Matzke, D., Ly, A., Selker, R., Weeda, W. D., Scheibehenne, B., Lee, M. D., & Wagenmakers, E.-J. (2017c). Bayesian inference for correlations in the presence of estimation uncertainty and measurement error. *Collabra: Psychology*, *3*(1), 25.
- Matzke, D., Verbruggen, F., & Logan, G. (2018). The stop-signal paradigm. In E.-J. Wagenmakers & J. T. Wixted (Eds.), *Stevens' handbook of experimental psychology and cognitive neuroscience, Volume five: Methodology* (4th ed., pp. 383–427). Wiley.
- Matzke, D., Curley, S., Gong, C. Q., & Heathcote, A. (2019). Inhibiting responses to difficult choices. *Journal of Experimental Psychology: General*, *148*, 124–142.
- Matzke, D., Logan, G. D., & Heathcote, A. (2020). A cautionary note on evidence-accumulation models of response inhibition in the stop-signal paradigm. *Computational Brain & Behavior*, *3*, 269–288.
- Matzke, D., Strickland, L., Sripada, C., Weigard, A., Puri, R., He, J. L., Hirst, R. J., & Heathcote, A. (submitted). *Stopping timed actions*. Manuscript submitted for publication. <https://psyarxiv.com/9h3v7/>
- Merchant, H., Harrington, D. L., & Meck, W. H. (2013). Neural basis of the perception and estimation of time. *Annual Review of Neuroscience*, *36*, 313–336.
- 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.
- Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, *24*, 167–202.
- Mittner, M., Boekel, W., Tucker, A. M., Turner, B. M., Heathcote, A., & Forstmann, B. U. (2014). When the brain takes a break: A model-based analysis of mind wandering. *Journal of Neuroscience*, *34*, 16286–16295.
- Navarro, D. J., & Fuss, I. G. (2009). Fast and accurate calculations for first-passage times in Wiener diffusion models. *Journal of Mathematical Psychology*, *53*, 222–230.
- Nigg, J. T. (2017). Annual Research Review: On the relations among self-regulation, self-control, executive functioning, effortful control, cognitive control, impulsivity, risk-taking, and inhibition for developmental psychopathology. *Journal of Child Psychology and Psychiatry*, *58*, 361–383.
- Nunez, M. D., Gosai, A., Vandekerckhove, J., & Srinivasan, R. (2019). The latency of a visual evoked potential tracks the onset of decision making. *NeuroImage*, *197*, 93–108.
- O'Reilly, R. C., Herd, S. A., & Pauli, W. M. (2010). Computational models of cognitive control. *Current Opinion in Neurobiology*, *20*(2), 257–261.
- Palestro, J. J., Bahg, G., Sederberg, P. B., Lu, Z.-L., Steyvers, M., & Turner, B. M. (2018a). A tutorial on joint models of neural and behavioral measures of cognition. *Journal of Mathematical Psychology*, *84*, 20–48.
- Palestro, J. J., Weichart, E., Sederberg, P. B., & Turner, B. M. (2018b). Some task demands induce collapsing bounds: Evidence from a behavioral analysis. *Psychonomic Bulletin & Review*, *25*, 1225–1248.
- 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.
- Ramautar, J. R., Kok, A., & Ridderinkhof, K. R. (2006). Effects of stop-signal modality on the N2/P3 complex elicited in the stop-signal paradigm. *Biological Psychology*, *72*, 96–109.
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, *20*, 873–922.

- Raud, L., & Huster, R. J. (2017). The temporal dynamics of response inhibition and their modulation by cognitive control. *Brain Topography*, *30*(4), 486–501.
- Rodriguez, C. A., Turner, B. M., & McClure, S. M. (2014). Intertemporal choice as discounted value accumulation. *PLoS One*, *9*, e90138.
- Rodriguez, C. A., Turner, B. M., Van Zandt, T., & McClure, S. M. (2015). The neural basis of value accumulation in intertemporal choice. *The European Journal of Neuroscience*, *42*, 2179–2189.
- Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychological Review*, *108*, 370–392.
- Rougier, N. P., Noelle, D. C., Braver, T. S., Cohen, J. D., & O'Reilly, R. C. (2005). Prefrontal cortex and flexible cognitive control: Rules without symbols. *Proceedings of the National Academy of Sciences of the United States of America*, *102*(20), 7338–7343.
- Scholkopf, B., & Smola, A. (2002). *Learning with kernels: Support vector machines, regularization, optimization, and beyond adaptive computation and machine learning*. MIT.
- Sebastian, A., Forstmann, B. U., & Matzke, D. (2018). Towards a model-based cognitive neuroscience of stopping – A neuroimaging perspective. *Neuroscience & Biobehavioral Reviews*, *90*, 130–136.
- Servant, M., Montagnini, A., & Burle, B. (2014). Conflict tasks and the diffusion framework: Insight in model constraints based on psychological laws. *Cognitive Psychology*, *72*, 162–195.
- Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychological Review*, *84*, 127–190.
- Simen, P., Vlasov, K., & Papadakis, S. (2016). Scale (in)variance in a unified diffusion model of decision making and timing. *Psychological Review*, *123*, 151–181.
- Skippen, P., Matzke, D., Heathcote, A., Fulham, W. R., Michie, P., & Karayanidis, F. (2019). Reliability of triggering inhibitory process is a better predictor of impulsivity than SSRT. *Acta Psychologica*, *192*, 104–117.
- Skippen, P., Fulham, W. R., Michie, P. T., Matzke, D., Heathcote, A., & Karayanidis, F. (2020). Reconsidering electrophysiological markers of response inhibition in light of trigger failures in the stop-signal task. *Psychophysiology*, *57*, e13619.
- Smallwood, J., Brown, K. S., Tipper, C., Giesbrecht, B., Franklin, M. S., Mrazek, M. D., Carlson, J. M., & Schooler, J. W. (2011). Pupillometric evidence for the decoupling of attention from perceptual input during offline thought. *PLoS One*, *6*, e18298.
- Stone, M. (1960). Models for choice-reaction time. *Psychometrika*, *25*, 251–260.
- Strickland, L., Loft, S., Remington, R. W., & Heathcote, A. (2018). Racing to remember: A theory of decision control in event-based prospective memory. *Psychological Review*, *125*, 851–887.
- Tanis, C. C., Heathcote, A., Zrubka, M., & Matzke, D. (2024). A hybrid approach to dynamic cognitive psychometrics. *Behavior Research Methods*.
- Tillman, G., Van Zandt, T., & Logan, G. D. (2020). Sequential sampling models without random between-trial variability: The racing diffusion model of speeded decision making. *Psychonomic Bulletin & Review*, *27*, 911–936.
- Trueblood, J. S., Heathcote, A., Evans, N., & Holmes, W. R. (2021). Urgency, leakage, and the relative nature of information processing in decision-making. *Psychological Review*, *128*, 160–186.
- Turner, B. M., & Sederberg, P. B. (2014). A generalized, likelihood-free method for posterior estimation. *Psychonomic Bulletin & Review*, *21*, 227–250.
- Turner, B. M., van Maanen, L., & Forstmann, B. U. (2015). Informing cognitive abstractions through neuroimaging: The neural drift diffusion model. *Psychological Review*, *122*, 312–336.
- Turner, B. M., Rodriguez, C. A., Norcia, T., Steyvers, M., & McClure, S. M. (2016). Why more is better: A method for simultaneously modeling EEG, fMRI, and behavior. *NeuroImage*, *128*, 96–115.
- Turner, B. M., Forstmann, B. U., Love, B. C., Palmeri, T. J., & van Maanen, L. (2017a). Approaches to analysis in model-based cognitive neuroscience. *Journal of Mathematical Psychology*, *76*, 65–79.

- Turner, B. M., Wang, T., & Merkle, E. C. (2017b). Factor analysis linking functions for simultaneously modeling neural and behavioral data. *NeuroImage*, *153*, 28–48.
- Turner, B. M., Schley, D. R., Muller, C., & Tsetsos, K. (2018). Competing theories of multialternative, multiattribute preferential choice. *Psychological Review*, *125*, 329–362.
- Turner, B. M., Rodriguez, C. A., Liu, Q., Molloy, M. F., Hoogendijk, M., & McClure, S. M. (2019). On the neural and mechanistic bases of self-control. *Cerebral Cortex*, *29*, 732–750.
- Ulrich, R., Schröter, H., Leuthold, H., & Birngruber, T. (2015). Automatic and controlled stimulus processing in conflict tasks: Superimposed diffusion processes and delta functions. *Cognitive Psychology*, *78*, 148–174.
- Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, *108*, 550–592.
- van Maanen, L., Brown, S. D., Eichele, T., Wagenmakers, E. J., Ho, T., & Serences, J. (2011). Neural correlates of trial-to-trial fluctuations in response caution. *Journal of Neuroscience*, *31*, 17488–17495.
- Vandekerckhove, J., Tuerlinckx, F., & Lee, M. D. (2011). Hierarchical diffusion models for two-choice response time. *Psychological Methods*, *16*, 44–62.
- Verbruggen, F., & Logan, G. D. (2008). Automatic and controlled response inhibition: Associative learning in the go/no-go and stop-signal paradigms. *Journal of Experimental Psychology: General*, *137*(4), 649–672.
- Verbruggen, F., Chambers, C. D., & Logan, G. D. (2013). Fictitious inhibitory differences: How skewness and slowing distort the estimation of stopping latencies. *Psychological Science*, *24*, 352–362.
- Verbruggen, F., McLaren, I. P. L., & Chambers, C. D. (2014). Banishing the control homunculi in studies of action control and behavior change. *Perspectives on Psychological Science*, *9*, 497–524.
- Verbruggen, F., Aron, A. R., Band, G. P. H., Beste, C., Bissett, P. G., Brockett, A. T., et al. (2019). A consensus guide to capturing the ability to inhibit actions and impulsive behaviors in the stop-signal task. *eLife*, *8*, e46323.
- Waller, D. A., Hazeltine, E., & Wessel, J. R. (2021). Common neural processes during action-stopping and infrequent stimulus detection: The frontocentral P3 as an index of generic motor inhibition. *International Journal of Psychophysiology*, *163*, 11–21.
- Weichart, E. R., Turner, B. M., & Sederberg, P. B. (2020). A model of dynamic, within-trial conflict resolution for decision making. *Psychological Review*, *127*, 749–777.
- Weigard, A., Heathcote, A., Matzke, D., & Huang-Pollock, C. (2019). Cognitive modeling suggests that attentional failures drive longer stop-signal reaction time estimates in attention deficit/hyperactivity disorder. *Clinical Psychological Science*, *7*, 856–872.
- Weigard, A., Matzke, D., Tanis, C., & Heathcote, A. (2023). Cognitive process modeling addresses context independence violations in the ABCD study stop-signal task. *Developmental Cognitive Neuroscience*, *59*, 101191.
- Wessel, J. R., & Aron, A. R. (2015). It's not too late: The onset of the frontocentral P3 indexes successful response inhibition in the stop-signal paradigm. *Psychophysiology*, *52*(4), 472–480.
- White, C. N., Ratcliff, R., & Starns, J. J. (2011). Diffusion models of the flanker task: Discrete versus gradual attentional selection. *Cognitive Psychology*, *63*, 210–238.
- Wiecki, T. V., & Frank, M. J. (2013). A computational model of inhibitory control in frontal cortex and basal ganglia. *Psychological Review*, *120*(2), 329–355.
- Wong-Lin, K., Eckhoff, P., Holmes, P., & Cohen, J. D. (2010). Optimal performance in a countermanding saccade task. *Brain Research*, *1318*, 178–187.