Received: 24 October 2018 | Revised: 25 March 2019 | Accepted: 6 May 2019

DOI: 10.1111/ejn.14435

RESEARCH REPORT

EJN European Journal of Neuroscience FENS WILEY

Switch costs in inhibitory control and voluntary behaviour: A computational study of the antisaccade task

Eduardo A. Aponte[1] | **Klaas E. Stephan**[1,2,3] | **Jakob Heinzle**[1]

[1]Translational Neuromodeling Unit (TNU), Institute for Biomedical Engineering, University of Zurich and ETH Zurich, Zurich, Switzerland

[2]Wellcome Centre for Human Neuroimaging, University College London, London, UK

[3]Max Planck Institute for Metabolism Research, Cologne, Germany

Correspondence
Eduardo A. Aponte, Translational Neuromodeling Unit (TNU), Institute for Biomedical Engineering, University of Zurich and ETH Zurich, Wilfriedstrasse 6, 8032 Zurich, Switzerland.
Email: aponte@biomed.ee.ethz.ch

Funding information
Universität Zürich; The René and Susanne Braginsky Foundation

Abstract

An integral aspect of human cognition is the ability to inhibit stimulus-driven, habitual responses, in favour of complex, voluntary actions. In addition, humans can also alternate between different tasks. This comes at the cost of degraded performance when compared to repeating the same task, a phenomenon called the "task-switch cost." While task switching and inhibitory control have been studied extensively, the interaction between them has received relatively little attention. Here, we used the SERIA model, a computational model of antisaccade behaviour, to draw a bridge between them. We investigated task switching in two versions of the mixed antisaccade task, in which participants are cued to saccade either in the same or in the opposite direction to a peripheral stimulus. SERIA revealed that stopping a habitual action leads to increased inhibitory control that persists onto the next trial, independently of the upcoming trial type. Moreover, switching between tasks induces slower and less accurate voluntary responses compared to repeat trials. However, this only occurs when participants lack the time to prepare the correct response. Altogether, SERIA demonstrates that there is a reconfiguration cost associated with switching between voluntary actions. In addition, the enhanced inhibition that follows antisaccade but not prosaccade trials explains asymmetric switch costs. In conclusion, SERIA offers a novel model of task switching that unifies previous theoretical accounts by distinguishing between inhibitory control and voluntary action generation and could help explain similar phenomena in paradigms beyond the antisaccade task.

KEYWORDS

antisaccades, response inhibition, SERIA model, switch costs, voluntary control

1 | INTRODUCTION

A hallmark of higher order cognition is the ability to alternate between different tasks and their corresponding actions, as well as between habitual and non-habitual responses (Isoda & Hikosaka, 2008). However, alternating between different tasks engenders reaction time (RT) and error rate (ER) switch costs (Kiesel et al., 2010). While inhibitory control of habitual actions (Aron, 2011) and flexible action selection (Monsell, 2003) have been investigated in great detail, the interplay between them has received much less attention (but see Hikosaka & Isoda, 2010). In particular, great effort has been devoted to developing computational models of action inhibition (Schall, Palmeri, & Logan, 2017) and task switching (Karayanidis et al., 2010;

Abbreviations: ER, Error rate; LME, Log model evidence; PP, Prosaccade trial probability; RT, Reaction time; TT, Trial type; WAIC, Watanabe–Akaike Information criterion.

Edited by Panayiota Poirazi. Reviewed by Balazs Hangya and Donatas Jonikaiti.

All peer review communications can be found with the online version of the article.

Schmitz & Voss, 2014). However, models which include both are still rare.

The antisaccade task (Hallett, 1978; Munoz & Everling, 2004) is an attractive experimental paradigm to study the above phenomena. In it, a habitual response – a (pro)saccade to a peripheral stimulus – needs to be overwritten by a non-habitual, voluntary action, that is, an (anti)saccade in the direction opposite to the stimulus. Behaviourally, switch costs in the mixed antisaccade task, in which pro- and antisaccade trials are alternated, have been studied in detail (Ansari, Derakshan, & Richards, 2008; Barton, Greenzang, Hefter, Edelman, & Manoach, 2006; Barton, Raoof, Jameel, & Manoach, 2006; Barton et al., 2002; Bojko, Kramer, & Peterson, 2004; Chan, Koval, Johnston, & Everling, 2017; Cherkasova, Manoach, Intriligator, & Barton, 2002; DeSimone, Weiler, Aber, & Heath, 2014; Ethridge, Brahmbhatt, Gao, McDowell, & Clementz, 2009; Fecteau, Au, Armstrong, & Munoz, 2004; Franke, Reuter, Breddin, & Kathmann, 2009; Heath, Gillen, & Samani, 2016; Heath, Starrs, Macpherson, & Weiler, 2015; Hunt & Klein, 2002; Lee, Hamalainen, Dyckman, Barton, & Manoach, 2011; Manoach, Lindgren, & Barton, 2004; Manoach et al., 2002, 2007; Mueller, Swainson, & Jackson, 2009; Pierce, McCardel, & McDowell, 2015; Rivaud-Pechoux, Vidailhet, Brandel, & Gaymard, 2007; Weiler & Heath, 2012a,2012b, 2014a,2014b). Despite the large number of studies, no unified picture of the cost of switching in this paradigm has emerged. Specifically, most studies indicate that switch prosaccades (i.e., correct prosaccades that follow an antisaccade trial) are slower than repeat prosaccades. However, while some studies indicate that switch antisaccades are faster than repeat antisaccades (e.g., Cherkasova et al., 2002), others report the opposite effect (e.g., Barton, Greenzang, et al., 2006), and yet others indicate no effect (e.g., Weiler & Heath, 2012a).

Two main explanations for switch costs in this task have been proposed. According to the task-set reconfiguration hypothesis (Rogers & Monsell, 1995; for an overview see Barton, Greenzang, et al., 2006), switch trials require the active reconfiguration of the task-set relevant to the new trial. This process is assumed to be an act of endogenous control that is not necessary in repeat trials, is time-consuming and can be prepared in advance of the peripheral stimulus. While intuitively appealing, this hypothesis is at odds with the observation that switch antisaccades are sometimes *faster* than repeat antisaccades (Cherkasova et al., 2002).

By contrast, the task-inertia hypothesis (Allport, Styles, & Hsieh, 1994; see Barton, Raoof, et al., 2006; Weiler, Hassall, Krigolson, & Heath, 2015) postulates that passive interference caused by non-dominant rules (antisaccades) leads to pro- but not antisaccade RT switch costs. In other words, antisaccades require the activation of a "non-dominant" rule, which interferes with prosaccades on the next trial. Because prosaccades are the "dominant" rule, no interference occurs after this task-set is activated. Again, this hypothesis is at odds with positive switch costs in switch pro- and antisaccade trials. In summary, none of these hypotheses offers a unified explanation of the conflicting findings in the antisaccade task.

One approach to reconcile conceptual theories and seemingly contradictory experimental evidence is the application of generative models to empirical data (Heinzle, Aponte, & Stephan, 2016; Karayanidis et al., 2010; Monsell, 2003). In this direction, we recently developed the *Stochastic Early Reaction, Inhibition and late Action* (SERIA) model of the antisaccade task (Aponte, Schobi, Stephan, & Heinzle, 2017). In essence, SERIA combines the "horse-race" model of the countermanding saccade task (Camalier et al., 2007; Logan, Cowan, & Davis, 1984) to explain the inhibition of habitual, fast prosaccades, with a second race between two voluntary actions that generate pro- and antisaccades. In contrast to previous models (Noorani & Carpenter, 2013), SERIA takes into account that prosaccades are not only reactive or habitual saccades, but can also be voluntary actions.

The main goal of our study was to investigate whether switch costs can be attributed to the inhibition of habitual responses and/or to the generation of voluntary saccades. Moreover, we investigated whether modelling supports and explains the predictions of the task inertia and/or the task reconfiguration hypotheses. With these goals in mind, we applied SERIA to two versions of the antisaccade task (Aponte, Tschan, Heinzle, & Stephan, 2018). In Task 1, the peripheral stimulus served as task cue, indicating whether a pro- or an antisaccade should be performed. In Task 2, subjects were cued about the task demands in advance of the peripheral stimulus. Following previous reports, we expected positive antisaccade RT switch costs in Task 1, in which task and direction cue overlapped (similar to the short delay condition in Hunt & Klein, 2002; Barton, Greenzang, et al., 2006; Ethridge et al., 2009). In Task 2, we expected either negative or non-significant antisaccade RT switch costs, as the task cue was presented much in advance of the peripheral target (Barton, Greenzang, et al., 2006; DeSimone et al., 2014; Ethridge et al., 2009).

Our results indicate that switch costs in the antisaccade task are explained by two distinct inter-trial effects. Specifically, modelling supports task-inertia like effects on inhibitory control of habitual actions, as well as task-set reconfiguration costs in the execution of voluntary actions. We show here that by distinguishing between inhibitory control and voluntary action generation, it is possible to develop a unified account of the cost of switching in the antisaccade task.

2 | METHODS

In this study, we analysed switch costs in the data reported in Aponte et al. (2018). Therefore, we provide here only a short summary of the experimental procedure. The data are available for download at https://doi.org/10.3929/ethz-b-00029

6409. The experiment was approved by the ethics board of the Canton of Zurich, Switzerland (KEK-ZH-Nr.2014-0246), and was conducted according to the Declaration of Helsinki.

2.1 | Participants

Twenty-five healthy male subjects participated in the experiment. All subjects had normal or corrected to normal vision and provided written informed consent to participate in the study.

2.2 | Apparatus

The experiment was conducted in a dimly illuminated room. Subjects sat 60 cm in front of a cathodic ray tube (CRT) screen (41.4 × 30 cm; Philips 20B40; refresh rate 85 Hz). CRT monitors have a very precise time response function (Ghodrati, Morris, & Price, 2015) that warrants accurate stimulus time presentation. Sprites were synchronized with the screen refresh times (11 ms) allowing for exact control over the presentation time of the stimuli.

Eye position was recorded at a sampling rate of 1,000 Hz with a remote, infrared eye tracker (Eyelink 1000; SR Research, Ottawa, Canada). Head position was stabilized using a chin rest. The experiment was controlled by in-house software written in Python (2.7) using the PsychoPy package (1.82.02; Peirce, 2007, 2008).

2.3 | Experimental design

Subjects took part in two tasks consisting of three blocks of 200 mixed pro- and antisaccade trials. Either 20%, 50% or 80% of the trials were prosaccade trials. Before the main experiment, subjects participated in a training block of 50 prosaccade trials followed by 50 antisaccade trials of each task. In this phase, but not in the main experiment, subjects received feedback about their performance.

In Task 1 (Figure 1), two red circles (radius 0.25°) were presented throughout the experiment at ± 12°. Each trial started with a central fixation cross (0.6 × 0.6°). Subjects were required to fixate for at least 500 ms, after which a random interval of 500–1,000 ms started. Subsequently, the fixation cross disappeared, and a green bar (3.48 × 0.8°) in either horizontal or vertical orientation was presented behind one of the red circles for 500 ms. Subjects were instructed to saccade to the red circle cued by a horizontal green bar (prosaccade trials) and to saccade to the un-cued circle in case of a vertical bar (antisaccade trials). The next trial started after 1,000 ms. Pro- and antisaccade trials were randomly interleaved, but the same sequence was presented to all subjects. The location (left or right) of the peripheral cue was randomly permuted. The number of pro- and antisaccade trials in each direction was the same.

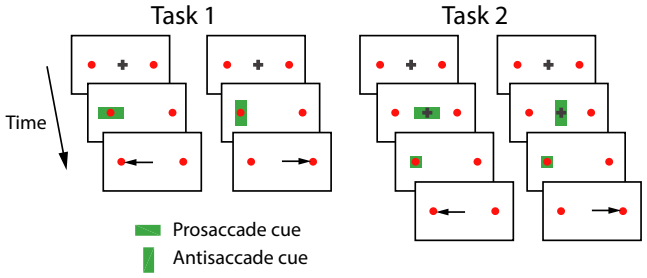

FIGURE 1 Experimental design. In both tasks, participants were instructed to first fixate a central cross. Task. 1: After a variable interval (500–1,000 ms), the trial type cue was presented behind one of the peripheral red circles for 500 ms. Depending on the orientation of the cue, a saccade had to be performed towards or away from the cued target. Task. 2: Before the peripheral stimulus was presented, subjects were cued for 700 ms about the task to be performed. After this period, the central fixation cross disappeared, and a neutral cue was presented behind one of the peripheral red circles for 500 ms. Depending on the orientation of the bar, a saccade towards or away from the cued target had to be performed. [Colour figure can be viewed at wileyonlinelibrary.com]

Task 2 differed in that subjects were cued about the trial type in advance of the peripheral stimulus. As in Task 1, subjects were required to initially fixate a grey cross for 500–1,000 ms. After this interval, either a horizontal or a vertical bar was displayed behind the fixation cross. The bar had the same dimensions and colour in both tasks. 700 ms later, the green bar and the fixation cross were removed, and a green square (1.74° × 1.74°) was presented behind one of the red circles for 500 ms. A horizontal bar indicated a prosaccade trial and a vertical bar indicated an antisaccade trial. After 1,000 ms, the next trial started.

2.4 | Data processing

Saccades were detected with the software provided by the eye tracker manufacturer, which uses a 22°/s velocity and a 3,800°/s^2 acceleration thresholds to define the onset of a saccade (Stampe, 1993). Only saccades larger than 2° were included in the analysis. Trials were rejected in case of eye blinks or if subjects failed to maintain fixation before the peripheral cue was presented. Saccades with a latency above 800 ms or below 50 ms were rejected as invalid. Antisaccades were also rejected if their RT was <90 ms. This follows the same analysis strategy as our previous studies (Aponte et al., 2017, 2018). Only trials that directly followed valid trials were included in the final analysis.

2.5 | Statistical analysis

As variables of interest, we investigated mean RT of correct saccades and mean ER. These were analysed with generalized linear mixed effects (GLME) models implemented in the programming language *R* (package *lme4*;

Bates, Mächler, Bolker, & Walker, 2015). Independent variables were prosaccade trial probability (PP) with levels 20%, 50% and 80%; trial type (TT); switch trial (SWITCH) with levels *switch* and *repeat*; and SUBJECT entered as a random effect. Significance was assessed with *F* tests with the Satterthwaite approximation to the degrees of freedom (Luke, 2017). For ER, the probit function was used as link function in the GLME. To test for significance, we used the Wald Chi-squared test implemented in the *car* package (Fox & Weisberg, 2011). When probabilities were investigated, we used a beta regression model implemented in the package *glmmADBM* (Fournier et al., 2012). Again, significance was assessed with Wald Chi-squared tests.

2.6 | The SERIA model

Briefly, SERIA (Aponte et al., 2017) models the race of four independent accumulators or units: an early (u_e), an inhibitory (u_i), a late prosaccade (u_p) and an antisaccade (u_a) unit. An action $A \in \{pro., anti.\}$ and its latency $T \in [0, \infty]$ are treated as random variables, whose distribution is a function of the hit times of each of the units, U_e, U_i, U_p and U_a respectively. Conceptually, SERIA can be decomposed into two different competitions (see Figure 2): First, the early unit, which models reactive, habitual responses, generates a prosaccade at time t if it hits threshold at time t (i.e., $U_e = t$) before all other units. An early response can be stopped by the inhibitory unit if the latter hits threshold at some earlier point. In that case, either a pro- or an antisaccade is generated, depending on the outcome of the second race between the late units. For example, a late prosaccade at time t is generated if the late prosaccade unit hits threshold at t (i.e., $U_p = t$) before the antisaccade unit (i.e., $U_a > t$).

Concretely, SERIA provides an explicit formula for the probability of an action A and its RT. First, a prosaccade at time t is generated when either the early unit u_e hits threshold at time t (i.e., $U_e > t$) before all other units. The probability of this event is

$$p(U_e = t)p(U_p > t)p(U_a > t)p(U_i > t). \quad (1)$$

Furthermore, a prosaccade at time t is triggered when the late prosaccade unit hits threshold at t before all other units

$$p(U_p = t)p(U_e > t)p(U_a > t)p(U_i > t) \quad (2)$$

or when an early response is stopped by the inhibitory unit (i.e., $U_e < t$ and $U_i < U_e$), and the late prosaccade unit hits threshold before the antisaccade unit

$$p(U_p = t)p(U_a > t)\int_0^t p(U_i = \tau)p(U_e > \tau)d\tau. \quad (3)$$

Similarly, an antisaccade at time t is generated when the antisaccade unit hits threshold at t ($U_a = t$), before all other units

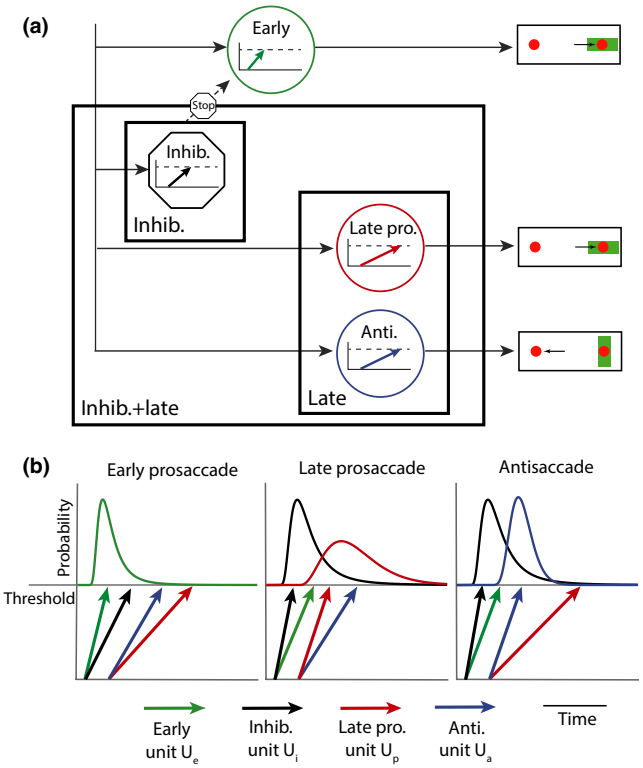

FIGURE 2 The SERIA model. (a) SERIA is a race model that incorporates four different units (displayed as circles): an early prosaccade unit (green), an inhibitory unit (black), a late prosaccade (red), and an antisaccade unit (blue). We hypothesized that the previous trial could affect the inhibitory unit (*inhib.*), the late units (*late*), or both (*inhib.+late*). These three hypotheses are represented by black frames around the units affected by the previous trial. (b) The RT distributions are a function of the hit time distributions of the four units. Early reactions, which are always prosaccades, occur when the early unit hits threshold before all other units. Late prosaccades occur mainly when the early unit is stopped by the inhibitory unit, and the late prosaccade unit hits threshold before the antisaccade unit. Similarly, antisaccades can only occur when the antisaccade unit hits threshold before the late prosaccade unit. Figure modified with permission from Aponte et al. (2018). [Colour figure can be viewed at wileyonlinelibrary.com]

$$p(U_a = t)p(U_e > t)p(U_p > t)p(U_i > t) \quad (4)$$

or when the antisaccade unit hits threshold before the late prosaccade unit after an early prosaccade has been stopped

$$p(U_a = t)p(U_p > t)\int_0^t p(U_i = \tau)p(U_e > \tau)d\tau. \quad (5)$$

To fit the model, we used a parametric form for the hit time distribution of each of the units: the hit times of the early (U_e) and inhibitory unit (U_i) were modelled with the inverse Gamma distribution, and the hit times of the late units (U_p and U_a) were modelled using the Gamma distribution. We selected these parametric forms following our previous studies (Aponte et al., 2017, 2018), where these distributions provided the best fist compared to other parametric forms.

Each unit can be fully characterized by two parameters controlling the mean and variance of the hit times. Accordingly, 8 parameters were required for the 4 units in a given condition.

2.7 | Model space

We aimed to answer two different questions through quantitative Bayesian model comparison (Kass & Raftery, 1995; Stephan, Penny, Daunizeau, Moran, & Friston, 2009) and qualitative predictive model checks (Gelman, Carlin, Stern, & Rubin, 2003): First, are models that include information about the previous trial superior in explaining experimental data compared to models that do not account for this factor? Second, can inter-trial effects be explained by changes in either the generation of voluntary saccades, inhibitory control or a combination of both?

To answer these questions, we fitted SERIA models that explain actions and RT not only as a function of the current trial type, but also as a function of the previous trial. For this, all trials were divided into four conditions, according to the cue displayed (pro- or antisaccade) and whether it was a switch or a repeat trial. Although a completely different set of parameters could operate in each condition, this seems biologically implausible. Rather, our goal was to identify which parameters could be fixed across conditions, while adequately fitting participants' behaviour. Based on our previous findings (Aponte et al., 2018), we fixed the subject-specific parameters of the early unit, the no-decision time, the probability of an early outlier and the delay of the late units (Aponte et al., 2017) across all conditions.

Initially, we evaluated a control model that could not account for any inter-trial effect. However, we allowed the inhibitory and the two late units to differ across pro- and antisaccade trials. This model had 12 free parameters for the late and inhibitory units ($2 \times 3 = 6$ per trial type). We refer to this model as the *no-switch* model.

Next, we considered the hypothesis that the late units (but not the inhibitory unit) could change on switch trials. Compared to the *no-switch* model, this model had $2 \times 2 \times 2 = 8$ additional parameters for switch and repeat trials. We refer to it as the *switch:late* model. By contrast, in the *switch:inhib*. model we allowed the inhibitory unit but not the late units to differ across switch and repeat trials. This required $2 \times 2 = 4$ extra parameters compared to the *no-switch* model. Finally, we combined the last two models into the *switch:inhib.+late* model, by permitting the late and the inhibitory units to differ across all conditions. Hence, this model required $(2 \times 2 \times 2) + (2 \times 2) = 12$ more parameters than the *no-switch* model.

2.8 | Model fitting

All models were fitted using the techniques described in our previous studies (Aponte et al., 2017, 2018). A set of parameters was estimated for each condition and subject. To regularize the estimates, we used a hierarchical model which imposed the population distribution as prior distribution (Aponte et al., 2018). Note that the population distribution was estimated simultaneously with the subject-specific parameters (Gelman et al., 2003). This means that while the model was fitted to the entire dataset at once, parameter estimates differed from subject to subject.

Samples from the posterior distribution were drawn using the Metropolis–Hasting algorithm. The evidence or marginal likelihood of a model was computed using thermodynamic integration (Aponte et al., 2016; Gelman & Meng, 1998) with 16 parallel chains ordered according to the temperature schedule in Calderhead and Girolami (2009). Besides the marginal likelihood, we report the Watanabe–Akaike information criterion (Gelman, Hwang, & Vehtari, 2014; Watanabe, 2010), which is an approximation of the out-of-sample predictive accuracy of a model. The sampling algorithm was run for 130,000 iterations. Only the last 30,000 samples were used to compute summary statistics. The implementation of the models and inference is available in the open source TAPAS toolbox (http://translationalneuromodeling.org/tapas/).

We were interested in several model-based statistics derived from the fits. First, we evaluated the probability of an inhibition failure, defined as the probability that the early unit hits threshold before all other units:

$$p(inhib.fail.) = \int_0^\infty p(U_e = t)p(U_i > t)p(U_p > t)p(U_a > t)dt. \quad (6)$$

Inhibition failures are fast, reflexive prosaccades, which are correct on prosaccade trials and errors on antisaccade trials.

We also report the conditional probability of a late prosaccade, defined as the probability that the late prosaccade unit hits threshold before the antisaccade unit:

$$p(late\ pro.) = \int_0^\infty p(U_p = t)p(U_a > t)dt. \quad (7)$$

The probability of a late prosaccade as defined above is independent of the race between the early and inhibitory units. It reflects the competition between voluntary actions independently of the influence of the inhibitory unit. Note that the conditional probability of an antisaccade is defined as

$$p(anti.) = 1 - p(late\ pro.). \quad (8)$$

We were also interested in the expected hit times of the late units, defined as

$$E[late\ pro.\ hit\ time] = \frac{1}{p(late\ pro.)} \int_0^\infty tp(U_p = t)p(U_a > t)dt \quad (9)$$

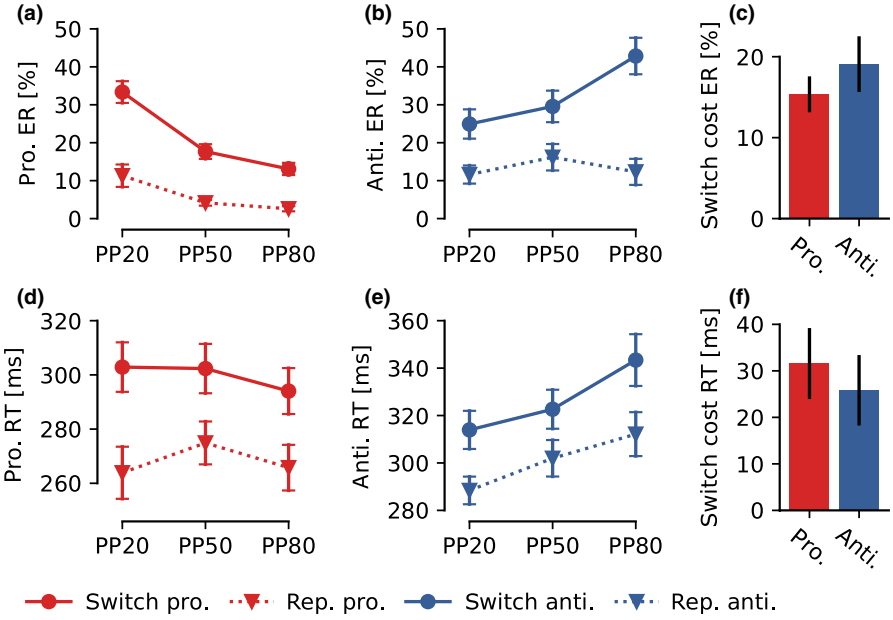

FIGURE 3 Error rate (ER) and reaction time (RT) in Task 1. (a) Mean ER on prosaccade trials. (b) Mean ER on antisaccade trials. (c) ER switch cost. (d) Mean RT on prosaccade trials. (e) Mean RT on antisaccade trials. (f) RT switch cost. Error bars display the sem. PP: prosaccade trial probability. [Colour figure can be viewed at wileyonlinelibrary.com]

and analogously so for antisaccades. This quantity is the expected hit time of the late prosaccade unit, conditioned on the antisaccade unit arriving at a later point. We report this statistic as it conveys an interpretable quantity that can be readily compared to experimental data.

3 | RESULTS

From the 25 participants recruited, two subjects were not included in the final analysis. One subject was excluded because of incomplete data and the second because in two blocks more than 50% of the trials were either invalid or directly followed an invalid trial.

In the following, we report Task 1 and Task 2 separately. First, classical statistical analyses of mean RT and ER are presented. These are followed by model-based analyses, in which we compare the *no-switch* and *switch* models using quantitative Bayesian model comparison. We then restrict our attention to *switch* models and explore them in detail, using posterior predictive fits (Gelman et al., 2003) to test when and why individual models fail to predict participants' behaviour.

3.1 | Task 1

In Task 1, roughly 4% of the trials were discarded.

3.1.1 | Error rate and reaction times

Mean RT, ER and switch costs are displayed in Figure 3. On average, participants made significantly more errors on anti- (22% ± 21%) than on prosaccade trials (14% ± 14%; $\chi^2 = (1,276) = 146.2$, $p < 10^{-5}$), and on switch (27% ± 19%) than on repeat trials (10% ± 13%; $\chi^2 = (1,276) = 406.6$, $p < 10^{-5}$). The antisaccade switch cost (19%) was significantly higher ($\chi^2 = (1,276) = 8.4$, $p = 0.003$) than the prosaccade switch cost (15%).

Regarding RT, antisaccades (313 ± 44 ms) were significantly slower than prosaccades (284 ± 45 ms; $F_{1,242} = 57.8$, $p < 10^{-5}$) and switch trials (313 ± 46 ms) were slower than repeat trials (285 ± 43 ms; $F_{1,242} = 53.6$, $p < 10^{-5}$). The interaction between TT and SWITCH was not significant ($F_{1,242} = 0.5$, $p = 0.463$). Therefore, the antisaccade switch cost (26 ms) did not significantly differ from the prosaccade switch cost (32 ms).

3.1.2 | SERIA – model comparison

All models were evaluated according to their log model evidence (LME), which corresponds to the accuracy or expected log-likelihood of a model adjusted by its complexity (Stephan et al., 2009). Table 1 displays the accuracy, LME and WAIC of all models in log units. The model with the highest evidence was the *switch:inhib.+late* model (ΔLME > 44 log units compared to all other models). It also obtained the highest accuracy and WAIC. Note that this model was heavily penalized (accuracy-evidence = 940) compared to the simpler models *no-switch* (accuracy-evidence = 782), *switch:late* (accuracy-evidence = 922) and *switch:inhib.* (accuracy-evidence = 834).

The predictive fits of all models are displayed in Figure 4. These represent the expected predictive distribution generated from posterior samples. Individual subjects' fits were averaged to generate a single prediction.

Visual inspection suggests that all models explained the most salient features of the data, including the shape of the

RT distributions. However, while the *no-switch* model failed to capture the distribution of switch prosaccades (Figure 4c), the *switch:inhib*. model failed to capture the distribution of late responses, and particularly so on prosaccade trials (Figure 4a,c). The *switch:late* model made a better job regarding late saccades, but it was less accurate in fitting early errors on antisaccade trials (Figure 4d). Finally, the *switch:late+inhib*. model was able to accommodate most features of the data.

Figure 5 displays the switch costs predicted by all *switch* models. Clearly, only model *switch:inhib.+late* was able to capture switch costs on both pro- and antisaccade trials, whereas models *switch*:*late* and *switch:inhib*. only explained ER and RT in one condition.

3.1.3 | SERIA – parameter estimates

According to SERIA, there are two types of errors on antisaccade trials: inhibition failures and late prosaccades. To disentangle how these two types of errors contributed to antisaccade switch costs, we first turned our attention to the probability of an inhibition failure (Equation 6). This is defined as the probability that the early unit hits threshold before all other units. The effect of switching on pro- (χ^2 (1,138) = 107.9, $p < 10^{-3}$) and antisaccades trials (χ^2 (1,138) = 229.2, $p < 10^{-5}$) was significant. When considered together, we found a significant interaction between TT and SWITCH (χ^2 (1,276) = 302.1, $p < 10^{-5}$). Concretely, prosaccade trials induced more inhibition failures on the next trial than antisaccades (pro. switch cost = −18%; anti. switch cost = 19%).

This suggests that the probability of an inhibition failure after a prosaccade trial does not depend on the upcoming trial type. To explore this hypothesis, we fitted a model in which the parameters of the inhibitory unit were a function of the previous trial. Thus, rather than switch costs, this model accounted for task-inertia effects. The evidence for this model was higher than the evidence for the *switch:inhib.+late* model (ΔLME = 12.1). In summary, there were fewer early saccades immediately after antisaccade trials than after prosaccade trials.

Next, we submitted the probability of late errors (Equations 7 and 8) on pro- (19 ± 15%) and antisaccade (4 ± 5%) trials to a single GLME. This revealed a significant interaction between SWITCH and TT (χ^2 (1,276) = 63.0, $p < 10^{-5}$). The late ER switching cost on prosaccade trials was 18%, whereas on antisaccade trials, it was less than 1%. When late ER on antisaccade trials was analysed separately, the factor SWITCH was not significant (χ^2 (1,138) = 0.1, $p < 0.81$).

Finally, we investigated the hit times of the late units (Equation 9). Switch late saccades (335 ± 42 ms) were significantly ($F_{1,248} = 81.9$, $p < 10^{-5}$) slower than repeat saccades (312 ± 36 ms). The late prosaccade RT switch cost (18 ms) was lower than the antisaccade RT cost (29 ms) which resulted in a significant interaction between TT and SWITCH ($F_{1,248} = 4.8$, $p < 0.028$). In conclusion, voluntary saccades were slower on switch trials than on repeat trials, revealing a cost of switching in voluntary behaviour.

TABLE 1 Model comparison Task 1

	Accuracy	LME	WAIC
no-switch	−15,673	−16,455	−16,026
switch:late	−15,231	−16,175	−15,665
switch:inhib.	−15,400	−16,235	−15,795
switch:inhib.+late	**−15,212**	**−16,153**	**−15,627**

Note: Expected log-likelihood or accuracy, log model evidence (LME) and Watanabe–Akaike Information Criterion (WAIC) are listed for the four models tested. The highest accuracy, LME and WAIC in the *switch* family are highlighted in bold.

3.2 | Task 2

In Task 2, around 9% of all trials were discarded. At most, 35% of all trials in a single block were excluded.

3.2.1 | Error rate and reaction time

In this condition (Figure 6), subjects made significantly fewer errors on pro- (2% ± 4%; Figure 6a) than on antisaccade trials (13% ± 13%; $\chi^2 = (1,276) = 297.4$, $p < 10^{-5}$; Figure 6b), and on repeat (5% ± 10%) than on switch trials (10% ± 12%; $\chi^2 = (1,276) = 77.4$, $p < 10^{-5}$). There was a significant interaction between SWITCH and TT ($\chi^2 = (1,276) = 6.3$, $p < 0.011$; Figure 6c) driven by larger switch costs on antisaccade trials (8%) than on prosaccade trials (3%).

Prosaccades (Figure 6d; 155 ± 26 ms) were faster than antisaccades (Figure 6e; 194 ± 30 ms; $F_{1,242} = 385.8$, $p < 10^{-5}$), but neither the effect of SWITCH ($F_{1,242} = 1.0$, $p = 0.314$) nor the interaction between SWITCH and TT ($F_{1,242} = 3.0$, $p = 0.079$) were significant (Figure 6f). Nevertheless, we submitted pro- and antisaccades to two separate GLME. As shown in Figure 6f, prosaccades were significantly slower on switch than on repeat trials (ΔRT = 5 ms; $F_{1,110} = 6.4$, $p = 0.012$), but there was no significant difference on antisaccade trials (ΔRT = −1 ms; $F_{1,110} = 0.2$, $p = 0.576$), although switch antisaccades were slightly faster than repeat antisaccades.

3.2.2 | SERIA – model comparison

In contrast to Task 1, the *no-switch* model obtained the highest LME (Table 2). The second best model was the *switch:inhib*. model, in which the inhibitory unit was allowed to change across all four possible conditions, but the late units were the same on switch and repeat trials. The *switch:inhib*. model also obtained the highest WAIC among the switch models.

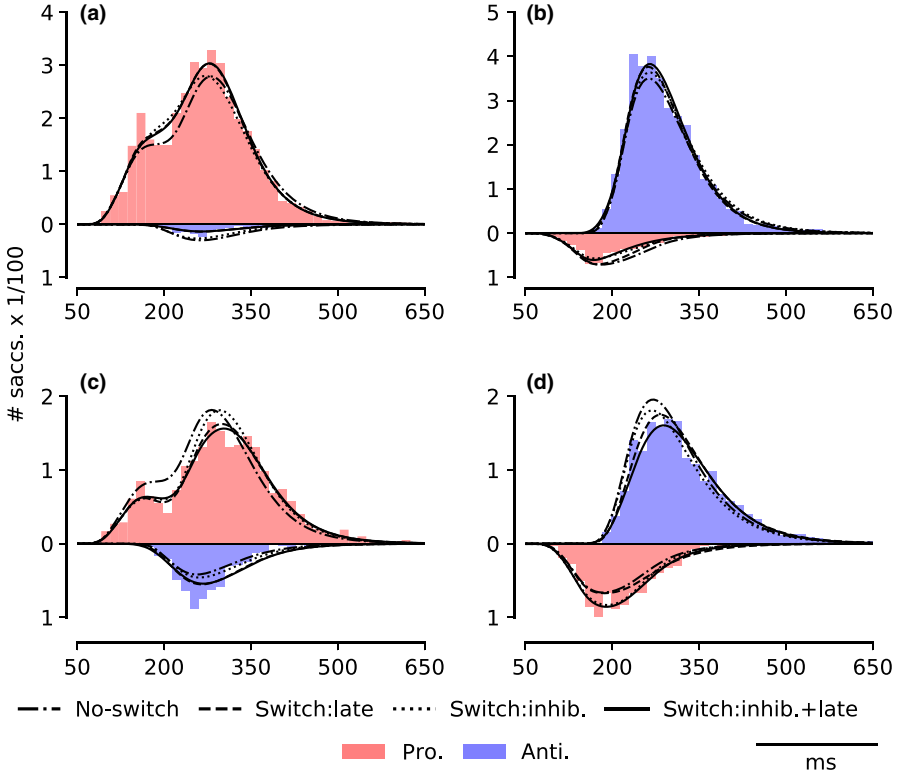

FIGURE 4 Histogram of the empirical reaction time (RT) distribution and model fits. Prosaccades are displayed in red and antisaccades in blue. Errors are displayed in the negative half plane. The posterior predictive distributions of models *no-switch*, *switch:late*, *switch:inhib*. and *switch:late+inhib*. are plotted in different line styles. (a) Prosaccade repeat trials. (b) Antisaccade repeat trials. (c) Prosaccade switch trials. (d) Antisaccade switch trials. [Colour figure can be viewed at wileyonlinelibrary.com]

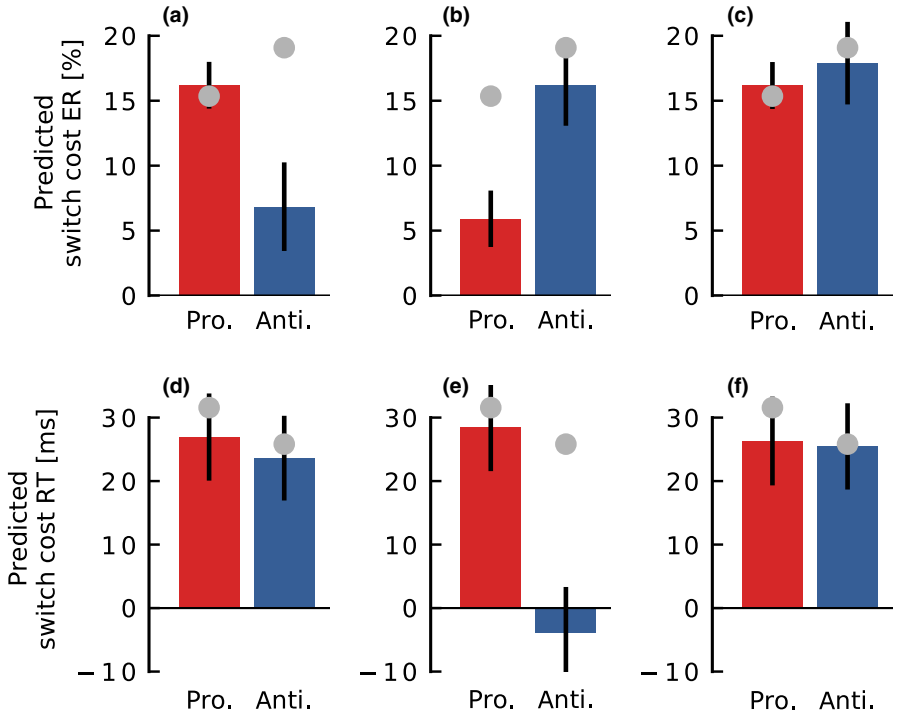

FIGURE 5 Predicted error rate (ER) and reaction time (RT) switch costs. (a–c) ER switch costs predicted by the *switch* models. Empirical switch costs (Fig. 3c and 3f) are displayed as grey circles. (a) *switch:late*. (b) *switch:inhib*. (c) *switch:inhib.+late*. While the *switch:late* model correctly predicted ER switch costs on prosaccade trials, antisaccade ER costs were clearly underestimated. By contrast, the *switch:inhib*. model captured ER costs on anti- but not on prosaccade trials. The *switch:inhib.+late* made a good job in both conditions. (d–f) RT switch costs predicted by the *switch* models. (d) *switch:late*. (e) *switch:inhib*. (f) *switch:inhib.+late*. The *switch:late* and *switch:inhib.+late* models captured RT costs in both pro- and antisaccade trials. Error bars depict the sem. of the model predictions. [Colour figure can be viewed at wileyonlinelibrary.com]

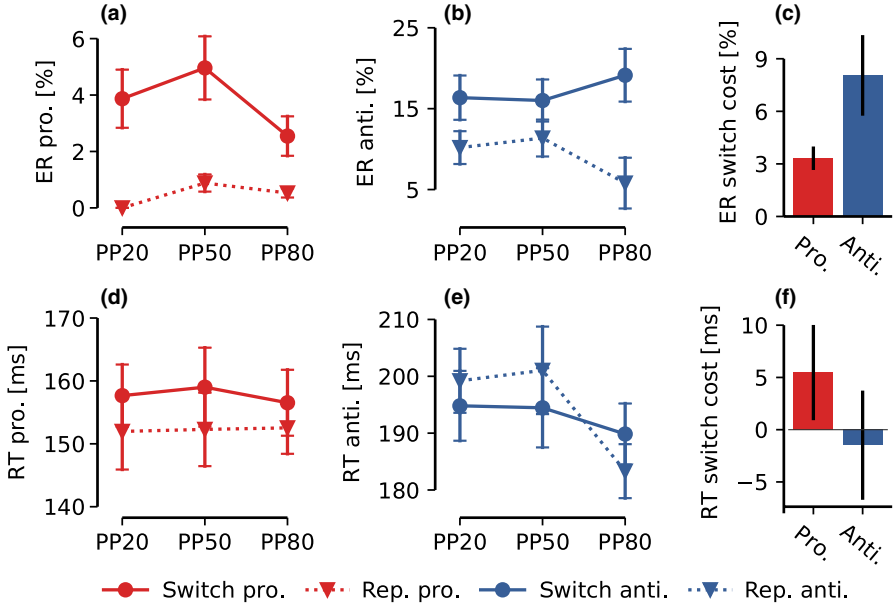

FIGURE 6 Error rate (ER) and reaction time (RT) in Task 2. (a) Mean ER on prosaccade trials. (b) Mean ER on antisaccade trials. (c) ER switch cost. (d) Mean RT on prosaccade trials. (e) Mean RT on antisaccade trials. (e) RT switch cost. Error bars display the sem. PP: prosaccade trial probability. [Colour figure can be viewed at wileyonlinelibrary.com]

All four models fitted RT and ER well (Figure 7), with no obvious difference between them. The *no-switch* model obtained the highest LME despite having the lowest accuracy because it was less penalized for its number of parameters (complexity) compared to the *switch* models. For example, the second best model (*switch:inhib.*) had 4 parameters more per block than the *no-switch* model. However, the *no-switch* model could not capture any of the effects of interest demonstrated by the classical analysis. This suggests that the penalization term in the LME is too strong given the subtlety of the switch costs in Task 2.

To make a more qualitative examination of the models, we compared their predictions with empirical observations. We start discussing switch costs in Task 2 based on the best model from the switch family (*switch:inhib.*). We come back to models *switch:late* and *switch:inhib.+late* below.

Qualitatively, the *switch:inhib.* model (Figure 8b) could reproduce the main behavioural findings: switch trials had higher ER (10 ± 11%) than repeat trials (6 ± 11%; $\chi^2(1,248) = 58.3$, $p < 10^{-5}$). Although the predicted switch cost was higher on anti- (5%) than on prosaccade trials (1%), the interaction between SWITCH and TT was not significant ($\chi^2(1,248) = 0.1$, $p < 0.75$) contrary to our behavioural analysis. Moreover, the model underestimated the ER switch cost on pro- and antisaccade trials (Figure 8b), as discussed in the next section. Regarding RT, the model captured the positive switch cost on prosaccade trials (5 ms, $F_{1,112} = 9.8$, $p = 0.002$), as well as the small negative cost on antisaccades trials ($F_{1,112} = 0.0$, $p = 0.834$).

3.2.3 | SERIA – model parameters

To understand how the *switch:inhib.* model was able to capture switch costs in Task 2 in the absence of changes in the late units, we plotted the probability of inhibition failures on switch and repeat trials (Figure 9a,b). As in Task 1, saccades that followed prosaccade trials were more likely to be inhibition failures, regardless of trial type (Figure 9c; interaction TT * SWITCH; $\chi^2(1,248) = 47.7$, $p < 10^{-5}$). Thus, switch antisaccade trials were more likely to be errors.

In general, prosaccade trials led to more inhibition failures on the next trial (regardless of trial type). Because there were more late prosaccades on switch than on repeat trials, switch prosaccades were slower and had higher ER than repeat prosaccades.

TABLE 2 Model comparison Task 2

	Accuracy	LME	WAIC
no-switch	−5,258	−6,008	−5,586
switch:late	**−5,194**	−6,092	−5,616
switch:inhib.	−5,234	**−6,047**	**−5,610**
switch:inhib.+late	−5,257	−6,253	−5,700

Note: Expected log-likelihood or accuracy, log model evidence (LME) and the Watanabe–Akaike Information Criterion (WAIC) are listed for the four models tested. The highest accuracy, LME and WAIC in the *switch* family are highlighted in bold.

3.3 | Prosaccade and antisaccade error rate switch cost

As illustrated above (Figure 8b), the *switch:inhib* model underestimated the ER switch cost. In the case of prosaccades, this model could not fully capture the eightfold increase in ER on switch trials (3.79%) compared to repeat trials (0.47%). Is it possible to explain this difference by extending the *switch:inhib.* model?

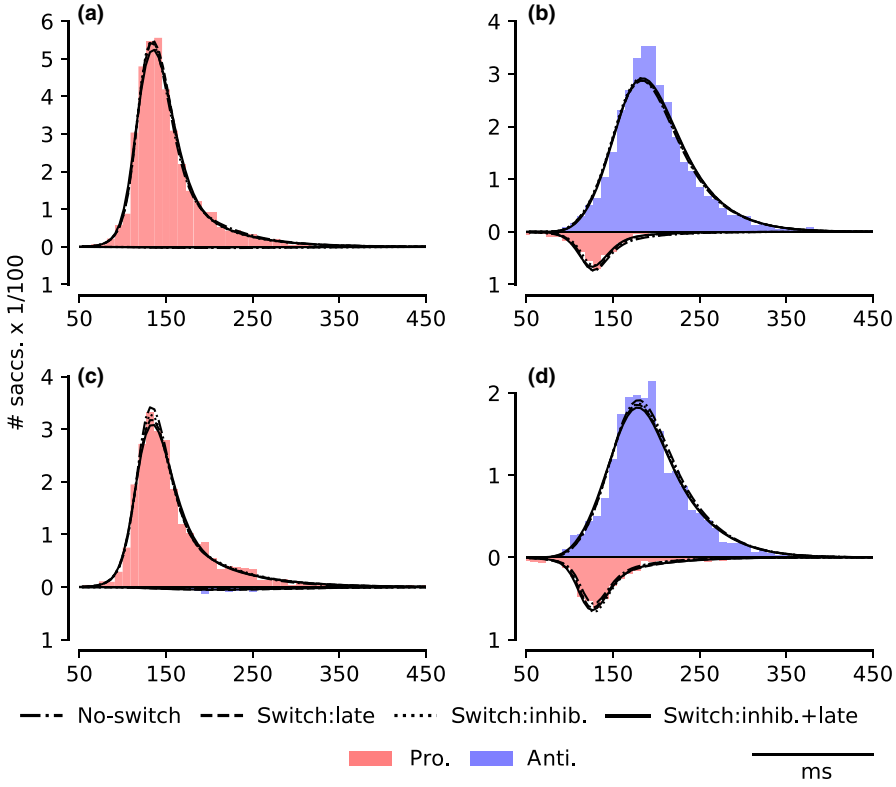

FIGURE 7 RT histograms and predictive model fits in Task 2. Similar to Fig. 4. (a) Prosaccade repeat trials. (b) Antisaccade repeat trials. (c) Prosaccade switch trials. (d) Antisaccade switch trials. With the exception of prosaccade switch trials (c), all models generated similar fits. [Colour figure can be viewed at wileyonlinelibrary.com]

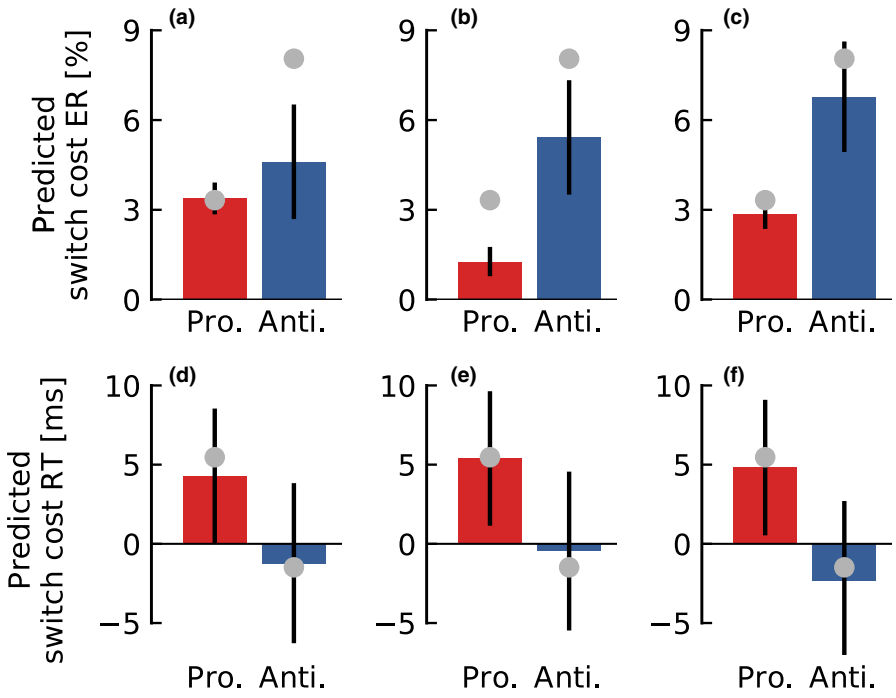

FIGURE 8 Model predictions. Top. Predicted ER switch cost. (a) *switch:late*. (b) *switch:inhib*. (c) *switch:inhib.+late*. Bottom. Predicted RT switch cost. (d) *switch:late*. (e) *switch:inhib*. (f) *switch:inhib.+late*. [Colour figure can be viewed at wileyonlinelibrary.com]

According to SERIA, a prosaccade error can almost only[3] occur when an early response is inhibited and the antisaccade unit hits threshold before the late prosaccade unit. Thereby, the prosaccade ER is *approximately* equal to

$$p_{pro.\ error} \approx (1 - p_{inhib.\ fail}) * p_{late\ error}. \quad (10)$$

In the *switch:inhib* model, the late units are assumed to not change across switch and repeat trials. Hence, an eightfold increase in the ER is only possible if there is an eightfold change in $1 - p_{inhib.\ fail}$ (Equation 10). Such a large change is impossible

[3]It is possible, although highly unlikely, that the antisaccade unit hits threshold before all three other units.

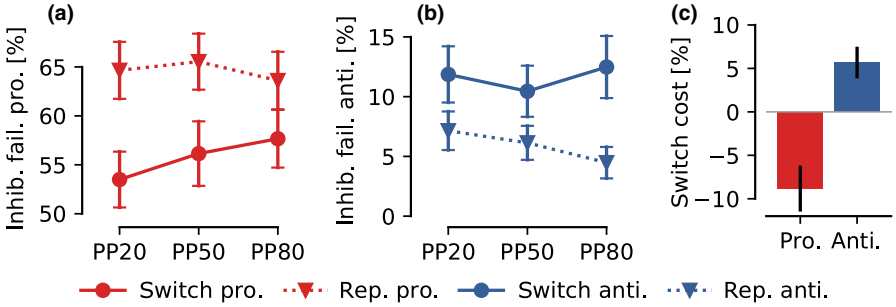

FIGURE 9 Inhibition failures in Task 2 according to the *switch:inhib*. model. (a) Predicted probability of an inhibition failure on prosaccade trials. (b) Inhibition failures on antisaccade trials. (c) Inhibition switch cost on pro- (-9%) and antisaccade (6%) trials. Error bars display the sem. [Colour figure can be viewed at wileyonlinelibrary.com]

because the predicted probability of an inhibition failures $p_{inhib.fail}$ is on average 0.6 (see Figure 9a). Thus, higher ER on prosaccade switch trials can only be explained by changes in the late units.

To account for this cost, we considered a model (*switch:inhib.+anti*.) in which we allowed the parameters of the antisaccade unit to differ between switch and repeat prosaccade trials. These parameters control the probability and RT of errors on prosaccade trials but have no influence on antisaccade trials. As displayed in Figure 10c,d, the ER predicted by the *switch:inhib.+anti*. model was 3.67% (switch) and 0.67% (repeat). When we considered again the interaction SWITCH*TT using the *switch:inhib.+anti*. model, this was significant (χ^2 (1,276) = 20.5, $p < 10^{-5}$). Nevertheless, the *switch:inhib.+anti*. model had a lower LME than the *switch:inhib*. model (ΔLME=67.0).

Regarding antisaccade trials, the ER switch cost was underestimated by the *switch:inhib*. model (empirical 8.1%, predicted 5.3%). However, as shown in Figure 11, this was mainly due to the PP80 condition, in which the empirical ER on repeat trials was lower than predicted by the model. Note that this condition is by design much less frequent than the others, and thereby the mean ER displays high uncertainty.

Taken together, our analyses demonstrate that, as in Task 1, switch voluntary prosaccades are more likely to be errors than repeat voluntary prosaccades. By contrast, there were no more late errors on switch than on repeat antisaccade trials. Rather, the antisaccade ER switch cost was completely explained by the increase on inhibition failures that follow prosaccade trials.

4 | DISCUSSION

Here, we investigated switch costs in the mixed antisaccade task with the help of a computational model. This allowed us to quantify to what extent task switching affects the inhibition of habitual responses (early prosaccades) and voluntary behaviour (late pro- and antisaccades). Modelling revealed two main effects: First, in Task 1 but not in Task 2, voluntary actions were slower on switch trials. Second, in both tasks, early prosaccades were less likely to be inhibited after prosaccade trials than after antisaccade trials. Can SERIA accommodate the predictions of the task-set reconfiguration and the task-inertia hypotheses? Does SERIA provide an alternative or more fine-grained explanation for these predictions? In the following, we discuss the answers to these questions.

4.1 | Switch costs in the antisaccade task

Previous findings in the mixed antisaccade task can be divided into three groups. Early studies (e.g., Barton et al.,

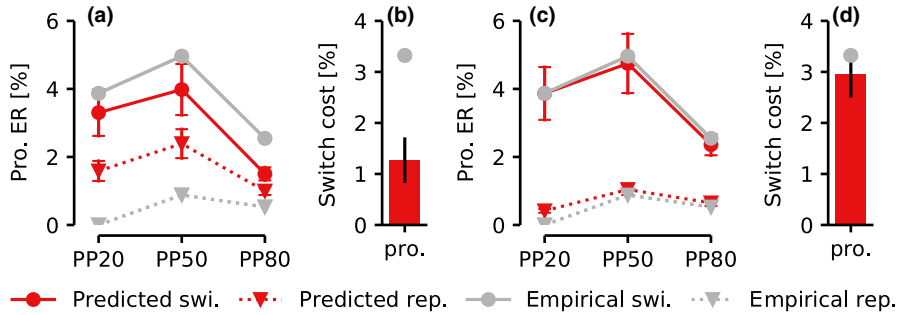

FIGURE 10 Predicted and empirical ER and switch cost on prosaccade trials. (a) *switch:inhib*. predictions. The *switch:inhib*. model accounts for the switch cost only through changes in inhibitory control. (b) *switch:inhib*. ER switch cost. Although this model does capture a fraction of the switch cost, it is limited by the proportion of inhibition failures on repeat trials. For visualization, the empirical switch cost is displayed as a grey circle. (c) *switch:inhib.+anti*. predictions. In the *switch:inhib.+anti*. model, the antisaccade unit can be different on prosaccade switch and repeat trials. In this case, the predicted ER on repeat trials is closer to the empirical ER. (d) *switch:inhib.+anti*. ER switch cost. Similar to panel B. Error bars display the sem. of the model predictions. [Colour figure can be viewed at wileyonlinelibrary.com]

2002; Cherkasova et al., 2002; Fecteau et al., 2004; Manoach et al., 2002) reported positive prosaccade RT switch costs, negative antisaccade RT costs, as well as higher ER on switch trials regardless of trial type. More recently, Heath and Weiler (e.g., Weiler & Heath, 2012a; Weiler et al., 2015) reported positive RT switch costs on prosaccade trials, and no RT switch costs on antisaccade trials. Again, all switch trials were characterized by higher ER. Finally, some studies have reported positive RT switch costs on pro- and antisaccade trials when the task cue is presented close in time to the peripheral stimulus (Barton, Greenzang, et al., 2006; Ethridge et al., 2009; Hunt & Klein, 2002).

Our empirical findings are well in line with these studies. Regarding Task 1, positive switch costs on pro- and antisaccade trials have been demonstrated in a similar design by Barton, Greenzang, et al. (2006); see also Hunt and Klein (2002). In Task 2, we found non-significant negative antisaccade RT costs and significant positive prosaccade RT costs. This is congruent with the unidirectional costs reported by Weiler and Heath (2012a).

Based on SERIA, we proposed three models or hypotheses to explain our results: (a) the *switch:inhib*. model in which only the parameters of the inhibitory unit can change across switch and repeat trials; (b) the *switch:late* model in which the late units but not the inhibitory unit can be different across conditions; and (c), the *switch:inhib.+late* model which combines both hypotheses.

Quantitative Bayesian model selection and qualitative posterior predictive checks (Gelman & Shalizi, 2013; Gelman et al., 2003) indicated that in Task 1 the *switch:inhib.+late* model accounted best for the data. In Task 2, contrary to the behavioural analysis, the model with the highest evidence did not allow for any switch cost. A possible reason for this discrepancy is that the effects of interest were subtle in Task 2 compared to Task 1. In other words, the improvement in accuracy of the *switch* models was not large enough to offset their additional complexity. Nevertheless, we proceeded to consider *switch* models in more detail, to understand how SERIA could explain switch costs in Task 2.

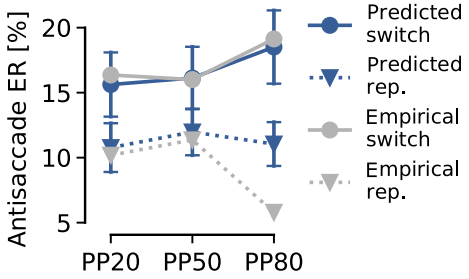

FIGURE 11 Predicted and empirical ER on antisaccade trials. Predictions are from the *switch:inhib*. model, in which the late units do not change between switch and repeat trials. The model overestimated the empirical ER only in repeat trials in the PP80 condition. Error bars display the sem. of the model predictions. [Colour figure can be viewed at wileyonlinelibrary.com]

In the *switch* family, the *switch:inhib*. model obtained the highest LME, as well as the highest out-of-sample predictive accuracy (WAIC). However, it is worth noting that all models captured the mean RT distributions well, but differed in subtle aspects related to task switching. This is to be expected because all models shared similar parameters. Qualitatively, the *switch:inhib*. model could fit RT switch costs on pro- and antisaccade trials, and, after an extension, it could fit prosaccade ER switch costs.

Altogether, SERIA demonstrated that in Task 1 there was a cost associated with switching between voluntary pro- and antisaccades. Concretely, voluntary saccades were slower on switch trials than on repeat trials. In Task 2, according to SERIA, voluntary saccades had the same latency on switch and repeat trials. This is compatible with the task-set reconfiguration hypothesis (Meiran, 1996; Rogers & Monsell, 1995), which states that switching between task-sets is time-consuming (Task 1), but can be done in advance of the response cue (Task 2).

Besides the switch cost associated with voluntary actions, we found that there was an inter-trial effect on inhibitory control in both tasks. Specifically, inhibition failures were more common after prosaccade trials than after antisaccade trials, regardless of the current trial type. This is predicted by the task-inertia hypothesis, according to which "switch costs should change as a function of the task that participants are switching from, not as a function of the task they are switching to" (Wylie & Allport, 2000). A second prediction of this hypothesis is confirmed by our modelling: inter-trial effects persist regardless of the delay between the task cue and the imperative stimulus (i.e., the peripheral target; Wylie & Allport, 2000).

The answer to our first question (Can SERIA accommodate the predictions of the task-inertia and task-set reconfiguration hypotheses?) is therefore positive. On one hand, asymmetric switch costs are explained by persistent or residual inhibition of habitual actions. On the other hand, higher RT on switch trials in Task 1 are caused by delays in the generation of voluntary actions.

The answer to our second question (How does SERIA explain these predictions?) is more nuanced. Although our modelling is compatible with the predictions of the task-inertia hypothesis, SERIA postulates a different mechanism for this inter-trial effect. Rather than passive interference between task-set rules (Weiler et al., 2015), our results indicate that the strong inhibition associated with an antisaccade trial reduces the probability of an inhibition failure on the next trial.

The mechanism described by SERIA differs also from the model proposed by Barton, Greenzang, et al. (2006). According to it, switch costs are partly caused by the generalized suppression of the response-system that "affects both the upcoming pro- and antisaccades." However, it is not obvious how generalized suppression can increase prosaccade

ER on switch trials compared to repeat trials. By contrast, in SERIA, switch prosaccades are more likely to be slow voluntary saccades (which explains the prosaccade RT switch cost in Task 2). More importantly, fewer inhibition failures also allow for more (late) errors on prosaccade switch trials (which explains the prosaccade ER switch cost). We come back to this point in the next section.

From a more biological perspective, increased inhibitory control after an antisaccade trial is associated with decreased activation of the frontal eye fields (FEF) in humans (Manoach et al., 2007). This difference might not extend to subcortical regions in the primate brain. In particular, Chan et al. (2017) found no differences in the superior colliculus (SC) between switch and repeat antisaccade trials. The physiological correlates of switching between voluntary saccades are less clear. The middle occipital gyrus and the inferior parietal lobule are more active on switch trials compared to repeat trials (Pierce & McDowell, 2017). In addition, Lee et al. (2011) found increased preparatory activity in the FEF and in the dorsal anterior cingulate cortex on switch trials compared to repeat trials. It has been hypothesized (Everling & Johnston, 2013) that the main role of the prefrontal cortex (PFC) in the antisaccade task is to encode task-set rules, rather than to inhibit prosaccades. This suggests that the cost of switching between voluntary actions could be related to neural activity in PFC. Interestingly, differential activation on repeat and switch trials in other tasks has been reported in the PFC as well as in the posterior parietal cortex (reviewed in Karayanidis et al., 2010).

So far, we have not discussed the negative or paradoxical antisaccade RT switch cost (Cherkasova et al., 2002). This only occurs when the task cue is presented in advance of the peripheral cue, as in Task 2. To explain this effect, in Supp. Material 1, we demonstrate that the *switch:inhib*. model can generate negative switch costs. It does so in the absence of changes in the late units across repeat and switch trials. This is possible because of the non-linear interactions between the antisaccade, the early and the inhibitory units.

In summary, SERIA can explain the plurality of behavioural findings in the antisaccade task: positive, unidirectional and paradoxical switch costs. Next, we discuss in more detail how the *switch:inhib*. model allows for asymmetric costs.

4.2 | Asymmetric costs

A key observation in the task switching literature is that switching from a non-habitual to a habitual response engenders higher costs than switching from a habitual to a non-habitual response (Allport et al., 1994; Wylie & Allport, 2000). SERIA provides a simple mathematical explanation for this phenomenon. The expected RT of dominant or habitual responses can be *approximated* as the weighted sum of the early and late components

$$E[habitual\ RT] = p_{early}E[early\ RT] + (1 - p_{early})E[late\ RT]. \quad (11)$$

The expected RT of non-habitual responses is given by

$$E[RT\ non\ habitual] = E[late\ RT]. \quad (12)$$

Accordingly, in a transition from a non-habitual to a habitual response, the probability of a late response increases, elevating the overall mean RT, even in the absence of voluntary action switch costs. In the case of a transition from a habitual to a non-habitual response, the RT of non-habitual responses should be equal to the RT of repeat trials. This is how the *switch:inhib*. model explains the positive RT switch cost on prosaccade trials as well as the absence of RT switch costs on antisaccade trials in Task 2. Note that this approximation is invalid under certain circumstances, as demonstrated in Supp. Material 1. There, we illustrate how the *switch:inhib*. model can generate negative antisaccade RT switch costs.

To our knowledge, no other computational model has been used to investigate the inhibition of habitual responses as a component of task switching, nor has this mechanism been suggested as an explanation for asymmetric costs. For instance, Schmitz and Voss (2014) extended the drift-diffusion model (Ratcliff, Smith, Brown, & McKoon, 2016) to explain task-switching costs. However, the goal of that study was to quantify the contribution of task-set reconfiguration and cue switching (Logan & Bundesen, 2003), without postulating any form of response inhibition. The same is true for the model of cue switch costs proposed by Logan and Bundesen (2003). A third example is the model in Altmann and Gray (2008), which relies on proactive interference to explain switch costs without postulating any form of inhibitory control.

It is worth mentioning that while the concept of "inhibition" plays a significant role in the task switching literature (reviewed in Koch, Gade, Schuch, & Philipp, 2010), this is usually defined as the "proactive interference resulting from having performed a competing task" (Koch et al., 2010). Here, we have used "inhibition of habitual responses" in the narrow sense of "motor inhibition" (Logan et al., 1984; see Schall et al., 2017). Specifically, inhibition, as used here, only affects habitual reactions.

Besides task-switching costs, other inter-trial effects have been observed in antisaccade paradigms. Alternating between left and right saccades (response switching) is known to affect performance in the mixed antisaccade task (Reuter, Philipp, Koch, & Kathmann, 2006). Trial number in a block affects antisaccade RT (Pierce et al., 2015), and training across sessions reduces anti- and prosaccade latency (Jamadar, Johnson, Clough, Egan, & Fielding, 2015). Interestingly, there is no evidence of task-inertia effects that extend beyond one trial (Hunt & Klein, 2002; Weiler & Heath, 2014b). While it would be interesting to study these

effects with SERIA, it has to be pondered that the number of model parameters increases exponentially with the number of factors considered. For example, modelling left–right switch costs in addition to task switching would require doubling the number of parameters. Studying these effects is an interesting future direction that would necessitate larger datasets.

A different approach to model "slow dynamics," such as learning, would require that the model parameters change from trial to trial. In a future extension of SERIA, some of its parameters could be treated as the states of a learning model, such as the hierarchical Gaussian filter (Mathys et al., 2014) or the Rescorla–Wagner model (Rescorla & Wagner, 1972). This could be used to study how trial type probability is learnt in a block.

5 | CONCLUSION

Using a model of antisaccade behaviour, we illustrated how conceptual theories of switch costs can profit from a rigorous formulation in computational terms, as seemingly contradictory hypotheses and findings can be formally unified under a more general theory. In particular, our analysis revealed that alternating between voluntary actions engenders task-set reconfiguration costs, whereas inter-trial inhibition of habitual responses can explain asymmetric switch costs.

ACKNOWLEDGEMENTS

This work was supported by the René and Susanne Braginsky Foundation (KES) and the University of Zurich (KES).

CONFLICT OF INTEREST

The authors declare no competing interests, financial or otherwise.

DATA ACCESSIBILITY

The data reported here are openly available at https://doi.org/doi:10.3929/ethz-b-000296409.

AUTHOR CONTRIBUTIONS

EAA designed the experiment, analysed the data, developed analytical tools and wrote the manuscript. KES provided the funding, designed the experiment and edited the manuscript. JH designed the experiment, analysed the data and wrote the manuscript.

ORCID

Eduardo A. Aponte https://orcid.org/0000-0001-8287-194X

REFERENCES

Allport, A., Styles, E. A., & Hsieh, S. (1994). Shifting intentional set: Exploring the dynamic control of tasks. In C. Umilta, & M. Moscovitch (Eds.), *Attention and Performance Series. Attention and Performance 15: Conscious and Nonconscious Information Processing* (pp. 266–290). Cambridge, MA: The MIT Press.

Altmann, E. M., & Gray, W. D. (2008). An integrated model of cognitive control in task switching. *Psychological Review*, *115*, 602–639.

Ansari, T. L., Derakshan, N., & Richards, A. (2008). Effects of anxiety on task switching: Evidence from the mixed antisaccade task. *Cognitive, Affective, & Behavioural Neuroscience*, *8*, 229–238.

Aponte, E. A., Raman, S., Sengupta, B., Penny, W. D., Stephan, K. E., & Heinzle, J. (2016). mpdcm: A toolbox for massively parallel dynamic causal modeling. *Journal of Neuroscience Methods*, *257*, 7–16.

Aponte, E. A., Schobi, D., Stephan, K. E., & Heinzle, J. (2017). The stochastic early reaction, inhibition, and late action (SERIA) model for antisaccades. *PLoS Computational Biology*, *13*, e1005692.

Aponte, E. A., Tschan, D. G., Heinzle, J., & Stephan, K. E. (2018). Inhibition and late errors in the antisaccade task: Influence of task design. *Journal of Neurophysiology*, *120*(6), 3001–3016.

Aron, A. R. (2011). From reactive to proactive and selective control: Developing a richer model for stopping inappropriate responses. *Biological Psychiatry*, *69*, 55–68.

Barton, J. J., Cherkasova, M. V., Lindgren, K., Goff, D. C., Intriligator, J. M., & Manoach, D. S. (2002). Antisaccades and task switching: Studies of control processes in saccadic function in normal subjects and schizophrenic patients. *Annals of the New York Academy of Sciences*, *956*, 250–263.

Barton, J. J., Greenzang, C., Hefter, R., Edelman, J., & Manoach, D. S. (2006). Switching, plasticity, and prediction in a saccadic task-switch paradigm. *Experimental Brain Research*, *168*, 76–87.

Barton, J. J., Raoof, M., Jameel, O., & Manoach, D. S. (2006). Task-switching with antisaccades versus no-go trials: A comparison of inter-trial effects. *Experimental Brain Research*, *172*, 114–119.

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, *67*, 1–48.

Bojko, A., Kramer, A. F., & Peterson, M. S. (2004). Age equivalence in switch costs for prosaccade and antisaccade tasks. *Psychology and Aging*, *19*, 226–234.

Calderhead, B., & Girolami, M. (2009). Estimating Bayes factors via thermodynamic integration and population MCMC. *Computational Statistics & Data Analysis*, *53*, 4028–4045.

Camalier, C. R., Gotler, A., Murthy, A., Thompson, K. G., Logan, G. D., Palmeri, T. J., & Schall, J. D. (2007). Dynamics of saccade target selection: Race model analysis of double step and search step saccade production in human and macaque. *Vision Research*, *47*, 2187–2211.

Chan, J. L., Koval, M. J., Johnston, K., & Everling, S. (2017). Neural correlates for task switching in the macaque superior colliculus. *Journal of Neurophysiology*, *118*, 2156–2170.

Cherkasova, M. V., Manoach, D. S., Intriligator, J. M., & Barton, J. J. (2002). Antisaccades and task-switching: Interactions in controlled processing. *Experimental Brain Research*, *144*, 528–537.

DeSimone, J. C., Weiler, J., Aber, G. S., & Heath, M. (2014). The unidirectional prosaccade switch-cost: Correct and error antisaccades differentially influence the planning times for subsequent prosaccades. *Vision Research*, *96*, 17–24.

Ethridge, L. E., Brahmbhatt, S., Gao, Y., McDowell, J. E., & Clementz, B. A. (2009). Consider the context: Blocked versus interleaved presentation of antisaccade trials. *Psychophysiology*, *46*, 1100–1107.

Everling, S., & Johnston, K. (2013). Control of the superior colliculus by the lateral prefrontal cortex. *Philosophical Transactions of the Royal Society of London. Series B, Biological sciences*, *368*, 20130068.

Fecteau, J. H., Au, C., Armstrong, I. T., & Munoz, D. P. (2004). Sensory biases produce alternation advantage found in sequential saccadic eye movement tasks. *Experimental Brain Research*, *159*, 84–91.

Fournier, D. A., Skaug, H. J., Ancheta, J., Ianelli, J., Magnusson, A., Maunder, M. N., … Sibert, J. (2012). AD Model Builder: Using automatic differentiation for statistical inference of highly parameterized complex nonlinear models. *Optimization Methods and Software*, *27*, 233–249.

Fox, J., & Weisberg, S. (2011). *An R companion to applied regression*, 2nd ed.. Thousand Oaks, CA: Sage.

Franke, C., Reuter, B., Breddin, A., & Kathmann, N. (2009). Response switching in schizophrenia patients and healthy subjects: Effects of the inter-response interval. *Experimental Brain Research*, *196*, 429–438.

Gelman, A., Carlin, J. B., Stern, H. S., & Rubin, D. B. (2003) *Bayesian data analysis*. London, UK: Chapman and Hall/CRC.

Gelman, A., Hwang, J., & Vehtari, A. (2014). Understanding predictive information criteria for Bayesian models. *Statistics and Computing*, *24*, 997–1016.

Gelman, A., & Meng, X. L. (1998). Simulating normalizing constants: From importance sampling to bridge sampling to path sampling. *Statistical Science*, *13*, 163–185.

Gelman, A., & Shalizi, C. R. (2013). Philosophy and the practice of Bayesian statistics. *British Journal of Mathematical and Statistical Psychology*, *66*, 8–38.

Ghodrati, M., Morris, A. P., & Price, N. S. (2015). The (un)suitability of modern liquid crystal displays (LCDs) for vision research. *Frontiers in Psychology*, *6*, 303.

Hallett, P. E. (1978). Primary and secondary saccades to goals defined by instructions. *Vision Research*, *18*, 1279–1296.

Heath, M., Gillen, C., & Samani, A. (2016). Alternating between pro- and antisaccades: Switch-costs manifest via decoupling the spatial relations between stimulus and response. *Experimental Brain Research*, *234*, 853–865.

Heath, M., Starrs, F., Macpherson, E., & Weiler, J. (2015). Task-switching effects for visual and auditory pro- and antisaccades: Evidence for a task-set inertia. *Journal of Motor Behavior*, *47*, 319–327.

Heinzle, J., Aponte, E. A., & Stephan, K. E. (2016). Computational models of eye movements and their application to schizophrenia. *Current Opinion in Behavioral Sciences*, *11*, 21–29.

Hikosaka, O., & Isoda, M. (2010). Switching from automatic to controlled behavior: Cortico-basal ganglia mechanisms. *Trends in Cognitive Sciences*, *14*, 154–161.

Hunt, A. R., & Klein, R. M. (2002). Eliminating the cost of task set reconfiguration. *Memory & Cognition*, *30*, 529–539.

Isoda, M., & Hikosaka, O. (2008). Role for subthalamic nucleus neurons in switching from automatic to controlled eye movement. *Journal of Neuroscience*, *28*, 7209–7218.

Jamadar, S. D., Johnson, B. P., Clough, M., Egan, G. F., & Fielding, J. (2015). Behavioral and Neural plasticity of ocular motor control: Changes in performance and fMRI activity following antisaccade training. *Frontiers in Human Neuroscience*, *9*, 653.

Karayanidis, F., Jamadar, S., Ruge, H., Phillips, N., Heathcote, A., & Forstmann, B. U. (2010). Advance preparation in task-switching: Converging evidence from behavioral, brain activation, and model-based approaches. *Frontiers in Psychology*, *1*, 25.

Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, *90*, 773–795.

Kiesel, A., Steinhauser, M., Wendt, M., Falkenstein, M., Jost, K., Philipp, A. M., & Koch, I. (2010). Control and interference in task switching–a review. *Psychological Bulletin*, *136*, 849–874.

Koch, I., Gade, M., Schuch, S., & Philipp, A. M. (2010). The role of inhibition in task switching: A review. *Psychonomic Bulletin & Review*, *17*, 1–14.

Lee, A. K., Hamalainen, M. S., Dyckman, K. A., Barton, J. J., & Manoach, D. S. (2011). Saccadic preparation in the frontal eye field is modulated by distinct trial history effects as revealed by magneto-encephalography. *Cerebral Cortex*, *21*, 245–253.

Logan, G. D., & Bundesen, C. (2003). Clever homunculus: Is there an endogenous act of control in the explicit task-cuing procedure? *Journal of Experimental Psychology: Human Perception and Performance*, *29*, 575–599.

Logan, G. D., Cowan, W. B., & Davis, K. A. (1984). On the ability to inhibit simple and choice reaction time responses: A model and a method. *Journal of Experimental Psychology: Human Perception and Performance*, *10*, 276–291.

Luke, S. G. (2017). Evaluating significance in linear mixed-effects models in R. *Behavior Research Methods*, *49*, 1494–1502.

Manoach, D. S., Lindgren, K. A., & Barton, J. J. (2004). Deficient saccadic inhibition in Asperger's disorder and the social-emotional processing disorder. *Journal of Neurology, Neurosurgery and Psychiatry*, *75*, 1719–1726.

Manoach, D. S., Lindgren, K. A., Cherkasova, M. V., Goff, D. C., Halpern, E. F., Intriligator, J., & Barton, J. J. (2002). Schizophrenic subjects show deficient inhibition but intact task switching on saccadic tasks. *Biological Psychiatry*, *51*, 816–826.

Manoach, D. S., Thakkar, K. N., Cain, M. S., Polli, F. E., Edelman, J. A., Fischl, B., & Barton, J. J. (2007). Neural activity is modulated by trial history: A functional magnetic resonance imaging study of the effects of a previous antisaccade. *Journal of Neuroscience*, *27*, 1791–1798.

Mathys, C. D., Lomakina, E. I., Daunizeau, J., Iglesias, S., Brodersen, K. H., Friston, K. J., & Stephan, K. E. (2014). Uncertainty in perception and the hierarchical Gaussian filter. *Frontiers in Human Neuroscience*, *8*, 825.

Meiran, N. (1996). Reconfiguration of processing mode prior to task performance. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 1423.

Monsell, S. (2003). Task switching. *Trends in Cognitive Sciences*, *7*, 134–140.

Mueller, S. C., Swainson, R., & Jackson, G. M. (2009). ERP indices of persisting and current inhibitory control: A study of saccadic task switching. *NeuroImage*, *45*, 191–197.

Munoz, D. P., & Everling, S. (2004). Look away: The anti-saccade task and the voluntary control of eye movement. *Nature Reviews Neuroscience*, *5*, 218–228.

Noorani, I., & Carpenter, R. H. (2013). Antisaccades as decisions: LATER model predicts latency distributions and error responses. *European Journal of Neuroscience*, *37*, 330–338.

Peirce, J. W. (2007). PsychoPy–Psychophysics software in python. *Journal of Neuroscience Methods*, *162*, 8–13.

Peirce, J. W. (2008). Generating stimuli for neuroscience using psychopy. *Frontiers in Neuroinformatics*, *2*, 10.

Pierce, J. E., McCardel, J. B., & McDowell, J. E. (2015). Trial-type probability and task-switching effects on behavioral response characteristics in a mixed saccade task. *Experimental Brain Research*, *233*, 959–969.

Pierce, J. E., & McDowell, J. E. (2017). Contextual effects on cognitive control and BOLD activation in single versus mixed saccade tasks. *Brain and Cognition*, *115*, 12–20.

Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion decision model. Current issues and history. *Trends in Cognitive Sciences*, *20*, 260–281.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. *Classical Conditioning II: Current Research and Theory*, *2*, 64–99.

Reuter, B., Philipp, A. M., Koch, I., & Kathmann, N. (2006). Effects of switching between leftward and rightward pro- and antisaccades. *Biological Psychology*, *72*, 88–95.

Rivaud-Pechoux, S., Vidailhet, M., Brandel, J. P., & Gaymard, B. (2007). Mixing pro- and antisaccades in patients with parkinsonian syndromes. *Brain*, *130*, 256–264.

Rogers, R. D., & Monsell, S. (1995). Costs of a predictable switch between simple cognitive tasks. *Journal of Experimental Psychology: General*, *124*, 207.

Schall, J. D., Palmeri, T. J., & Logan, G. D. (2017). Models of inhibitory control. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *372*, 20160193.

Schmitz, F., & Voss, A. (2014). Components of task switching: A closer look at task switching and cue switching. *Acta Psychologica*, *151*, 184–196.

Stampe, D. (1993). Heuristic filtering and reliable calibration methods for video-based pupil-tracking systems. *Behavior Research Methods, Instruments, & Computers*, *25*, 137–142.

Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J., & Friston, K. J. (2009). Bayesian model selection for group studies. *NeuroImage*, *46*, 1004–1017.

Watanabe, S. (2010). Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory. *Journal of Machine Learning Research*, *11*, 3571–3594.

Weiler, J., Hassall, C. D., Krigolson, O. E., & Heath, M. (2015). The unidirectional prosaccade switch-cost: Electroencephalographic evidence of task-set inertia in oculomotor control. *Behavioral Brain Research*, *278*, 323–329.

Weiler, J., & Heath, M. (2012a). Task-switching in oculomotor control: Unidirectional switch-cost when alternating between pro- and antisaccades. *Neuroscience Letters*, *530*, 150–154.

Weiler, J., & Heath, M. (2012b). The prior-antisaccade effect influences the planning and online control of prosaccades. *Experimental Brain Research*, *216*, 545–552.

Weiler, J., & Heath, M. (2014a). Oculomotor task switching: Alternating from a nonstandard to a standard response yields the unidirectional prosaccade switch-cost. *Journal of Neurophysiology*, *112*, 2176–2184.

Weiler, J., & Heath, M. (2014b). Repetitive antisaccade execution does not increase the unidirectional prosaccade switch-cost. *Acta Psychologica*, *146*, 67–72.

Wylie, G., & Allport, A. (2000). Task switching and the measurement of "switch costs". *Psychological Research*, *63*, 212–233.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of the article.

How to cite this article: Aponte EA, Stephan KE, Heinzle J. Switch costs in inhibitory control and voluntary behaviour: A computational study of the antisaccade task. *Eur J Neurosci*. 2019;50:3205–3220. https://doi.org/10.1111/ejn.14435