



## BIROn - Birkbeck Institutional Research Online

Sexton, Nicholas J. and Cooper, Richard P. (2017) Task inhibition, conflict, and the n-2 repetition cost: A combined computational and empirical approach. *Cognitive Psychology* 94 , pp. 1-25. ISSN 0010-0285.

Downloaded from: <https://eprints.bbk.ac.uk/id/eprint/17947/>

*Usage Guidelines:*

Please refer to usage guidelines at <https://eprints.bbk.ac.uk/policies.html>  
contact [lib-eprints@bbk.ac.uk](mailto:lib-eprints@bbk.ac.uk).

or alternatively



ELSEVIER

Contents lists available at [ScienceDirect](http://www.sciencedirect.com)

## Cognitive Psychology

journal homepage: [www.elsevier.com/locate/cogpsych](http://www.elsevier.com/locate/cogpsych)

# Task inhibition, conflict, and the n-2 repetition cost: A combined computational and empirical approach <sup>☆</sup>

Nicholas J. Sexton <sup>\*</sup>, Richard P. Cooper

Centre for Cognition, Computation and Modelling, Department of Psychological Sciences, Birkbeck, University of London, United Kingdom

## ARTICLE INFO

## Article history:

Accepted 13 January 2017

Available online 21 February 2017

## ABSTRACT

Task inhibition (also known as backward inhibition) is an hypothesised form of cognitive inhibition evident in multi-task situations, with the role of facilitating switching between multiple, competing tasks. This article presents a novel cognitive computational model of a backward inhibition mechanism. By combining aspects of previous cognitive models in task switching and conflict monitoring, the model instantiates the theoretical proposal that backward inhibition is the direct result of conflict between multiple task representations. In a first simulation, we demonstrate that the model produces two effects widely observed in the empirical literature, specifically, reaction time costs for both (n-1) task switches and n-2 task repeats. Through a systematic search of parameter space, we demonstrate that these effects are a general property of the model's theoretical content, and not specific parameter settings. We further demonstrate that the model captures previously reported empirical effects of inter-trial interval on n-2 switch costs. A final simulation extends the paradigm of switching between tasks of asymmetric difficulty to three tasks, and generates novel predictions for n-2 repetition costs. Specifically, the model predicts that n-2 repetition costs associated with hard-easy-hard alternations are greater than for easy-hard-easy alternations. Finally, we report two behavioural experiments testing this hypothesis, with results consistent with the model predictions.

© 2017 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

## 1. Introduction

A fundamental requirement of much human behaviour in a dynamic world is the ability to flexibly switch between tasks in response to environmental context. Thus when cooking one may need to alternate between slicing onions and regularly stirring a simmering pan. The control processes involved in regulating this form of behaviour have been studied in the laboratory using so-called task-switching paradigms, where participants typically alternate performing each of two or more possible tasks afforded by a stimulus, for example classifying digit stimuli by magnitude (high/low) or parity (odd/even) (see [Kiesel et al., 2010](#); [Vandierendonck, Liefooghe, & Verbruggen, 2010](#), for reviews).

In a typical task-switching experiment the ordering of tasks is such that each trial requires either a switch or a repetition of the preceding task. An early observation was that task switches were associated with costs, in terms of response times and error rates, compared to task repeats. Numerous explanations for the switch cost sought to attribute it to cognitive processes required on switch trials, but not repeat trials, such as an executive mechanism responsible for reconfiguring the cognitive system appropriately for the new task (e.g., [Rogers & Monsell, 1995](#); [Rubinstein, Meyer, & Evans, 2001](#); [Sohn & Anderson,](#)

<sup>☆</sup> This work was supported by the Economic and Social Research Council [grant number ES/J500021/1].

<sup>\*</sup> Corresponding author.

E-mail address: [n.sexton@mail.bbk.ac.uk](mailto:n.sexton@mail.bbk.ac.uk) (N.J. Sexton).

2001). However, alternative accounts suggested that the switch cost does not directly reflect control processes, but results from interference between the two task representations (e.g., Allport, Styles, & Hsieh, 1994; Allport & Wylie, 2000; Altmann & Gray, 2008; Yeung & Monsell, 2003) with a possible middle ground being that interference is itself mediated by control processes (e.g., the suggestion by Goschke, 2000, that inhibition of tasks is regulated by response conflict).

Residual inhibition of the now-relevant task-set is a common explanatory concept for a number of interference-based accounts of the switch cost (Allport et al., 1994; Allport & Wylie, 2000; Gilbert & Shallice, 2002; Mayr & Keele, 2000). However, task interference can also be conceptualised in activation-only terms (Altmann & Gray, 2008; Yeung & Monsell, 2003), and non-inhibitory, non-interference based accounts also exist. In these accounts, switch costs are a consequence of cue-switching, and may not necessarily reflect a task switch (Logan & Bundesen, 2003; Logan & Schneider, 2010). Thus, while it remains plausible that task switching involves task inhibition, switch costs alone do not provide strong evidence for such processes (for recent reviews, see Gade, Schuch, Druery, & Koch, 2014; Koch, Gade, Schuch, & Philipp, 2010).

More promising behavioural evidence for task inhibition comes from a procedure in which participants switch between three possible tasks (e.g., Mayr & Keele, 2000). For example, in response to single-digit stimuli, excluding 5, participants may switch between classifying the stimuli according to parity (odd/even), magnitude (greater/less than 5) or centrality (near/far from 5) (e.g., Schuch & Koch, 2003). In most examples of this paradigm, task repeats are avoided, hence each trial ( $n$ ) represents a switch compared to the previous ( $n-1$ ) trial. The main dependent measures are the contrast in RT and error rates between 'n-2 repeats' – trials in which a task re-occurs after a single intervening trial (i.e., given switching between three tasks,  $A$ ,  $B$  and  $C$ , the third trial in the task sequence  $ABA$ ) – and 'n-2 switches', in which the task has not been performed for two or more trials (such as the third trial in  $CBA$  sequences). Assuming persistent activation (i.e., priming) of task sets, one might expect n-2 repeats to be facilitated by more recent performance of the task, compared to n-2 switches. Instead, the robust finding is that n-2 repeats are significantly slower, and more error prone, than n-2 switches. To date, these n-2 repetition costs have been resistant to alternative explanations, with a consensus that the effect indeed reflects task inhibition (Koch et al., 2010). One suggestion is that task inhibition, also referred to as 'backward inhibition', forms part of a cognitive mechanism for sequential task control (Mayr & Keele, 2000). Such a mechanism would have the functional benefit of reducing cross-task interference from highly active, yet goal-irrelevant, task-sets in multiple-task situations.

However, while task inhibition has commonly been invoked in verbal descriptions of task switching, formal accounts of the mechanism remain at an early stage (e.g., Grange & Juvina, 2015). To date, a formal model has yet to be specified which shows how task inhibition affects performance in both two and three-task paradigms, thus clarifying the contribution (if any) of backward inhibition to the  $n-1$  switch cost. The remainder of this article describes and evaluates such a model. The model extends the interactive activation model of Gilbert and Shallice (2002) with a simple task inhibition mechanism sensitive to task conflict. It represents a formalisation of the idea that task inhibition occurs in response to conflicting task representations (Goschke, 2000; Koch et al., 2010).

In a first pair of simulations, the model demonstrates that backward inhibition can explain the main qualitative pattern – costs for both  $n-1$  switches and  $n-2$  repeats – and that these effects are a property of the model architecture and not of specific parameter settings (given certain theoretically-motivated constraints). A third simulation explores the effects of variable inter-trial intervals on  $n-1$  and  $n-2$  costs, replicating existing empirical results. Finally with respect to the simulation work, and as backward inhibition in the model is sensitive to inter-task conflict, we explore a version of task switching where three tasks are of different (i.e., asymmetric) difficulty. In such a procedure, the inter-task conflict generated by hard-easy-hard alternations may be contrasted with that of easy-hard-easy alternations. The model makes clear behavioural predictions for differential effects on  $n-2$  repetition costs. We then report two novel three-task switching experiments which successfully test these predictions. In the final part of the paper, the implications of the model are discussed, including the functional benefit provided by the backward inhibition mechanism. Specifically, we argue that backward inhibition facilitates task control by automatically reducing cross-task interference, thereby reducing or eliminating the need for attentionally-demanding deliberate control in some situations.

## 2. Existing computational accounts of task switching

To date, a variety of computational accounts of task switching have been proposed, drawing either on the principles of Interactive Activation and Competition (IAC) (Brown, Reynolds, & Braver, 2007; Gilbert & Shallice, 2002) or developed within the ACT-R architecture (Altmann & Gray, 2008; Sohn & Anderson, 2001, see Grange & Houghton, 2014, for a recent review and theoretical comparison). Only one has extended to three-task switching and  $n-2$  effects (Grange & Juvina, 2015; Grange, Juvina, & Houghton, 2013).

### 2.1. The IAC model of Gilbert and Shallice (2002)

Gilbert and Shallice (2002) proposed an Interactive Activation and Competition (IAC) model of task switching based on earlier IAC models of the Stroop task (Cohen, Dunbar, & McClelland, 1990; Cohen & Huston, 1994). Within the model stimuli are represented as single units for each possible alternative across its two possible dimensions ('word' and 'ink colour'), with possible responses also represented as single units. The two alternative tasks thus form two segregated processing pathways. The weights of the associative links from stimulus input to response units are greater in the word reading than the colour

naming pathway, reflecting greater experience with this task. Two ‘task demand’ units provide control of the current task, analogous to task-sets. These have excitatory connections with all response units for their respective task and inhibitory connections with the alternative task response units.<sup>1</sup> As is typical with IAC models, lateral inhibition is implemented between units at the same level, meaning that units within a level compete to become the most highly active, while simultaneously suppressing activation in the others. The overall effect of task demand unit activation, therefore, is general facilitation of the respective task demand unit pathway and inhibition of the alternative task pathway. In the absence of intervention from higher-level control processes, the network produces responses from the stronger, word reading pathway. To perform task switching, a top-down control input is supplied to the relevant task-demand unit on each trial. As more biasing is required to perform the weaker task than the stronger one, the top-down-control input is stronger for colour naming than word reading.

On a typical trial, stimulus input units are activated for each processing pathway (e.g., the red colour unit and the green word unit representing the word ‘green’ incongruously displayed in red), with activation propagating through both pathways. Simultaneously, a single task-demand unit becomes active through top-down-control excitation. This biases processing in favour of the current task, while suppressing processing in the competing task.

In the model, interference from previous tasks occurs due to residual activation of the task-demand units from the previous trial. Repeat trials are therefore facilitated, as the relevant task-demand unit is already highly active. Interference occurs on switch trials, as residual task-demand activation now facilitates processing for the irrelevant task, and a greater period of processing is required for top-down control to re-activate the relevant task-demand unit against irrelevant task interference, producing switch costs. Importantly, as the task-demand unit for the weaker task requires greater activation to perform the task than that for the well-learned task, a greater degree of residual activation is carried over to the following trial, producing greater switch costs for hard-easy switches than for easy-hard switches. This phenomenon is known in the empirical task-switching literature as an asymmetric switch cost (see [Kiesel et al., 2010](#), for a review).

The model of [Gilbert and Shallice \(2002\)](#) has not been extended to the n-2 repetition paradigm. However, while the model lacks any explicit control processes, its top-down input is regarded as a proxy for the output of higher-level control processes which bias processing in the production of responses – an arrangement analogous to the two-level contention scheduling and supervisory systems proposed by [Norman and Shallice \(1986\)](#). In theoretical and practical terms, the model therefore makes a suitable platform for implementing a putative task inhibition mechanism.

## 2.2. The IAC model of [Brown et al. \(2007\)](#)

[Brown et al. \(2007\)](#) have proposed a model of task-switching that, while conceptually similar to the model of [Gilbert and Shallice \(2002\)](#), is substantially more complex.

It incorporates two additional cognitive control mechanisms. Firstly, an incongruity detector signals conflict between co-active incompatible responses, and in response sends additional excitation to the currently active task unit. Secondly, a change detector responds to trial-by-trial changes in either the task units or the response units, by removing a ‘tonic arousal signal’ – an effect which slows all responses and lasts for a number of trials. Through slowing of responses, this latter mechanism also leads to increased accuracy. Thus, the model reacts to sequential conflict (generated between current trial processing and residual activation from previous trials) by slowing all responses, while within-trial conflict (between simultaneously active responses) is resolved by reinforcing the activation of the currently active task unit. While this model is noteworthy in integrating conflict-monitoring mechanisms with a task switching model, the authors do not explore the model’s performance in the n-2 repetition paradigm, and instead focus on sequential stimulus congruency and response repetition effects. The model fits the authors’ target data, reproducing a complex behavioural pattern, however it introduces a high level of complexity and requires hand-setting of 21 free parameters.

## 2.3. The model of [Altmann and Gray \(2008\)](#)

In the model of [Altmann and Gray \(2008\)](#), based on the ACT-R cognitive architecture, task switching is conceptualised as a memory retrieval process. Task cues are encoded as an episodic memory trace, used as the basis for retrieving the relevant stimulus dimension (e.g., odd/even), identifying the correct category (odd) and making a response (left). As is standard in all ACT-R models, the time taken to access representations (‘chunks’) from memory is based on their activation. ACT-R’s base-level-learning (BLL) equation, which governs chunk activation, ensures that the activation of a chunk increases each time it is accessed, after which it decays according to a power law. Given this formulation, task repeats are faster than task switches because the relevant task representations have been accessed more recently – a form of priming.

As with all ACT-R models, the activation of chunks in declarative memory flows associatively. Thus, activation of a task dimension (parity) spreads to associated dimensions (odd/even), and then to the associated responses (left/right). Conceptually, then, the explanation for switch costs in the model is similar to that of the model of [Gilbert and Shallice \(2002\)](#). It is due to residual activation/priming of recently performed tasks and the related representations (although in the Gilbert and

<sup>1</sup> The model also involves a variably-weighted connection from stimulus input units to task-demand units – stimuli occurring in the previous trial evoke the task that occurred on that trial. While this mechanism, simulating a short-term priming effect, accounts for an item-specific component of the switch cost ([Allport & Wylie, 2000](#)), we do not consider this component of the switch cost further in this article.

Shallice model there is a greater emphasis on between-task interference, which occurs due to lateral inhibition from residually active irrelevant representations).

#### 2.4. The task inhibition models of Grange and colleagues

Given the absence of inhibitory processes within the ACT-R architecture, Grange et al. (2013) argued that the model of Altmann and Gray (2008) cannot account for n-2 repetition costs, and would instead predict n-2 repetition facilitation, for similar reasons as for n-1 repeats. Grange et al. (2013) therefore implemented a model within ACT-R by modifying the BLL equation to include an early, inhibitory component (i.e., self-inhibition). This component is large immediately after chunk retrieval, making chunks harder to retrieve immediately, but dissipates passively over time, leading to an inverted-U shaped activation curve. In this model, the locus of the n-2 repetition cost is the process of retrieving target representations from memory – for a certain period this process is slower for more recently retrieved target representations, as in the n-2 repetition condition.

Using this model, Grange and Juvina (2015) predicted that n-2 repetition costs should be specific to relatively novel tasks: when tasks become relatively automatised after a large amount of practice, n-2 repetition costs should be reduced or eliminated. This is due to the increase in the base-level activation with learning for each task representation, which makes well-learned tasks easier to retrieve. In an experiment with 10 participants completing 6100 trials over five days, the average n-2 repetition cost did significantly decrease with practice. However, the model predicted a smaller reduction in costs than was observed empirically. To fit the data, Grange and Juvina (2015) further assumed that cue-target associations also become stronger as a result of practice.

While modelling task inhibition as a form of self-inhibition provides a good fit to the data, whether it represents a viable theoretical proposal remains unresolved (see Grange et al., 2013; Koch et al., 2010, for a range of arguments).

However, Grange and colleagues themselves concede that a limitation of their model is that it cannot simultaneously account for both empirically observed n-1 and n-2 effects of task switching; the model with the inhibitory version of the BLL equation predicts n-1 switch *facilitation* (or equivalently, n-1 repetition costs) rather than n-1 switch costs.

#### 2.5. Interim summary

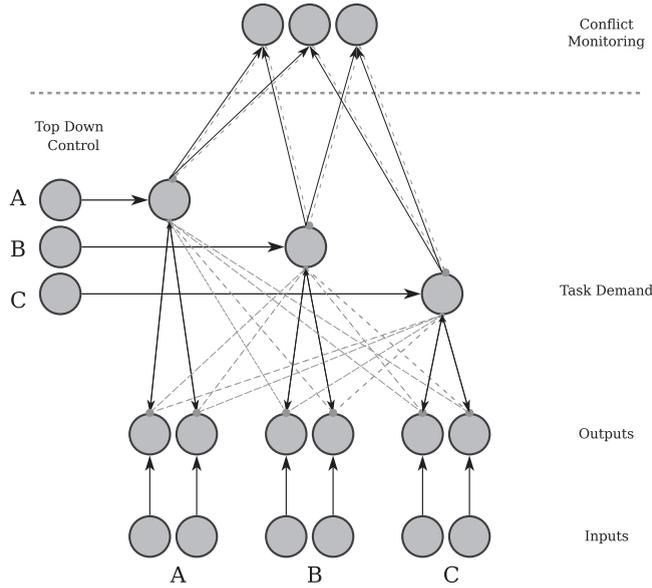
None of the models discussed above provides a complete account of n-1 and n-2 effects in task switching. The model of Altmann and Gray (2008) explains n-1 switch costs through priming, but the dynamics of its activation equations suggest that the model would (incorrectly) predict n-2 repetition facilitation and n-2 switch costs. Grange and colleagues (Grange & Juvina, 2015; Grange et al., 2013) present a model of the n-2 repetition paradigm by introducing a short-term inhibitory component to the ACT-R BLL equation, but the model does not fully explain how or why this occurs, and with the modified activation equations the model incorrectly predicts n-1 switch facilitation and n-1 repetition costs. An ACT-R model that is capable of producing both empirical n-1 and n-2 effects remains a future prospect. The model of Gilbert and Shallice (2002) successfully simulates a range of n-1 switch effects, but has not yet been adapted to the n-2 paradigm. It also lacks any high-level cognitive mechanisms to detect conflict or inhibit task sets. Thus, while it incorporates lateral inhibition between units at the same level (i.e., between incompatible responses), it is unlikely this is sufficient to produce the n-2 repetition cost (see Grange et al., 2013). Finally, while the model of Brown et al. (2007) includes a range of conflict monitoring mechanisms and would seem well suited to addressing the n-2 repetition cost paradigm and the wide range of empirical findings concerning the role of conflict, the complexity of the model is a disadvantage for the current purpose (i.e., of specifying the minimal cognitive mechanisms necessary and sufficient for backward inhibition). Even if the model of Brown et al. (2007) is able to produce the empirical pattern of behaviour, the wide range of cognitive mechanisms postulated within it may be superfluous to any explanation (i.e., they may be sufficient, but not necessary), and without a systematic approach to fitting the large number of free parameters, it is unclear what other patterns of behaviour the model is capable of fitting (cf. Roberts & Pashler, 2000). In the following section we therefore present an alternative model of task switching which incorporates a backward inhibition mechanism and which, as is shown in subsequent sections, is capable of accounting for both n-1 and n-2 effects.

### 3. Model description

The present theoretical perspective conceptualises backward inhibition as a cognitive control mechanism that operates by biasing processing between multiple task pathways, on the basis of conflict at the level of task representations. The model extends the interactive activation model of Gilbert and Shallice (2002), which in itself is analogous to a lower-level contention scheduling system (Cooper & Shallice, 2000; Norman & Shallice, 1986), with an additional task inhibition mechanism.

#### 3.1. Model architecture

The model architecture is illustrated in Fig. 1 The portion of the figure below the dashed line is equivalent to the model of Gilbert and Shallice (2002) applied to three tasks. Specifically, it is implemented as an interactive activation model in which



**Fig. 1.** Architecture of the model, capable of switching between three tasks (A, B & C). Excitatory connections are represented as lines in solid black with sharp arrows, inhibitory connections are dashed grey lines with circular arrows. Arrowheads show the direction of the connection. Not shown are within-level connections (e.g., lateral inhibition).

processing in each unit of the model is allowed to bias processing at other levels (McClelland, 1993). The model has three sets of input and output units, corresponding to each of three tasks (referred to in Fig. 1 as tasks A, B and C). For each task, two input units correspond to the two possible input values. For example, the input stimulus ‘9’, affording the three tasks *parity*, *magnitude* and *centrality*, might be represented as *odd* (left), *high* (right) and *peripheral* (right). Input units have feed-forward connections to the corresponding output unit. There are six output units, representing two possible responses to each of the three tasks. Within the set of output units, units that correspond to the same response are mutually excitatory, while units that correspond to alternative responses are mutually inhibitory.

The model therefore implements an experimental procedure in which responses for all tasks are mapped to the same set of response keys (cf. Gade & Koch, 2007b).

The strength of the connection weights between input and output units represents the strength of ‘bottom-up’ processing in the model: a well-learned or prepotent task (e.g., word reading of Stroop stimuli) has stronger connection weights than a weaker task (e.g., colour naming).<sup>2</sup>

### 3.2. Task demand units

As in the model of Gilbert and Shallice (2002), task processing is biased by task-demand units – one per task – which have excitatory (i.e., positively weighted) connections to their respective output units and inhibitory (i.e., negatively weighted) connections to the output units associated with the other tasks. These connections are bi-directional, so a response activated bottom-up by a strong input connection will tend to activate the task with which the response is associated. The currently relevant task unit also receives a ‘top-down control’ input, from units represented on the left of Fig. 1, simulating the level of deliberate control required to perform each task to a reasonable level. As a simplification, the model omits the modifiable connections from input to task demand units, which Gilbert and Shallice (2002) used to simulate the influence of item-specific priming. While it should be stressed that task demand units are ordinary units, with input and activation calculated in the same way as all other units in the model (excepting conflict monitoring units), they take input from a range of sources within the model. Explicitly, the input ( $I_{td}$ ) may be expressed in the following terms:

$$I_{td} = S_{tdc} + \sum_c \alpha_c \omega_c + \sum_o \alpha_o \omega_o \tag{1}$$

where  $S_{tdc}$  is the top down control strength parameter for the respective task,  $\alpha_c$  and  $\omega_c$  are the activation and weights for the two conflict units connected to the unit, and  $\alpha_o$  and  $\omega_o$  are the activation and connection weights of the six output units in the model.

<sup>2</sup> Unless otherwise specified, the weights of connection are fixed and take the default values used by Gilbert and Shallice (2002). Moreover, while training of connection weights in the model did not occur, previous work has demonstrated that a similar architecture can learn this arrangement of connection weights when presented with training sets biased toward particular tasks (Cohen et al., 1990).

### 3.3. Conflict monitoring units

The current model extends previous models by introducing an upper level of conflict monitoring units, operating similarly to those in the model of Botvinick, Braver, Barch, Carter, and Cohen (2001). The input ( $I_c$ ) to these units is somewhat different to elsewhere in the model. Each monitors the conflict (i.e., simultaneous activation) between two Task Demand (TD) units, according to Eq. 2. In this equation, the activations of each task demand unit ( $\alpha_1, \alpha_2$ ) are linearly rescaled from  $(-1, 1)$  to  $(0, 1)$  to ensure the minimum activation value is zero. The input to conflict units ( $I_c$ ) on each cycle is then calculated as the product of these rescaled values, multiplied by a gain parameter ( $\gamma_c$ ), plus a constant negative bias ( $\beta_c$ ).

$$I_c = \gamma_c \left( \frac{\alpha_1 + 1}{2} \right) \left( \frac{\alpha_2 + 1}{2} \right) + \beta_c \quad (2)$$

Thus, in the absence of conflict, the conflict units receive a constant negative baseline input ensuring their activation decreases. If the activation values of two TD units are both greater than zero, the mutually connected conflict unit will receive an input greater than this baseline.

Unlike the model of Botvinick et al. (2001), conflict units bias model processing interactively, via inhibitory connections to both respective task demand units multiplied by a weight parameter ( $\omega$ ).<sup>3</sup> In sum, the model instantiates a theoretical position similar to the proposal by Koch et al. (2010), i.e., that task inhibition is recruited by conflict generated during task processing.

### 3.4. Activation calculation

The equation for calculating unit activation is unchanged from the model of Gilbert and Shallice (2002). For each unit  $i$ , the change in activation value on each model cycle,  $\Delta\alpha_i$ , is calculated as follows:

$$\Delta\alpha_i = \begin{cases} \sigma I_i (\alpha_{max} - \alpha_i) + \mu & \text{if } I_i \geq 0 \\ \sigma I_i (\alpha_i - \alpha_{min}) + \mu & \text{if } I_i < 0 \end{cases} \quad (3)$$

where  $\alpha_i$  is the unit's current activation,  $I_i$  is its net input, and  $\sigma$ ,  $\alpha_{max}$  and  $\alpha_{min}$  are parameters affecting step size (0.0015 for all simulations reported here), and maximum (1.0) and minimum ( $-1.0$ ) unit activation values respectively.  $\mu$  is a noise term, drawn from a Gaussian distribution, with a mean of 0 and standard deviation of 0.006.<sup>4</sup>

### 3.5. Processing within the model

On the first trial of each block, all units are initialised with zero activation. On subsequent trials, TD units carry over 20% of their activation from the final step of the previous trial (as in the original model of Gilbert & Shallice, 2002) and conflict units carry over 50% of their activation. This models the effects of residual task inhibition. All other units are initialised as for the first trial.

In a simulated trial, one input unit in each task pathway (representing a trivalent stimulus) and a top-down control unit (representing the currently cued task) are set to 1. Activation then iteratively propagates throughout the model. As response units suppress their non-congruent competitors via lateral inhibition, they exhibit a 'winner-takes-all' property, in which small differences in activation between the 'winning' and 'losing' units become decisive as the losing units are actively suppressed. A response is therefore taken as the most active output unit, when its activation exceeds that of the next most active, non-congruent output unit by a response threshold of 0.15. The number of cycles taken for this to occur is the simulated response time (RT).

## 4. Simulations

The model was evaluated in four simulations using an analog of a behavioural paradigm allowing for calculation of both  $n-1$  and  $n-2$  switch effects. All simulations were run on variably sized blocks of trials, where each trial was one of three tasks (A, B or C). Each trial was also categorised into a switch condition on the basis of the number and type of task switches on the preceding two trials, with the dependent variable being the simulated RT and error rate for each condition. In the general case there are five types of triplet: OSW (where there is a switch after trial  $n-2$  but trial  $n-1$  and trial  $n$  are the same – ABB, ACC, BAA, BCC, CAA, CBB); 1SW (where there is no switch from trial  $n-2$  to trial  $n-1$  but a switch from trial  $n-1$  to trial  $n$  – AAB, AAC, BBA, BBC, CCA, CCB); 2SW (where each trial is a switch and no task is repeated – ABC, ACB, BAC, BCA, CAB, CBA); ALT (where the tasks alternate – ABA, ACA, BAB, BCB, CAC, CBC); and BLK (where one task repeats three times – AAA, BBB, CCC). Paralleling prior empirical work, the  $n-1$  switch cost was calculated as the difference between 1-switch (1SW) and

<sup>3</sup> Given that unit activation varies between  $-1$  and  $1$ , for conflict units exclusively a threshold of zero was applied, such that only above-zero unit activations affect task demand units (as inhibition, given that the weight of the connection is negative). This prevents 'negative conflict' from exciting task demand units via inhibitory connections.

<sup>4</sup> These parameter values were the defaults used by Gilbert and Shallice (2002).

0-switch blocks (OSW), in which the final trial was a task switch (e.g., AAB) versus a task repeat (e.g., ABB). The n-2 repetition cost was calculated as the difference between final trial RT on blocks where the final trial, with respect to the second preceding (n-2) trial, was an alternating (ALT) switch (e.g., ABA) versus a non-alternating (or two-switch: 2SW) switch (e.g., CBA). If the correct response was not made within 500 cycles the trial was classified as an error and the block discarded from the RT analysis.<sup>5</sup>

Running the model requires a number of parameter values to be specified. In addition to those shared with the model of Gilbert and Shallice (2002), which took the default values used in that previous work, the conflict monitoring layer introduces: (a) an additional *conflict squashing* parameter that controls the amount of residual conflict activation carried over from one trial to the next within a block (set to 50% for all simulations); and (b) three further parameters (*gain*, *bias*, and *weight*) that determine the steepness of the response function and the strength of the effect of the conflict units (as described above). In simulation 1, these parameters were hand-set to assess whether the model was capable of producing the empirically observed pattern of behaviour. In simulations 2, 3 and 4, the model's parameters were explored more systematically.

#### 4.1. Simulation 1: Basic effects

Simulation 1 aimed to explore the qualitative effect of the proposed inhibitory control mechanism. Two versions of the model were therefore compared: a no-inhibition version, where the weight of inhibitory conflict task-demand connections was zero (effectively, a three-task version of the Gilbert & Shallice (2002) model) and a model where the value of the weight parameter ( $\omega$ ) was  $-20.0$ . Conflict unit gain ( $\gamma$ ) and bias ( $\beta_c$ ) were  $50.0$  and  $-10.0$ , respectively, for both models.

##### 4.1.1. Method

In order to evaluate model performance over a large block of consecutive trials, paralleling a typical behavioural paradigm, each version of the model was run on 100 blocks consisting of 600 trials with re-initialisation of the model occurring between blocks but not between trials. The sequence of tasks was generated randomly, with all three tasks having an equal probability of occurring on any given trial. As in previous empirical research, each trial was categorised according to the task sequence formed by it and its two preceding trials as described above.

##### 4.1.2. Results

Mean response times and error rates for each sequence of switches were calculated for each block. Fig. 2 plots the mean and 95% confidence intervals for all blocks, for both versions of the model (without and with inhibition) and triplet type. Confidence intervals allow direct inference of statistically significant differences from the figure.

Firstly, consider the no-inhibition version of the model. For mean reaction times, there is no statistically significant difference between repeats where the previous switch was recent (OSW) or more distant (BLK), inferred from 95% confidence intervals for each condition which include the other condition's mean. Similarly, 1SW is significantly slower than OSW, i.e., an n-1 switch cost. These two findings replicate the results of Gilbert and Shallice (2002). There is also little difference between 1SW and 2SW conditions, while the ALT condition is faster than 2SW, i.e., a slight n-2 repetition facilitation (or equivalently, a negative n-2 repetition cost). Due to the stochastic nature of the data, more variance is present in the error rate data, as reflected by larger confidence intervals. Nevertheless, a similar pattern is observed: No significant difference between BLK or OSW repeats, a significant n-1 switch cost, but no significant difference between any of the switching conditions (1SW, 2SW or ALT).

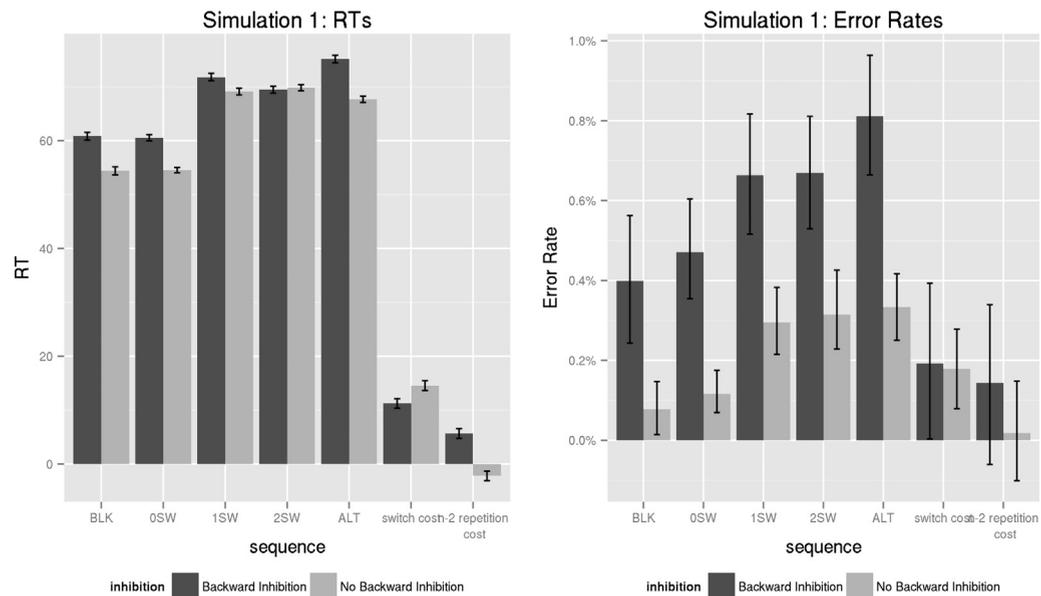
Next, consider the backward inhibition model. For reaction times, again, no significant differences are observed between different repeat conditions (BLK or OSW). As in the no-inhibition model, 1SW sequences are slower than OSW, producing a significant n-1 switch cost. However the switch cost is reduced compared to the no-inhibition model, due to larger OSW RTs and smaller 1SW RTs than the no-inhibition model. Effectively, switches are facilitated, at the cost of slower repeat trials. However, ALT switches are significantly slower than 2SW sequences, i.e., a positive n-2 repetition cost, in contrast to n-2 facilitation in the no-inhibition model.

##### 4.1.3. Discussion

Simulation 1 demonstrates that the basic task switching model (analogous to the model of Gilbert & Shallice, 2002), which incorporates lateral inhibition at the response and task demand level, but no backward inhibition, predicts n-1 switch costs, but negative n-2 repetition costs, for both reaction times and error rates. This finding strengthens the conclusion of Grange et al. (2013), who argued against lateral inhibition as a plausible mechanism for backward inhibition. They argued that lateral inhibition alone was incapable of producing lasting, residual inhibition of task demand units on the basis of activity on the n-2 trial. In contrast, the backward inhibition model predicts both n-1 switch costs (albeit reduced in magnitude compared with the no-inhibition model) and n-2 repetition costs, for both RTs and error rates.

With respect to the error rate data, although the pattern of results is similar to the response time data, larger variability in the data leads to few statistically significant differences. In particular, the 95% confidence interval for the n-2 repetition cost

<sup>5</sup> A more generous 'timeout' threshold of 500 cycles was used in preference to the 400 cycles used by Gilbert and Shallice (2002), however this change had negligible effect on model behaviour and is not considered further.



**Fig. 2.** Simulation 1 results: RTs and error rates for sequences of task switches, and derived costs. Error bars represent bootstrapped 95% confidence intervals.

overlaps zero, suggesting the model does not predict statistically significant effect for error rate data. Although the size of the confidence intervals is somewhat arbitrary, given the nature of simulation (i.e., a smaller CI could be obtained simply by running more blocks, or increasing the size of each block), the point remains that more statistical power is available for the response time data, paralleling the empirical literature.

It is also notable that one effect of backward inhibition is a general increase in the error rate, irrespective of triplet type. On the face of it, this may raise concerns about the function of a backward inhibition mechanism: what benefits would such an error-causing mechanism bring to a cognitive system? This question is deferred until the general discussion, however it is worth noting at this point that the model's error rate, even with backward inhibition, is generally below 1% – substantially below empirically observed error rates. Thus, while the model accounts well for the empirically observed pattern of reaction times, as a general model of speeded response generation/execution it remains somewhat incomplete in comparison with other models which account for a speed/accuracy tradeoff (e.g., [Brown et al., 2007](#)).

#### 4.2. Simulation 2: Generality of the effects

It has been argued that fitting a model to an empirical data pattern is not informative if the model could fit any arbitrary pattern of behaviour ([Roberts & Pashler, 2000](#)). Simulation 2 therefore pursues a strategy similar to that advocated by [Pitt, Kim, Navarro, and Myung \(2006\)](#), whereby a large portion of the three-dimensional parameter space defined by *gain*, *bias*, and *weight* was systematically scanned and partitioned into regions that yielded qualitatively different behaviours.<sup>6</sup> Thus we sought to understand where in parameter space the behaviour of a model would be qualitatively similar to human participants (i.e., RT costs for n-2 repetitions and n-1 switches) and where it would produce alternative patterns (e.g., n-1 repetition costs and n-2 switch costs, or costs for both n-1 and n-2 switches, or costs for both n-1 and n-2 repetitions). More specifically, in simulation 2 the three parameters were independently varied across their ranges, and for each resulting point within parameter space, a model with those values was tested on the paradigm. The resulting dependent variables were generated and plots produced to show effect sizes for each effect at each point in parameter space.

##### 4.2.1. Method

Simulation 2 varied the three parameters of the conflict system: *gain* ( $\gamma$ , 0–100); *bias* ( $\beta_c$ , –40 to 0); and *weight* ( $\omega$ , –30 to 0). As described above, *gain* and *bias* affect the rate at which conflict unit activation builds up, and decays, respectively. *Weight* affects the amount of biasing that conflict units exert on Task Demand (TD) units. A weight of zero is therefore functionally equivalent to a model with no conflict mechanism (and thus, only lateral inhibition between TD units). The effect that the task inhibition/conflict mechanism has on behaviour is assessed by comparing stronger levels of *weight* with this baseline. For simplicity, simulations were run on blocks of three trials, with the dependent measure being RT in model cycles

<sup>6</sup> Note that for simplicity, the *conflict squashing* parameter was not varied in these simulations. However, additional simulations revealed that the effect of varying this parameter was to change the magnitude of the resulting effect sizes but not the qualitative pattern of results.

for the final trial in each triplet. The model was re-initialised after each block. Thus, each block was defined by the sequence of task switches as 0SW, 1SW, 2SW or ALT. Also for simplicity, given the stochastic nature of the error rate data, and the finding from simulation 1 that error rates largely parallel RTs, only RT data was considered. Mean RT switch costs and n-2 repetition costs, in model cycles, were calculated from 3000 three-trial blocks of each switch condition and for each point within parameter space.

#### 4.2.2. Results

Dependent variables were compared for each point using a Welch two-samples t-test, with the resulting  $t$  converted to an effect size ( $r$ ) and plotted in Fig. 3 (upper panel: n-1 switch cost; centre panel: n-2 repetition cost). Dark grey voxels in the upper two panels represent no data: with this combination of parameter settings, the model is unable to correctly complete any trials in at least one condition. The intersection of both empirical effects (n-1 switch costs and n-2 repetition costs), calculated as the geometric mean of both effect sizes, for points in parameter space with both effects in the empirically observed direction, is shown in Fig. 3 (lower panel).

#### 4.2.3. Discussion

As shown in Fig. 3 (upper panel), n-1 switch costs are robustly predicted over a wide region of the model's parameter space, except for a small region in the upper right of the plot for stronger *weight* values. Here, a high *gain* and weak *bias* means that activation of conflict units increases irrespective of the degree of actual conflict. Where *bias* is close to zero, the activation of the conflict units quickly saturates, particularly for strong *gain* values. Where *weight* values are also high, the model ceases to function correctly (i.e., it cannot correctly complete any trials). Conversely, behaviour in the bottom left of each plot (i.e., strong *bias* and low *gain*) is relatively uniform. With these settings conflict unit activation decreases irrespective of input, and thus little biasing of model processing occurs.

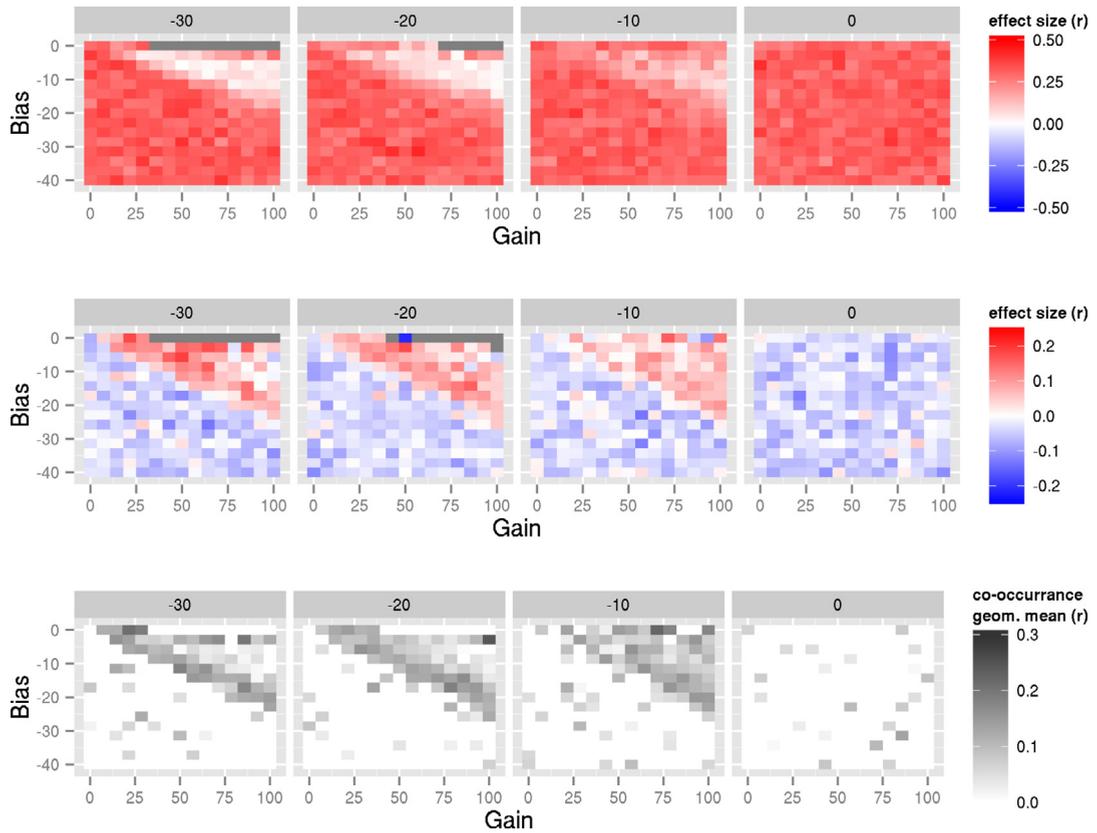
N-2 repetition costs (Fig. 3, centre panel) are less robust than switch costs, partly because the difference in sequence types (occurring on the n-2 trial) must affect processing after an intervening trial (i.e., the n-1 trial). Nevertheless, systematic effects occur, and for stronger *weight* values inhibition of TD units by the conflict units produces the empirically observed effect for a contiguous region of parameter space. Note also that, replicating the results of simulation 1, the model does not produce n-2 repetition costs for near-zero *weight* values, demonstrating that lateral inhibition of task representations alone is insufficient to produce n-2 repetition costs.

In addition to the above, a large region of parameter space (below a top left - bottom right diagonal) results in n-2 repetition *facilitation*. Here, the combination of strong *bias* and low *gain* means that conflict decays too quickly, with the units insufficiently sensitive to their inputs to produce residual conflict effects. Only in the region above this diagonal, where *bias* is weaker and/or *gain* stronger, are n-2 repetition costs consistently produced for non-zero *weight* values.

The two empirically observed phenomena, costs for n-1 switches and n-2 repetitions (Fig. 3 lower panel) co-occur in a well-defined region for non-zero *weight* values. Informally, this behavioural pattern is obtained with the constraints that the activation of conflict units must increase given conflict and decrease given lack of conflict. Outside these regions, other behaviour (e.g., switch costs but n-2 repetition facilitation) may be understood either in terms of implementational failure of the model (the parameter settings do not implement the theory of backward inhibition) or in terms consistent with theory (TD unit processing must be biased by inhibitory connections from conflict units). Therefore, the empirical pattern is produced by any set of parameters in which the model functions according to these theoretical constraints.

#### 4.3. Simulation 3: Variable intertrial intervals

One of the key empirical findings suggesting that n-2 repetition costs are an inhibitory phenomenon, rather than merely reflecting interference from a residually active irrelevant task, comes from studies which manipulated the intertrial interval (RCI, response-cue interval). When RCIs are manipulated blockwise, n-2 repetition costs tend to decrease with longer intervals (Koch, Gade, & Philipp, 2004; Meiran, Chorev, & Sapir, 2000). This finding would be consistent with n-2 repetition costs being caused either by residual activation (i.e., of the B task in an ABA sequence), or residual inhibition (of the A task in the same sequence), which dissipates over time and is thus greater for shorter RCIs. Gade and Koch (2005), however, hypothesized that if task inhibition is primarily affected by decay of the n-1 task, n-2 repetition costs should mainly be affected by manipulation of the RCI preceding trial  $n$  ( $RCI_n$ ) – that is, task inhibition should decrease with longer intervals before the return to the inhibited task. Alternatively, if task inhibition were mainly affected by conflict between co-active tasks, n-2 repetition costs should mainly be affected by the RCI preceding trial  $n-1$  ( $RCI_{n-1}$ ) – that is, task inhibition should decrease with longer intervals immediately before the trial that causes the 'A' task to be inhibited. In their experiments, when  $RCI_{n-1}$  and  $RCI_n$  could each be either short (i.e., 0.1 s) or long (1.4 s or 1.9 s in experiments 1 and 2 respectively), n-2 repetition costs were affected by  $RCI_{n-1}$  but not  $RCI_n$ . This finding has been taken to support the view that n-2 repetition costs are an instance of cognitive inhibition, and linked to intertrial conflict generated on the n-1 trial (Koch et al., 2010). As the present model represents an implementation of the inter-task conflict hypothesis, it should readily produce this same empirical pattern. This was tested in simulation 3.



**Fig. 3.** Simulation 2 results. Maps within each horizontal panel correspond to increasing levels of the *weight* parameter, from  $-30$  (far left) to zero (i.e., baseline, far right). Upper panel:  $n-1$  switch cost effect size ( $r$ ). Central panel:  $n-2$  repetition cost effect size ( $r$ ). Lower panel: geometric mean of  $n-1$  switch cost and  $n-2$  repetition cost effect sizes for voxels with positive costs only. Dark grey voxels in the upper two figures represent no data: the model does not correctly complete trials in these regions.

#### 4.3.1. Method

In the model, the concepts of residual activation and inhibition are modelled by two features. As in [Gilbert and Shallice \(2002\)](#), residual task activation is simulated by a proportion of task demand unit activation (determined by the task demand squashing parameter,  $\phi$ , valid between 0 and 1), which is carried forward to subsequent trials. Inhibition is caused by the residual activation of conflict units, modelled by a proportion of conflict unit activation (conflict squashing parameter,  $\chi$ ) carried forward to subsequent trials. Thus, residual activation in the model,  $\alpha_r$ , resulting from the previous trial, is calculated as:

$$\alpha_r = \begin{cases} \alpha_{td} \times \phi & \text{for task demand units} \\ \alpha_c \times \chi & \text{for conflict units} \end{cases} \quad (4)$$

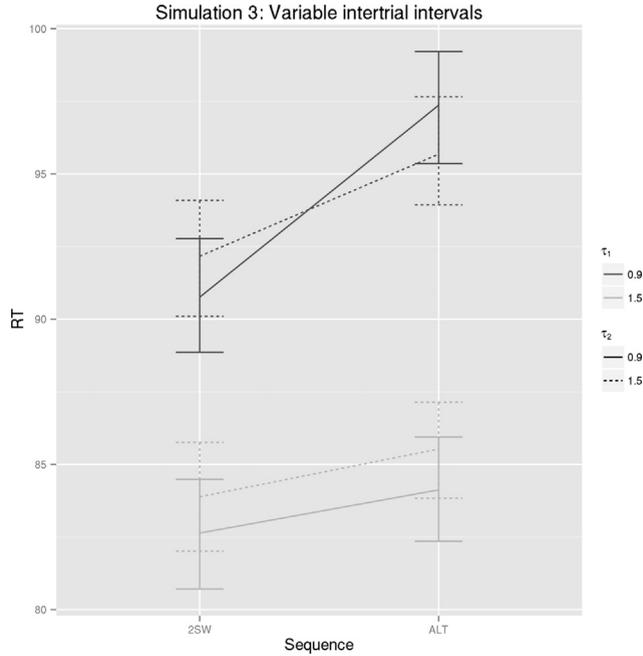
where  $\alpha_{td}$  is the task demand unit activation, and  $\alpha_c$  is the final conflict unit activation at the end of the preceding trial.

These calculations assume that RCI is constant and do not consider how the decay of activation over time will be affected by RCI. This simulation assumes that Eq. (4) calculates decay after unit time. More generally, decay after time  $\tau$  may be calculated as:

$$\alpha_r = \begin{cases} \alpha_{td} \times \phi^\tau & \text{for task demand units} \\ \alpha_c \times \chi^\tau & \text{for conflict units} \end{cases} \quad (5)$$

Thus, the model parameters which affect residual activation and its decay are the task demand squashing parameter  $\phi$ , the conflict squashing parameter  $\chi$ , and  $\tau_{n-1}$  and  $\tau_n$ , representing the RCI preceding trial  $n-1$  and trial  $n$ , respectively.

Simulation 3 varied the  $\tau$  parameters while maintaining  $\phi$  and  $\chi$  at their default values (0.8 and 0.5 respectively). Values of  $\tau$  of 0.9 and 1.5 were used for the short and long RCIs. The simulation was run on the four different combinations of  $\tau_{n-1}$  and  $\tau_n$  for 2000 3-trial blocks (1000 of each ABA and CBA sequences).



**Fig. 4.** Simulation 3: RTs for 2SW and ALT conditions. Black/grey lines represent short/long values of  $\tau_{n-1}$ , while solid/dashed lines represent short/long values of  $\tau_n$ . Error bars represent bootstrapped 95% confidence intervals.

4.3.2. Results

The RTs for trial 3 of 2SW and ALT sequences are plotted in Fig. 4. As can be observed from the figure, for a given set of parameters, ALT trials have longer RTs than 2SW trials. Hence all parameterisations produce an n-2 repetition cost. The  $\tau_{n-1}$  parameter is represented by line colour: short/long values represented by black/grey lines respectively. It is evident that shorter values of  $\tau_1$  produce a greater n-2 repetition cost, in addition to producing longer RTs overall. The  $\tau_n$  parameter is represented by line style, with solid/dashed lines representing short/long values. Different values of this parameter have a relatively small effect on both overall RTs and n-2 repetition costs.

This graphical analysis was assessed quantitatively using a 3-way ANOVA ( $sequence \times \tau_{n-1} \times \tau_n$ ). The main effect of sequence (n-2 repetition costs) was significant,  $F(1, 7301) = 23.09, p < .0001$ , as was the main effect of  $\tau_{n-1}$ ,  $F(1, 7301) = 213.88, p < .0001$ , however the main effect of  $\tau_n$  was not,  $F(1, 7301) = 0.83, p = .36$ . The effect of interest, that of RCI on n-2 repetition costs, was assessed by the two  $\tau \times sequence$  interactions. First the  $\tau_n \times sequence$  interaction was not significant,  $F(1, 7301) = 1.13, p = .29$ . However, the  $\tau_{n-1} \times sequence$  interaction was highly significant,  $F(1, 7301) = 6.71, p = .009$ . Finally, the  $\tau_{n-1} \times \tau_n$  interaction,  $F(1, 7301) = 1.14, p = .29$ , was not significant, and neither was the three-way  $\tau_{n-1} \times \tau_n \times sequence$  interaction,  $F(1, 7301) = 1.43, p = .23$ . Thus, for these two values of  $\tau$ , the model confirms the hypothesis and reproduces the empirical finding, that n-2 repetition costs are affected by the RCI preceding the n-1 ( $\tau_{n-1}$ ), but not the n ( $\tau_n$ ), trial.

To determine the respective effects of  $\tau_{n-1}$  and  $\tau_n$  on n-2 repetition costs more generally, a further simulation was run in which the two were varied factorially, from 0.3 to 2.4 in increments of 0.3. For each of these 64 combinations, 1000 blocks of three trials were run for both ALT and 2SW sequences. The mean RTs are represented as 3D plots in Fig. 5a and b, and the resulting n-2 repetition costs are plotted in Fig. 5c. Graphically, it is evident that the n-2 repetition cost decreases with longer values of  $\tau_{n-1}$ , while it is not systematically affected by changes in  $\tau_n$ . To quantify this model behaviour, linear regression was used to estimate the effect of  $\tau_{n-1}$  and  $\tau_n$  on trial 3 RT, for both 2SW and ALT sequences. Estimates for  $\beta$  coefficients for  $\tau_{n-1}$  and  $\tau_n$  (and their bootstrapped 95% confidence intervals) are given in Table 1. Note that the confidence intervals for  $\beta(\tau_{n-1})$  do not overlap, while both confidence intervals for  $\beta(\tau_n)$  include the other estimated  $\tau_n$  coefficient, and indeed are nearly identical. The estimated coefficients yield the following regression equations, which may be combined to calculate n-2 repetition costs:

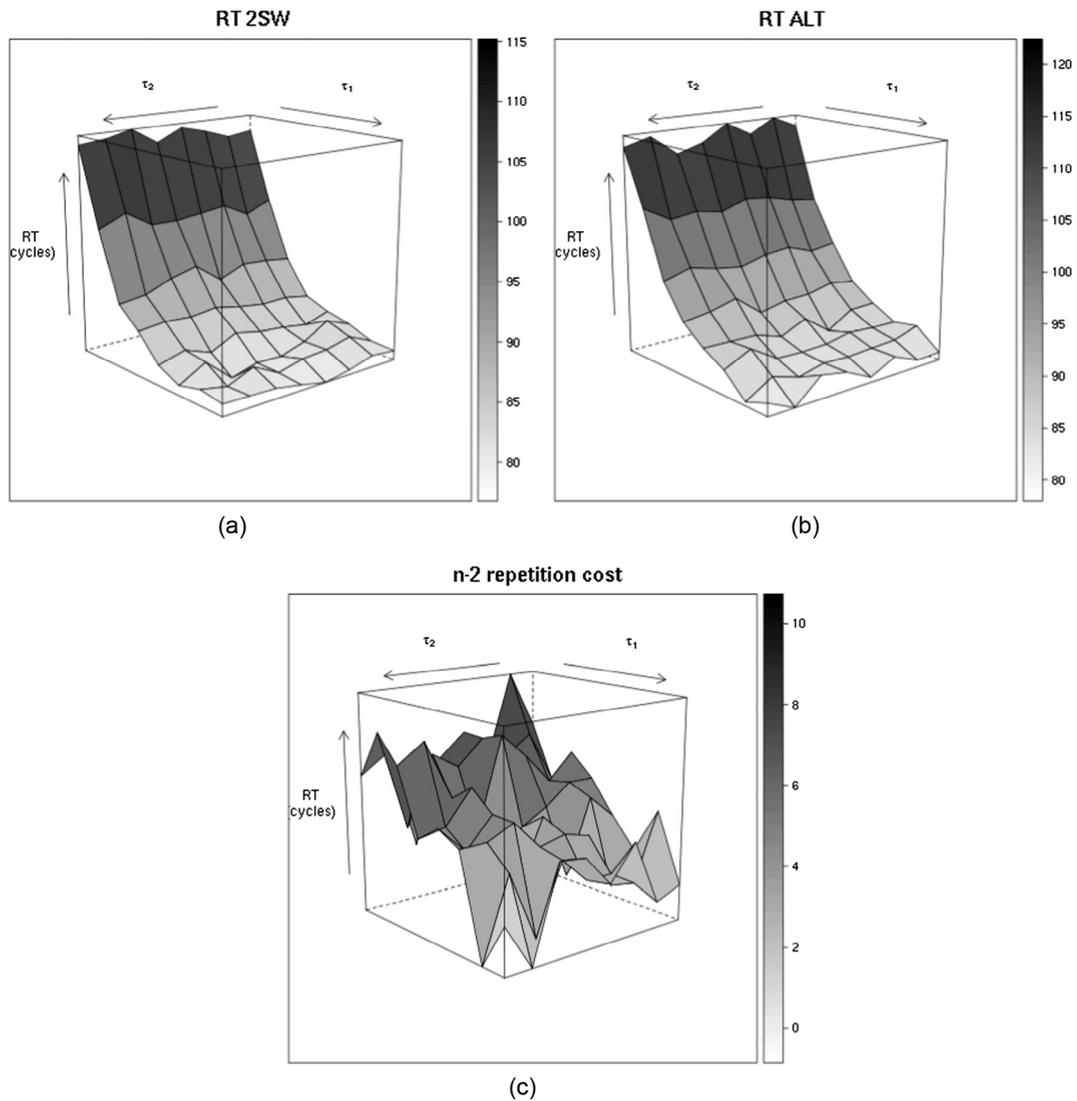
$$2SW RT = 105.22 - 12.45\tau_{n-1} + 0.14\tau_n$$

$$ALT RT = 112.56 - 14.57\tau_{n-1} + 0.14\tau_n$$

...

$$n-2 \text{ repetition cost} = 7.34 - 2.12\tau_{n-1}$$

Importantly, the n-2 repetition cost is dependent on  $\tau_{n-1}$ , while the  $\tau_n$  terms cancel out.



**Fig. 5.** Effects on simulated RT (expressed in model cycles) of varying RSI preceding  $n-1$  and  $n$  trials ( $\tau_{n-1}$ ,  $\tau_n$ ) factorially. (a) 2SW RTs. (b) ALT RTs. (c)  $n-2$  repetition costs.

**Table 1**

Regression coefficients [and 95% confidence intervals] for RT for 2SW and ALT blocks.

Sequence	Intercept	$\beta(\tau_{n-1})$	$\beta(\tau_n)$
2SW	105.32 [104.59–106.04]	–12.45 [–12.81 to –12.10]	0.14 [–0.20 to +0.49]
ALT	112.56 [111.81–113.31]	–14.57 [–14.93 to –14.20]	0.14 [–0.22 to +0.49]

#### 4.3.3. Discussion

The model's predictions clearly parallel the behavioural pattern observed by [Gade and Koch \(2005\)](#). Most critically, longer values of  $\tau_{n-1}$  reduce RTs of both 2SW and ALT trials, however the former are reduced by less than the latter. This has the effect of decreasing  $n-2$  repetition costs.  $\tau_n$  has no reliable effect on RTs. The general increase in RTs for both 2SW and ALT trials for shorter values of  $\tau_{n-1}$ , but not  $\tau_n$ , is also observed by [Gade and Koch \(2005\)](#), and is consistent with blockwise RCI manipulations (e.g., [Grange & Houghton, 2009](#); [Koch et al., 2004](#)).

This simulation demonstrates that a characteristic empirical pattern is predicted by the model with no changes to its architecture, beyond the assumption that residual activation decays over time (and hence decays more with longer intertrial intervals). It also suggests that the claim that backward inhibition is a residual effect caused by conflict occurring when switching away from a task is a viable theoretical proposal, and that this does indeed predict the observed empirical pattern.

In short, it validates the verbally specified model of Koch et al. (2010), demonstrating that the theoretical proposals are indeed consistent with the data.

#### 4.4. Simulation 4: Asymmetric tasks

Although simulation 1 demonstrates that a conflict mechanism is sufficient for explaining n-2 repetition costs in human participants, it does not provide any insight into whether this occurs as a result of slowing or facilitating particular trials. Additionally, it does not address why a cognitive system for task switching should employ a conflict detection mechanism in order to regulate performance. These questions may be addressed by considering switching between tasks of asymmetric difficulty. The model of Gilbert and Shallice (2002) demonstrated that asymmetric n-1 switch costs occur because a greater amount of task demand activation is required to execute a harder task than an easier task, and therefore a greater amount of this activation is carried over to the next trial. This in turn causes greater cross-task interference (i.e., conflict). Consequently switch costs are greater for hard-easy switches than for easy-hard switches. Simulation 4 extends the asymmetric task paradigm to three task switching, manipulating the between-trial conflict by using two tasks of identical, fixed, intermediate difficulty, while varying the difficulty of the third task.

##### 4.4.1. Method

In the model of Gilbert and Shallice (2002), task difficulty is specified by two parameters – *stimulus input strength (SIS)*, representing the automatic, bottom-up activation of a response by a stimulus (greater for stronger tasks), and *top-down control strength (TDCS)*, specifying the control needed to ensure the task is performed (greater for weaker tasks). Simulation 4 varied the *SIS* and *TDCS* parameters of task B together with the *weight* parameter of the conflict layer to create a three-dimensional space. The *bias* and *gain* parameters of the conflict layer were fixed at  $-10.0$  and  $75.0$  respectively, while the task parameters (*TDCS*, *SIS*) of tasks A and C were held at their default values.

For asymmetric tasks, each task sequence (e.g., OSW) has various permutations – (e.g., ABB, BAA and BCC). Here, we considered only switches from task B (variable *SIS* and *TDCS*) to task A (fixed). Hence, OSW sequences are all BAA, 1SW are BBA, 2SW are CBA and ALT are ABA, with only the n-1 task being of variable difficulty (except for the OSW condition). By varying the parameters of task B, simulation 4 therefore tests the effect on behaviour for both hard-easy-hard (HEH) and easy-hard-easy (EHE) switches.

As top-down control provides a constant positive input to the cued task demand unit, a variable *TDCS* is a confounding factor in that the same degree of task inhibition has a stronger influence on the processing of units with a low *TDCS* (i.e., easier tasks) than those with a high *TDCS* (i.e., harder tasks). Thus, rather than using a single weight value for all connections from inhibitory conflict units to TD units (as in simulation 1), in this simulation the influence of conflict units on TD units was calculated as the product of the *weight* parameter and the *TDCS* parameter for each TD unit.<sup>7</sup> The input to task demand units ( $I_{td}$ ) may be expressed as the following modified equation:

$$I_{td} = S_{tdc} + \sum_c \alpha_c \omega_c S_{tdc} + \sum_o \alpha_o \omega_o \quad (6)$$

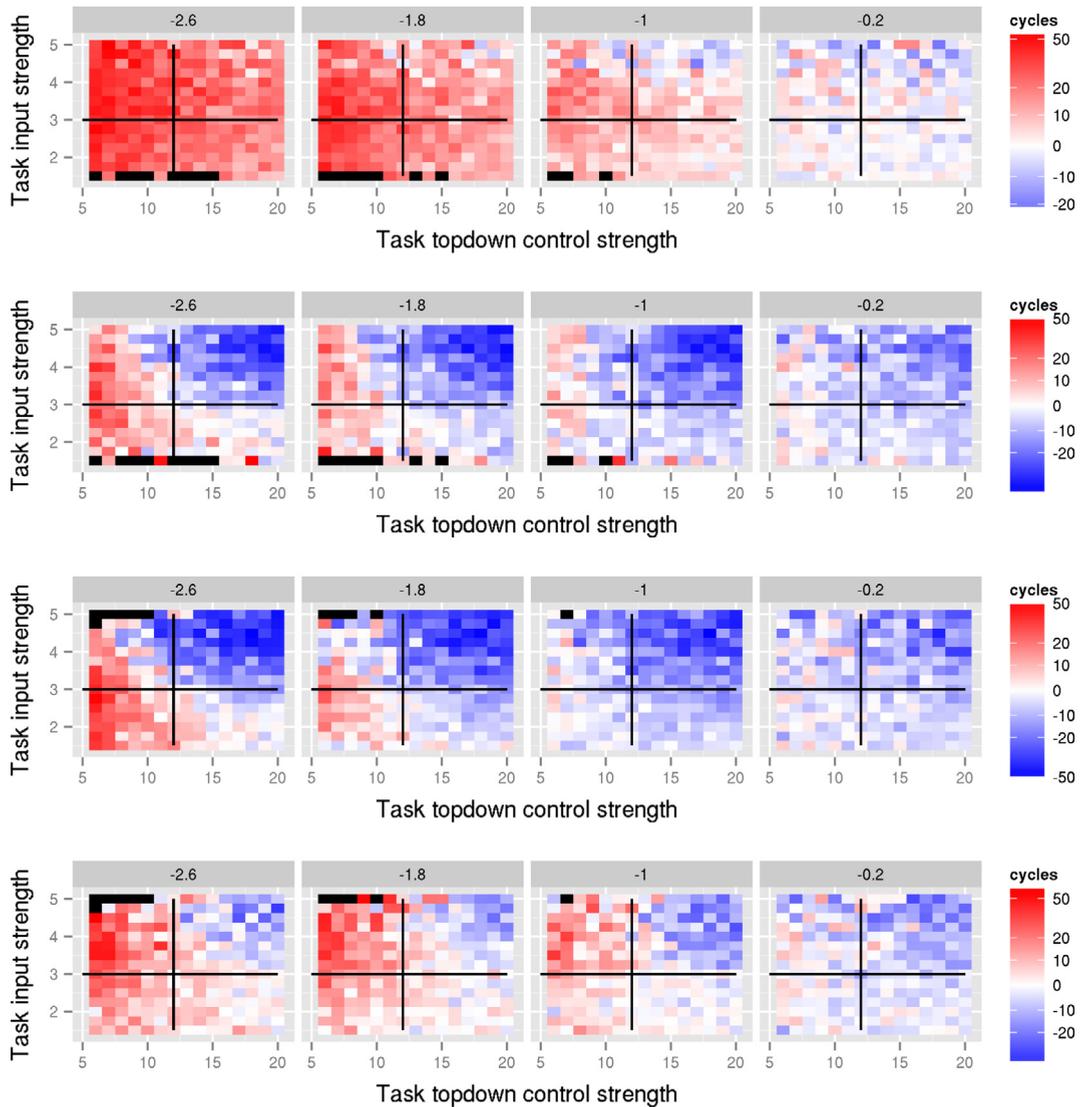
where  $S_{tdc}$  is the top down control strength parameter for the respective unit,  $\alpha_c$  and  $\omega_c$  are the activations and weights for the two conflict units connected to the unit, and  $\alpha_o$  and  $\omega_o$  are the activations and connection weights of the six output units in the model.

By varying *SIS* and *TDCS* of task B factorially, such that either may be greater or less than that for task A, the resulting two-dimensional parameter space is divided into four quadrants. The upper-left represents the region in which task B is stronger, but less controlled, than task A, as in a stronger task B (e.g., word reading). In the lower-right, B is weaker, but more controlled than task A, indicating a weaker task B (e.g., colour naming). In the upper-right both the input and control strength are greater for task B, hence the task has more control than is needed to perform the task. Finally, in the lower-left a weak task is coupled with insufficient control.

##### 4.4.2. Results and discussion

To illustrate the effect of the conflict/task-inhibition mechanism on performance, Fig. 6 plots the basic RTs (panels 1–4) for the trial 3 of a 3-trial block, relative to a baseline where the *weight* of conflict-task demand connections is zero, while Fig. 7 plots the resultant modulation of n-1 switch costs (upper panel), n-2 repetition costs (lower panel). The plot of relative switch costs (Fig. 7, upper panel) shows that stronger *weight* values produce smaller switch costs, especially for HEH switches (upper quadrants). The effect is qualitatively modulated by the strength of inhibitory biasing: for the weakest *weight*, the effect is minimal, while for stronger *weight* values, the reduction in switch costs is greatest where the n-1 task receives a high level of activation (top right quadrant) due to both a high level of control and high input strength. This is because greater task-demand activation on the n-1 task leads to more residual conflict on the n (switch) trial, and the resulting task inhibition reduces interference, leading to a reduction in relative switch cost. Interestingly, this selective reduction

<sup>7</sup> It was determined through simulation that this modified architecture produces the same pattern of results as observed in simulation 1. However, if we consider that in simulation 1 *TDCS* is identical for all three tasks, we can see that the effect is merely to multiply the weight value of all three conflict-task demand connections by a common constant, effectively rescaling the *weight* axis of parameter space but having no effect on the pattern of results.

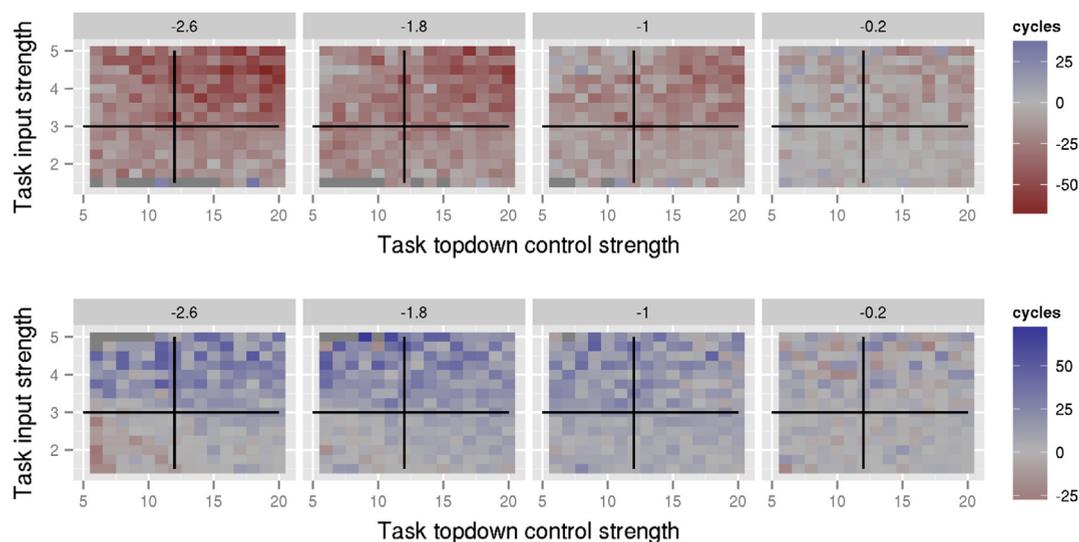


**Fig. 6.** Simulation 4: Modulation of RT for individual OSW, 1SW, 2SW and ALT conditions (from top to bottom), for successive values of *weight*, from strong (left) to weak (right). All values are relative to a zero weight baseline. Black squares represent no data, and the model's failure to produce correct responses with these parameter settings.

in switch costs exaggerates the switch cost asymmetry (i.e., it reduces costs more for EH than HE switches) suggesting that in a task-switching system with such a mechanism, a component of the switch cost asymmetry may be attributable to task inhibition.

To understand the basis for this asymmetry, we next consider the effect of the backward inhibition mechanism on 1SW and OSW sequences individually. The simplest case is the 1SW (e.g., BBA) condition (Fig. 6, panel 2), where the switch occurs on trial  $n$ . Here, lower *weight* levels produce RT facilitation, particularly for switches from highly controlled tasks (e.g., HE switches, right hand quadrants). On the switch trial, more conflict is generated in activating the A task demand unit when B is harder than when it is easier, due to greater residual task-demand activation when performing the harder task on trial  $n-1$ . This is the reason for the asymmetric  $n-1$  switch cost (i.e., larger costs for HE than EH switches) in the model of Gilbert and Shallice (2002). Due to this asymmetry in conflict, processing on trial  $n$  is facilitated when task B is harder than task A. When B is easier than A, more activation of task demand unit A is required to produce a response. However, because task demand unit B is still residually active, task inhibition counteracts the activation by top-down control of task demand unit A, leading to interference on trial  $n$ , and slower responses.

In the task-repeat OSW condition (e.g., BAA, Fig. 6, panel 1), stronger *weight* values produce longer RTs on the final trial, especially following switches from easier tasks (left quadrants). As for the switch on trial  $n$  in the 1SW sequence, conflict occurs between the A and B task demand units on the switch trial, excepting that for OSW sequences this occurs on the



**Fig. 7.** Simulation 4: Modulation of  $n-1$  switch costs (upper panel) and  $n-2$  repetition costs (lower panel) for successive values of *weight*, from strong (left) to weak (right). All values are relative to a zero *weight* baseline.

$n-1$  trial. Due to persistent conflict, the A/B conflict unit is still residually active on trial  $n$ , inhibiting the A task. Thus, slowing in OSW trials is due to residual conflict from previously occurring task switches. At stronger *weight* values, there is greater slowing for EH (top left) than for HE (bottom right) switches. Because OSW is the baseline condition for establishing the  $n-1$  switch cost, slower responses lead to a reduction in the EH switch cost.

Taken together, the reduction in switch cost, which is greater for EH switches, occurs for two reasons: firstly, 1SW trial facilitation, particularly for switches from more controlled tasks (including HE switches); and secondly, OSW trial interference, particularly for switches from less controlled tasks (including EH switches). Overall, the switch cost is reduced for both HE and EH switches, but the effect is greater for EH switches, exaggerating the switch cost asymmetry.

In contrast to the effects on  $n-1$  switch costs, stronger *weight* values produce larger  $n-2$  repetition costs (Fig. 6, lower panel). This effect is modulated by input control strength, and is larger for HEH alternations (i.e., upper quadrants) than EHE alternations. In the ALT condition (Fig. 6 panel 4), the effect is modulated by  $n-1$  task difficulty. For low *weight* values, conflict units generally facilitate performance. For intermediate values they selectively impair HEH alternations (top left quadrant), while for the strongest values they also begin to slow EHE alternations.

The reason for these asymmetric effects on ALT trials is because more conflict is generated on the  $n-1$  trial for HEH sequences than EHE sequences, for the same reasons that HE switches generate more conflict than EH switches in 1SW trials above. In the current model, greater conflict on the  $n-1$  trial means that the  $n-2$  task is inhibited more on HEH than EHE switches, and therefore re-activating this task on trial  $n$  of ALT sequences takes more time.

In the 2SW condition (Fig. 6 panel 3), mild and intermediate *weight* values produce generalised facilitation. For stronger *weight* values, the effect on RTs is highly modulated by top-down control, with interference and facilitation caused by low-control and high-control  $n-1$  trials, respectively, and with most pronounced effects for under- and over-activated trials (i.e., bottom left and top right quadrants). The fact that facilitation dominates irrespective of *SIS* or *TDCS* for all but the strongest *weight* values suggests that the mechanism beneficially reduces the amount of control required to achieve good performance when frequently switching between multiple tasks. In contrast to the effect on ALT trials, on a 2SW trial the difficulty of the  $n-1$  task has a much less pervasive effect on trial  $n$  performance, because the task demand unit for trial  $n$  has not recently been inhibited by conflict units.

Together, the 2SW and ALT results explain the larger  $n-2$  repetition costs observed in HEH alternations than in EHE alternations – it is a composite of stronger facilitation for switches from easier/more controlled tasks in the 2SW condition, and greater interference when switching from those same tasks in the ALT condition.

In general, the effects of backward inhibition on all four types of trial can be understood in terms of the conflict generated on switch trials by simultaneously active task demand units, and the asymmetry in activation dynamics between easy-hard and hard-easy switches. Because the model regulates task inhibition on the basis of detected conflict, the model predicts an asymmetric effect on the size of the  $n-2$  repetition cost. More specifically it predicts that the cost is larger for HEH than EHE alternations, and that this is a residual effect from the additional conflict generated by HE switches compared with EH switches occurring on the  $n-1$  trial. This prediction was directly tested in the following experiments.

## 5. Experiment 1

This experiment tested the prediction of simulation 4 – that n-2 repetition costs are modulated by asymmetric task difficulties: for ABA triplets (i.e., the n-2 repetition condition) the n-2 repetition cost is larger for hard-easy-hard (HEH) triplets than for easy-hard-easy (EHE) triplets.

To date, only one study has systematically examined n-2 repetition costs for asymmetric tasks. [Arbuthnott \(2008\)](#) used a digit classification variant, in which the three tasks were judgments about magnitude (low/high), parity (odd/even) and whether the digit was prime (prime/multiple), in increasing order of difficulty, assessed through RT and error rates. In her procedure, the order of tasks was fully randomised for each participant, and each trial was classified by the sequence of switches in relation to previous tasks – possibilities were OSW (e.g., BAA), 1SW (e.g., BBA), 2SW (e.g., CBA) or ALT (e.g., ABA). Thus, a single procedure was capable of obtaining values for (n-1) switch costs (the RT difference between 1SW and OSW trials) and n-2 repetition costs (the RT difference between ALT and 2SW trials). Arbuthnott reasoned that if the asymmetric n-1 switch cost was due to residual inhibition on the switch trial (that is, greater for hard-easy (HE) switches than easy-hard (EH) switches, because executing the H task requires greater inhibition of the E task than vice versa, and this carries over to subsequent trials), and if this same mechanism was responsible for backward inhibition, then it would lead to greater n-2 repetition costs for EHE than HEH alternations. Alternatively, if the asymmetric switch cost was due to residual priming of the harder task (i.e., greater priming for HE than EH switches), then this would have no differential effect on n-2 repetition costs for EHE compared to HEH alternations.<sup>8</sup>

In two experiments, involving either separate or overlapping response sets respectively, [Arbuthnott \(2008\)](#) observed asymmetric n-2 repetition costs, with greater costs for EHE triplets than HEH triplets. That is, the n-2 task received greater backward inhibition when it was easy than when it was hard. This was consistent with Arbuthnott's hypothesis that backward inhibition was caused by the same mechanism as the n-1 switch cost, namely residual inhibition of the preceding task. However, the effect on RTs was not robust, and only reached statistical significance for one pairing of tasks (magnitude-prime switches, not parity-prime switches), and then for non-overlapping response sets only. Additionally, unexpected effects occurred, such as the reversal of direction of the n-1 switch cost asymmetry for magnitude-prime tasks; that is, switching to the prime (harder) task incurred a greater n-1 switch cost than switching to the magnitude (easier) task, a finding inconsistent with the literature on two-task switch costs. It is difficult to see why, if asymmetric backward inhibition is the same mechanism as responsible for asymmetric n-1 switch costs, the direction of the latter effect should be reversed when the former effect is in the hypothesised direction. Moreover, this pattern of findings is difficult to fully explain using only 'inhibition' as an explanatory mechanism, without recognising the potential dynamic interaction of excitatory and inhibitory processes.

### 5.1. Methodological concerns

In addition to the difficulties in interpreting results, there are three potential concerns regarding the procedure used by [Arbuthnott \(2008\)](#).

Firstly, and most simply, at least one study suggests that procedures which include immediate task repeats may reduce or eliminate n-2 repetition costs ([Philipp & Koch, 2006](#)). In investigating modulations of the n-2 repetition cost, therefore, it seems prudent to design task sequences to avoid immediate task repetitions, at least in one condition.

The second issue relates to the possibility of graded response congruence effects in the digit classification task, in which single digits are classified according to three different criteria: In this case, parity (odd or even), magnitude (greater or less than 5), or prime (prime or non-prime), with two possible responses signalled by pressing left or right response keys. [Arbuthnott \(2008\)](#) used the following mapping of stimuli to responses: for magnitude, 2, 3, 4/6, 7, 9; for parity, 3, 7, 9/2, 4, 6; for prime, 2, 3, 7/4, 6, 9. Given six stimuli three possible dimensions, the stimuli have different levels of congruence between the correct response, and the responses to the alternative tasks afforded by the stimulus. For example, for magnitude judgements, the stimulus 6 has a correct response (right) congruent with both the odd-even task's response and the prime-nonprime task's response (i.e. fully congruent), 7 is congruent on only one of the irrelevant tasks (semi-congruent), while 4 is incongruent with both irrelevant tasks (fully incongruent). Thus, specific stimulus items are conflated with the degree of response conflict generated on each task. In fact, there is recent empirical evidence that this type of graded response congruence indeed affects reaction times ([Longman, Lavric, Munteanu, & Monsell, 2014](#); [Schneider, 2014](#)) in three-task versions of the digit classification task. While differing levels of stimulus response congruence might be assumed to counterbalance for each task for symmetric difficulties, when difficulty varies by task, stimulus response congruence also varies per task, thus confounding stimulus congruence with task difficulty.

The final issue is the assumption that the task sets involved do not overlap. It is questionable, however, whether prime number judgement is really independent from the other two tasks. For example, if a digit is even it is guaranteed to be non-prime unless it is 2, providing a heuristic method for determining prime/nonprime status using an easier odd/even judgement. If responding to a prime trial involved even partial activation of the parity task-set, then any task-set inhibition

<sup>8</sup> Note that this reasoning assumes that asymmetric n-1 switch costs are exclusively due to either inhibition or activation (priming) of task sets. In models of task switching incorporating inhibitory connections ([Brown et al., 2007](#); [Gilbert & Shallice, 2002](#)), activation and inhibition are dynamically interactive, and thus not easily separated.

occurring on the next trial may be applied to both prime and parity task-sets. In sum, in three-task switching paradigms that incorporate only two possible responses to each task, pairings of task difficulty may be confounded with response congruence and response set overlap from the switched-away-from task. Avoiding these problems within a three-task switching paradigm requires tasks with (at least) three possible responses.

## 5.2. Method

Participants classified letters according to the dimensions of colour (red, green or blue), alphabetical position (beginning, middle, or end) and font (bold, regular, italic). Importantly, each task has three possible responses to each stimulus, rather than two. This has the advantages that, firstly, response repetition effects can be eliminated (even if response repetitions never occur, there are still two possible responses to each task), and secondly, as the three stimulus dimensions (colour, alphabet, font) can be manipulated independently, irrelevant-dimension congruency effects can be controlled such that all stimulus dimensions are always mutually incongruent.

### 5.2.1. Participants

39 participants were recruited through the Birkbeck psychology participant database, of whom 22 were female. Informed consent was obtained from each participant. The median participant age was 28 years, with a range of 16–46 years. All participants spoke English fluently, although nine reported a first language using a non-Roman alphabet. Participants received their choice of either £7.50 cash or course credit.

### 5.2.2. Design

The experiment compared n-2 repetition costs for each of three possible pairings of tasks. The relative ordering of task difficulty was determined by comparing reaction times and error rates. For each task pairing, n-2 repetition costs could then be determined independently for easy-hard-easy (EHE) or hard-easy-hard (HEH) alternations. For example, n-2 repetition costs for the colour-alphabet pairing are calculated as the difference between colour-alphabet-colour and font-alphabet-colour sequences (EHE), and alphabet-colour-alphabet and font-colour-alphabet sequences (HEH). Thus, the experimental design had two factors: switch direction (EHE vs. HEH) and task pairing (colour-font vs. colour-alphabet vs. alphabet-font).

### 5.2.3. Tasks

Participants performed one of three tasks in response to a target letter stimulus. The target letter, rendered in the Latin Modern Roman font, appeared in the centre of the screen, approximately 20 mm × 30 mm in size. Simultaneously with the target stimulus, a cue presented above the target, 'colour', 'alphabet' or 'font' in lower case and approximately 7 mm in height, indicated which task to perform. The response to all three tasks was made by pressing one of the left, down, or right cursor keys on a standard PC keyboard, thus the set of possible responses was fully overlapping. The colour task was to identify the colour of the target, indicating red, green, or blue, by pressing left, down or right respectively. The font task was to indicate whether it was rendered in a bold (left), regular (down), or italic (right) font. The alphabet task was to indicate whether the letter occurred toward the beginning, (**d, e, f, g, h**; left), middle, (**k, l, m, n, o**; down) or end (**s, t, u, v, w**; right) of the alphabet.

### 5.2.4. Procedure

The procedure was administered using PsychoPy version 1.82.01 (Peirce, 2009) on a GNU/Linux PC. Stimuli were presented on a 380 mm × 300 mm LED monitor, placed approximately 750 mm from the participant. With the exception of the target letter, unless otherwise stated, all text was displayed in white against a black background.

Participants were instructed to perform the tasks 'as quickly and accurately as possible'. Instructions for each task were worded as follows: 'In the colour task, press the key corresponding to whether the letter appears in red, green, or blue'; 'In the alphabet task, press the key corresponding to where the letter appears in the alphabet: toward the beginning (d e f g h), middle (k l m n o), or end (s t u v w)'; 'In the font task, press the key corresponding to whether the font of the letter is bold, regular, or italic'. Simultaneously with the instruction, the three response mappings (red/green/blue, beginning/middle/end, bold/regular/italic) appeared in the bottom left, bottom centre, and bottom right of the screen, to correspond with the response key mappings.

After receiving instructions for each task, participants performed a practice block of 10 trials. Participants then performed 20 practice trials where the task switched, as in the main procedure. For all practice blocks, specific feedback was provided on incorrect responses (e.g., 'Oops! the correct response was: left. Press space to continue').

Following the practice blocks and a break screen, in which participants could rest as long as desired, participants performed six blocks of 103 trials, with a break screen between each block. The order of tasks was randomised, with the constraint that immediate task repeats were not allowed (hence, given a specific preceding task, the two other switch tasks may each occur with 50% likelihood). Stimulus dimensions were randomised, subject to the constraint that all three possible response keys were represented in each stimulus (e.g., a bold, green letter *t*) to avoid stimulus congruency effects.

**Table 2**  
Mean (SD) of RTs (milliseconds) and error rates for each task.

Task	RT	Error rate
Colour	1544 (391)	3.74% (2.83%)
Font	1774 (411)	5.86% (3.61%)
Alphabet	1849 (500)	7.98% (5.99%)

Error feedback throughout the experimental blocks was considered an important element of the procedure,<sup>9</sup> and was provided in the following ways. Firstly, immediate feedback was given following incorrect responses. This was in the format of an additional message which appeared on screen ('Oops!') and a reminder of the response mappings for all three tasks, appearing in the bottom left, bottom centre, and bottom right, of the screen. This message remained on screen for 4000 ms, before trials resumed. This additional pause following an incorrect response was intended to break any rhythm in responding, as well as providing a mild penalty for an incorrect response. Secondly, overall feedback was given following each block on the accuracy rates for each task. This was in the format of e.g., 'colour accuracy: 95%'. This message appeared in white if accuracy was 90% or higher, but in red if accuracy dropped below 90%. If the accuracy for any task was below 90%, the message 'Please take more care!' also appeared on screen, while if all were above 90%, the message read 'Well done!'.

### 5.3. Results

A single participant with a mean RT more than 2.5 SDs above the group mean was excluded from the analysis, for generally very slow responses (overall mean RT 4131 ms, compared to a group mean of 1782, SD 553 ms). In addition, error trials and the subsequent two trials, and trials with very slow RTs (over 20 s) were all excluded from the RT analysis. After trimming, the mean proportion of trials retained for analysis was 83.2%, this proportion did not vary substantially across cells of the experimental design, with the least trials retained for colour-alphabet switches (81.4%) and the most for font-alphabet switches (84.9%). The means and SDs for the remaining trials for each task are presented in Table 2.

The asymmetry of the three tasks was established using a one-way, within-subjects ANOVA. For RTs, the effect of task was significant,  $F(1.66, 61.57) = 23.1, p < .001$ , Greenhouse-Geisser corrected for non-spherical data. Pairwise t-tests (Holm-Bonferroni corrected for multiple comparisons) indicated significant RT differences between colour and font,  $t(37) = 5.39, p < .001$ , colour and alphabet,  $t(37) = 5.43, p < .001$ , and a marginally significant difference for font and alphabet,  $t(37) = 1.89, p = .067$ . Parallel tests for error rates revealed the same pattern: the effect of task was significant,  $F(1.73, 64.16) = 13.95, p < .001$ , Greenhouse-Geisser corrected, with significant pairwise differences between colour and font,  $t(37) = 3.38, p = .003$ , colour and alphabet,  $t(37) = 4.87, p < .001$ , and font and alphabet,  $t(37) = 2.40, p = .022$ .

However, individual differences did exist in the ordinal difficulty of the three tasks. In increasing difficulty, assessed by RTs/error rates, were *colour-font-alphabet* (21 participants for RTs/20 for error rates) *colour-alphabet-font* (8/10), *font-alphabet-colour* (2/2), *font-colour-alphabet* (2/2), *alphabet-colour-font* (4/4) and *alphabet-font-colour* (2/2). Nevertheless, given the statistically significant differences in task difficulty, it was concluded that the three tasks were indeed of asymmetric difficulty, with the colour task easier than the font task, which in turn was easier than the alphabet task. The analysis then proceeded to test the main hypothesis.

Mean and standard deviations of n-2 repetition costs (for RTs and error rates) for each task pairing and direction of alternation are given in Table 3. The analysis of n-2 repetition costs consisted of a  $3 \times 2$  within-subjects ANOVA, with factors of task pairing (colour-font vs. colour-alphabet vs. alphabet-font) and switch direction (EHE vs. HEH), to test the hypothesis that n-2 repetition costs for hard-easy-hard triplets would be greater than for easy-hard-easy triplets. For RT, there was a significant main effect of switch direction,  $F(1, 37) = 4.441, p = .042$ , partial  $\eta^2 = 0.022$ , consistent with the main hypothesis. The main effect of task pairing was not significant,  $F(2, 74) = 1.06, p = .353$ , and nor was the task  $\times$  switch direction interaction ( $F(2, 74) = 0.218, n.s.$ ). For error rates, the main effect of switch direction failed to reach significance ( $F(1, 37) = 2.815, p = .102$ ). The main effect of task pairing was also not significant ( $F(2, 74) = 0.064, n.s.$ ). However, the task  $\times$  switch direction interaction was significant,  $F(1.98, 73.19) = 3.400, p = .039$ , partial  $\eta^2 = 0.034$ , Greenhouse-Geisser corrected.

Given the statistically marginal nature of the results, the individual differences in task difficulty present one area for potential concern. For example, the larger n-2 repetition costs observed for HEH alternations might be partially driven by the 18 participants for whom task difficulty did not conform to the group mean. The results were therefore analysed again according to relative task difficulty, with the three tasks for each participant ranked as primary, secondary, and tertiary tasks

<sup>9</sup> In a pilot experiment, in which feedback on errors was provided only during the practice phase, participants appeared to regulate their performance by maintaining a relatively uniform response time across tasks, with large differences instead appearing in error rates. This may have occurred for a number of reasons. Participants may sometimes have been unaware of their errors; they may not have been motivated to try to reduce their errors for harder tasks, or the uniform rate of stimulus presentation may have encouraged participants into a rhythmic response pattern, leading to a forced response in a relatively constrained response time window, rather than responding to the dual constraints of responding quickly and accurately. In the present experiment, where reaction time was the primary dependent variable, it was expected that participants should maintain a high level of accuracy across all three types of trial, adjusting the level of control appropriate for each task, and thus taking more time on harder trials.

**Table 3**

Experiment 1: Mean (SD) n-2 repetition costs for RT and error rates by task pairings and switch direction.

Task pairing	Easy-hard-easy		Hard-easy-hard	
	RT (ms)	Error rate (%)	RT (ms)	Error rate (%)
Colour-Font	7 (202)	−0.95 (4.01)	98 (292)	1.82 (4.96)
Colour-Alphabet	16 (249)	−0.45 (4.67)	100 (266)	0.82 (4.57)
Font-Alphabet	−10 (211)	−0.96 (4.23)	31 (263)	−0.47 (5.51)

on the basis of mean RTs, in ascending order of difficulty. For RT, there was no statistically significant main effect of task,  $F(2, 76) = 1.02, p = .36$ , however the main effect of switch direction was significant,  $F(1, 38) = 4.58, p = .039, \eta_p^2 = 0.030$ . The task  $\times$  direction interaction was not significant,  $F(2, 76) = 1.09, p = .34$ . A parallel analysis of error rates revealed no statistically significant effects (main effect of task:  $F(1, 38) = 1.88, p = .16$ ; main effect of direction:  $F(1, 38) = 0.00, p = .95$ ; interaction:  $F(2, 76) = 1.10, p = .34$ ).

#### 5.4. Discussion

The results, a significant main effect of switch direction for RTs, and a nonsignificant trend for error rates, with greater n-2 repetition costs for HEH triplets than EHE triplets, contrast with those of [Arbuthnott \(2008\)](#), who reported greater inhibition of the easier task (larger n-2 repetition costs for EHE switches). They are, however, consistent with the main hypothesis and the prediction of the model.

How might one account for the differences in direction of effect between this experiment and the findings of [Arbuthnott \(2008\)](#)? As discussed above, there are a number of differences in task design. Specifically, the current procedure uses tasks with three possible responses, allowing the procedure to remove stimulus congruency and varying response set overlap as possible confounding variables. Additionally, the three tasks (font, alphabet, and colour) were chosen to avoid potential overlap in their task-sets.

One feature of the present results is that the magnitude of the n-2 repetition costs observed are rather small. However, the tested hypothesis concerned the modulation of the n-2 repetition cost, with no prediction as to the magnitude of the n-2 repetition cost overall. One potential concern is that the absence of an n-2 repetition cost might suggest the absence of task inhibition, precluding its modulation in the present experiment. While there is reason to believe n-2 repetition cost of zero does not represent zero task inhibition, merely reduced task inhibition (cf. [Grange et al., 2013](#), and as also predicted by the present model, simulation 1), small task inhibition effects may be one reason for the small effect size and marginally significant results. It is therefore reassuring that the additional analysis conducted for ordinal task difficulty, in which individual differences in task difficulty are accounted for, produced a slightly larger effect size for reaction times than the first analysis. If the procedure were modified to produce larger n-2 repetition costs, it is reasonable to expect a larger effect size. The next experiment therefore aimed to obtain stronger support for the model by modifying the procedure to produce larger n-2 repetition costs overall.

## 6. Experiment 2

### 6.1. Introduction

One experimental factor known to produce larger n-2 repetition costs is the use of non-transparent or abstract task cues, where the relationship between the cue and the task is arbitrary rather than bearing a resemblance (e.g., using geometric shapes to represent which task to perform next, as in [Houghton, Pritchard, & Grange, 2009](#)). Experiment 2 featured a modified procedure with non-transparent cues with the aim of obtaining larger n-2 repetition costs.

### 6.2. Method

#### 6.2.1. Participants

40 participants were recruited through the Birkbeck psychology participant database, of whom 28 were female. Participants received £7.50 in cash.

#### 6.2.2. Design

The design was identical to experiment 1.

#### 6.2.3. Tasks

The tasks were identical to experiment 1, with the exception that tasks were cued by a frame appearing simultaneously to, and enclosing the target letter. The frame could take the shape of a circle, square or hexagon. The association between frame shape and target task was randomised for each participant.

**Table 4**  
Mean (SD) of RTs (milliseconds) and error rates for each task.

Task	RT	Error rate
Colour	1866 (498)	5.49% (3.63%)
Font	2079 (446)	7.45% (4.43%)
Alphabet	2095 (490)	6.66% (5.31%)

**Table 5**  
Experiment 2: Mean (SD) n-2 repetition costs for RT and error rates by task pairings and switch direction.

Task pairing	Easy-hard-easy		Hard-easy-hard	
	RT (ms)	Error rate (%)	RT (ms)	Error rate (%)
Primary-Secondary	112 (250)	0.50 (5.85)	184 (303)	0.40 (6.42)
Primary-Tertiary	79 (247)	1.30 (6.03)	255 (315)	2.31 (7.72)
Secondary-Tertiary	77 (387)	0.05 (5.31)	153 (308)	1.22 (6.81)

#### 6.2.4. Procedure

As participants were required to learn cue-task associations, they received additional practice trials (compared to experiment 1), performing 15 trials for each task, before completing a practice switching block of 20 trials. The main experiment consisted of 8 blocks of 60 trials.

#### 6.3. Results

Of the 40 participants, data was not analysed from two participants who failed to complete it, and five who exceeded an error rate of 20%, leaving data from 33 participants. The data were analysed as for experiment 1. For consistency with the previous experiment, trials exceeding 20s were excluded from the reaction time analysis. However, the pattern of data was robust to stricter RT thresholds of 10 s and 5 s. Additionally, error trials and the two subsequent trials were also excluded from the RT analysis. Mean RTs and error rates for the three tasks are presented in Table 4.

Unlike experiment 1, the difference in difficulty of the three tasks was equivocal. For RT, ANOVA suggested significant differences in task difficulty,  $F(2, 64) = 7.70, p = 0.010$ , however pairwise t-tests suggested significant differences between colour and alphabet ( $t(32) = 3.27, p = 0.007$ ) and colour and font ( $t(32) = 3.29, p = 0.007$ ) pairings, but not alphabet and font ( $t(32) = 0.26, p = 0.796$ ). For error rates, ANOVA revealed overall significant differences between tasks,  $F(2, 64) = 4.44, p = 0.016$ , but pairwise t-tests suggested only significant differences between colour and alphabet ( $t(32) = 3.13, p = 0.011$ ) but not colour and font ( $t(32) = 1.52, p = 0.277$ ) or alphabet and font ( $t(32) = 1.37, p = 0.277$ ), all Holm-Bonferroni corrected for multiple comparisons.

Further inspection of the data revealed substantial individual differences in task difficulty. The ordinal ranking of task difficulty, in increasing difficulty, assessed by RTs/error rates, were *colour-alphabet-font* (11 participants for RTs/6 for error rates) and *colour-font-alphabet* (10/12), followed by *font-alphabet-colour* (5/6), *font-colour-alphabet* (4/3), *alphabet-colour-font* (2/5) and *alphabet-font-colour* (1/1).

The remaining analysis therefore proceeded according to ordinal task difficulty, as was also conducted for the previous experiment. On this basis, n-2 repetition costs (RT and error rates) for the critical conditions are shown in Table 5. The mean proportion of trials retained was 81.3%. Like the previous experiment, this proportion was consistent across individual cells, with the least trials retained for tertiary-primary switches (79.9%) and the most for tertiary-secondary switches (82.8%).

Once again, the analysis of n-2 repetition costs consisted of a  $3 \times 2$  within-subjects ANOVA, with factors of task pairing (primary-secondary vs. primary-tertiary vs. tertiary-secondary) and switch direction (EHE vs. HEH), testing the hypothesis that n-2 repetition costs for hard-easy-hard triplets would be greater than for easy-hard-easy triplets. For RTs, the results revealed that the main effect of task was not significant ( $F(2, 64) = 0.59, p = 0.557$ ), but the main effect of direction was ( $F(1, 32) = 5.92, p = 0.021, \eta_p^2 = 0.031$ ). The interaction was not significant ( $F(2, 64) = 0.64, p = 0.53$ ). Pairwise t-tests revealed these differences reached significance for the primary-tertiary task pairing ( $t(60.49) = 2.52, p = 0.014$ ) but not for primary-secondary ( $t(61.75) = 1.05, p = 0.298$ ) or secondary-tertiary ( $t(60.93) = 0.89, p = 0.379$ ).

No significant effects were found in the analysis of error rates (main effect of task:  $F(2, 64) = 0.79, p = .46$ ; main effect of direction:  $F(1, 32) = 0.81, p = .37$ ; interaction:  $F(2, 64) = 0.17, p = .85$ ).

#### 6.4. Discussion

Once again, as predicted, a significant effect of switch direction was found, with a greater n-2 repetition cost for HEH than EHE triplets. Although the finding was only significant for the task pairing with the greatest differences in difficulty for each participant (the primary-tertiary pairing), the trend was observed for each of the task pairings.

As anticipated, mean reaction times and error rates were greater for experiment 2 than experiment 1, reflecting the increased difficulty introduced by non-transparent task cueing. Unexpectedly, non-transparent task cueing appeared to modify the relative difficulty of each of the tasks, introducing substantial individual differences. However, as our hypothesis concerns the relative difficulty of tasks, rather than differences between specific tasks, the inferences from these results remain valid.

While the results are inconsistent with those of [Arbuthnott \(2008\)](#), presumably for the reasons noted in the discussion of experiment 1, they are as predicted by the model. Notably, the model predicts that HEH should produce greater n-2 repetition costs than EHE, for the same reason that the model of [Gilbert and Shallice \(2002\)](#) predicts switch cost asymmetries when switching between two tasks. Firstly, executing a harder task against interference from easier, but irrelevant, stimulus dimensions requires a greater degree of biasing by task-demand units than executing an easier task against interference from a harder task. Secondly, if the next trial is a switch, HE switches have a greater switch cost than EH switches because more task demand activation is carried over ([Gilbert & Shallice, 2002](#)). Finally, in a system with a backward inhibition mechanism the simultaneous activation of task demand units on the n-1 trial means more conflict is detected on the n-1 trial for HEH switches than EHE switches. As a result the n-2 task receives more inhibition. In contrast, on a two-switch trial, the difficulty of the n-1 task has a much less pervasive effect on performance on trial n. This overall effect, greater n-2 repetition costs for HEH trials, is precisely what was observed.

## 7. General discussion

We have presented a model of task switching that extends the basic model of [Gilbert and Shallice \(2002\)](#) with the addition of a conflict monitoring layer, and shown that the augmented model accounts for two well-established classes of empirical effects – n-1 switch-costs and n-2 repetition costs – across a theoretically justified portion of its parameter space. Additionally, the model predicts a key empirical finding – that the n-2 repetition cost is affected by the intertrial interval preceding trial n-1, but not trial n ([Gade & Koch, 2005](#)). This latter effect has been widely accepted as evidence that the n-2 repetition cost genuinely reflects task inhibition stemming from conflicting task representations, rather than a task interference effect akin to n-1 switch costs. Moreover, we have demonstrated that the model predicts that the magnitude of the n-2 repetition cost will be modulated by task difficulty (being greater for hard-easy-hard switches than easy-hard-easy switches). This prediction was then confirmed by two novel empirical studies. We structure this discussion of the findings around three broad issues: the theoretical interpretation of our empirical results; the relationship between the proposed model and existing models of task switching; and limitations of the proposed model.

### 7.1. Empirical findings and theoretical interpretation

Our empirical results extend the growing literature on n-2 repetition costs by showing that the effect is modulated by the relative difficulty of the tasks involved. The typical size of the basic n-2 repetition cost (e.g., 31–35 ms, as reported by [Mayr & Keele, 2000](#)) is smaller than the typical size of the n-1 switch cost (e.g., 200 ms, as reported by [Rogers & Monsell, 1995](#)), and consequently it is harder to obtain positive evidence for the cost.<sup>10</sup> The relative difference in size is not surprising given that n-2 repetition costs relate to effects that persist across two switches (and a longer temporal window), as opposed to one, but our finding that the n-2 repetition cost is also sensitive to task difficulty (albeit in the opposite direction to the only previous study of asymmetric repetition costs; though see the earlier discussion for potential reasons why) further suggests that considerable methodological care is required to isolate such costs.

At a theoretical level, the empirical findings argue against attributing switch costs (solely) to reconfiguration or cue retrieval, but rather support the role of cross-task interference in the switch cost, as the interference/conflict asymmetry between easy-hard and hard-easy switches plays an important role in producing asymmetric n-2 repetition costs on subsequent trials.

N-2 repetition costs are typically attributed to residual task inhibition. In the model presented here, task demand units receive both inhibitory and excitatory inputs, hence ‘task inhibition’ may be an overly simplistic label. However, consistent with the original hypothesis, units sensitive to conflict at the level of task representations inhibit those same task representations. Rather than persistent inhibition, it is the activation of conflict units that persists across trials. Thus, conflict generated on switch (n-1) trials causes inhibition when returning to the original task on trial n. The effect is to effectively smooth performance in switch trials, at the cost of interference when resuming a recently abandoned task. As a side-effect, the reduction in switch cost contributes to the switch cost asymmetry, suggesting that one component of the switch cost may be due to task inhibition that is modulated by inter-task conflict.

A beneficial effect of the proposed mechanism is facilitation of performance following under-controlled tasks, as seen in the 2SW condition of simulation 4. This suggests that conflict units insulate switching performance against deterioration when top-down control is lower than ideal, such as in the case of distraction or divided attention. This would be beneficial if deployment of top-down control is effortful or somehow limited in capacity, as is often assumed. One possibility is that

<sup>10</sup> There is evidence of substantial individual differences in n-2 repetition costs, with [Grange and Juvina \(2015\)](#) recently reporting costs of up to 200 ms for certain individuals.

conflict units might serve a role of 'smoothing' performance if control is diverted (e.g., in the case of distraction). However, it is notable that the benefit of conflict units is most observed in the 2SW condition – i.e. when rapidly switching between different task-sets. The trade-off is weaker performance when repeating the same task. Thus, conflict units might serve the function of an intermediate control layer – 'dumb' units that are unselective/uncontrolled as to the target of inhibition, but that effectively facilitate performance in contexts requiring control, such as task switching. Such units might provide an automatic, low-level control layer, reserving top-down attentional biasing for the 'heavy lifting' of deliberative control. Additionally, the effect of conflict units is heavily modulated by top-down control. Specifically, in switching conditions (1SW, 2SW) conflict units, and backward inhibition more generally, protects performance following a highly controlled task (such as a simple task with a high cost of failure – imagine carrying an antique vase across a polished floor), effectively protecting subsequent tasks against distraction.

A further theoretical question is how the proposed backward inhibition mechanism fits within an overall cognitive architecture. Should it properly be regarded as an executive function, perhaps belonging to a supervisory attentional system (Norman & Shallice, 1986)? Alternatively, Mayr and Keele (2000) prefer an account in which task-set inhibition is a relatively automatic process on the basis that once established, task-set inhibition appears to be resistant to high-level intervention, such as is afforded by a preparation interval. In the model presented here, the task inhibition mechanism is relatively low-level, being automatically triggered by task conflict. However, given its role in facilitating rapid switching between tasks, it may be more properly conceptualised as one sub-component of a more general 'task switching' executive function. Whether all such functions should be viewed as residing within a putative supervisory attentional system remains a separate issue.

### 7.2. Relationship with other models of backward inhibition

Our model implements a backward inhibition mechanism as a general purpose task-inhibition mechanism which responds to task conflict through a general inhibitory influence on task representations. This has a disproportionately large effect on representations which are no longer being excited from another source (such as top-down control) hence it reduces cross-task interference.

In contrast, the model of Grange and colleagues (Grange & Juvina, 2015; Grange et al., 2013) conceptualises backward inhibition as the self-inhibition of task-sets, as part of a memory retrieval process. Grange and Juvina (2015) generated a specific behavioural prediction of their model – the change in  $n-2$  repetition costs following practice – and confirmed this prediction empirically. Although the authors argued that inhibition is proportional to the interference in the system, and thus related to conflict while not being triggered by it, the model does not appear to easily accommodate the robust empirical finding that backward inhibition is modulated by manipulations of inter-task conflict (Koch et al., 2010). Additionally, as discussed earlier it does not seem to generalise to reproducing established empirical effects from two-task switching (e.g., switch costs) without changes to the model. In the present model, it has been assumed that the same cognitive mechanism, without any re-parameterisation, is responsible for producing  $n-1$  switch costs and  $n-2$  repetition costs. The empirical observation of  $n-1$  switch costs, rather than facilitation, is a key piece of evidence against a self-inhibitory account of task inhibition (Koch et al., 2010). However, one possibility raised by Grange et al. (2013) is that task inhibition may be strategically recruited by the human participant in three-task switching paradigms, but not in two-task paradigms, as it provides a selective advantage in situations where task repeats are not possible. This view has some empirical support. While some studies have shown that  $n-2$  repetition costs occur when  $n-1$  repeats are possible (Arbuthnott, 2005), others have shown that the  $n-2$  repetition cost is significantly reduced in such circumstances, (e.g., Philipp & Koch, 2006), with yet others suggesting that consistent  $n-1$  switch and  $n-2$  repetition costs tend not to occur in the same block of trials (Philipp & Koch, 2006, for review). While reconfiguration remains an intriguing possibility, current evidence offers only weak support, and the cognitive system responsible for such strategic (re)configuration remains to be specified in theoretical terms.

Arguably, the strategic recruitment of task inhibition has some parallels in the present model in the *weight* parameter, which controls the strength of inhibitory biasing. However, it is an empirical question whether this parameter models something fixed or variable in a human cognitive system. In other words, does the conflict system exert more or less biasing on task representations in different contexts, and if so, how is it regulated? This possibility is a topic for future study.

### 7.3. Limitations of the present model

A further difference between our model and that of Grange and colleagues (Grange & Juvina, 2015; Grange et al., 2013) relates to the conceptualisation of backward inhibition. In the Grange et al. model, backward inhibition is a process of self-inhibition as part of the cue-task memory retrieval mechanism that is absent from our model. One reason for the difference in theoretical emphasis lies in the different experimental paradigms. The set-shifting paradigm of Mayr and Keele (2000), favoured by Grange and colleagues, emphasises cue-task memory retrieval in order to select the correct target, determined by singleton status along the cued dimension. In contrast, the item classification paradigm favoured by Gade, Koch and colleagues and used in our own experiments emphasises switching between different stimulus-response mappings according to the cued task, in response to a single target. While the bulk of empirical work on backward inhibition suggests conflict

occurring at the task and response levels is most critical for backward inhibition (e.g., Gade & Koch, 2007a; Philipp, Jolicoeur, Falkenstein, & Koch, 2007; Schuch & Koch, 2003), conflict occurring during the cue-processing stage is known to also affect n-2 repetition costs (e.g., Arbutnott & Woodward, 2002; Houghton et al., 2009). Indeed, the typical increase in n-2 repetition costs as a result of increased cue-processing difficulty forms part of the rationale for the present experiment 2. In the present work, however, we have not attempted to model the cue-encoding stages of the task (e.g., translating a particular polygon into a representation of the 'colour' task) which presumably include some form of working memory retrieval. Instead, the model assumes that the necessary cue processing stages of the task have already taken place and that once encoded, the activation of the target cue representation is continuous and discrete (i.e., either fully on or off). Clearly, an end-to-end model of task processing must include a cue encoding stage. While models of cue-encoding have been proposed within the task-switching literature (e.g., Logan & Bundesen, 2003; Schneider & Logan, 2005), the nature of the cue-encoding process remains an open theoretical question. Recent research has suggested that cue-encoding itself consists of multiple stages or processes (Schneider, 2016). For now, integrating a theory and model of cue-encoding with the present model remains a future prospect. However, without modelling these stages, our model does not fully explain how conflict at the cue encoding level contributes to the n-2 repetition cost. Each model, therefore, appears to account for slightly different aspects of backward inhibition. It remains an open question whether backward inhibition is a single mechanism that can be deployed at any stage of cognitive processing, or whether backward inhibition is a somewhat generic process that occurs independently at multiple stages of processing. Answering this question, and the possibility of reconciling both types of cognitive model, is one goal of future empirical and modelling research.

Within the present implementation, conflict units are sensitive to simultaneously active task-demand units, rather than response units. This arrangement was motivated by modelling convenience rather than representing a strong theoretical claim. As the response and task demand layer have reciprocal connections, response units tend to activate relevant task-demand units and vice versa. Hence, conflict calculated at the response and task-demand levels would be expected to correlate. While some data suggests that task conflict and response conflict may be separable (e.g., Steinhauser & Hübner, 2009), clear evidence of a dissociation between the two types of conflict has yet to emerge. More critically, a range of evidence suggests response processes have a critical role in backward inhibition (Koch et al., 2010). Accommodation of these findings within the current model would seem to require an elaboration of the model's response processes, at least. Alternatively, these phenomena may be better explained by a model in which task inhibition is triggered by response, not task, conflict. Development of such a model, and detailed behavioural comparisons on a range of simulated experimental paradigms, elucidating the role of response or task conflict in task inhibition, is a goal for future research.

One further area of difference between the model and the present empirical data is the simulation of error trials. In human participants, a trade-off is observed between speed and accuracy: it appears that participants regulate their level of performance to give either fast, but more error-prone responses, or slower, more accurate responses. Indeed, providing error feedback in the empirical work was designed to encourage participants to regulate their performance appropriately for each task. In the model, however, error trials tend to occur with very long RTs, while the fastest responses tend to be the most accurate. This is because responses are not produced until one response unit's activation is sufficiently more than that of the next largest, incompatible unit. Ordinarily, correct responses may occur quickly. However, due to a combination of residual task-demand activation (irrelevant task interference) and noise, errors may occur if an incorrect response unit is able to exceed that of the correct unit. In this case, intense competition typically occurs, with the activation of both units at a similar level, before one unit 'wins' and its activation exceeds the response threshold. Thus, although the simulations have been successful in predicting the modulation of n-2 repetition costs, the model fails to capture this aspect of behaviour due, in our view, to a simplified model of response production.

#### 7.4. Conclusion

Through combining two previous cognitive models – task switching and conflict monitoring – we have produced a novel model of a backward inhibition mechanism. Primarily, the model demonstrates the viability of the proposal that backward inhibition is triggered by conflict between task representations, specifically on the n-1 trial (Koch et al., 2010). In the model, a series of conflict-detector units bias processing between competitive task representations, when task switching, through a generalised inhibition of all conflicting task representations. This is sufficient to produce n-2 repetition costs, given that the relevant task representation for each task is being excited, presumably as the result of a cue-driven memory retrieval process.

The model reproduces observed findings related to both n-1 switch costs and n-2 repetition costs, and additionally predicts that n-2 repetition costs will be modulated by asymmetries in task difficulty. This prediction was supported by two experiments which adapted the commonly used item classification paradigm, and in so doing demonstrated n-2 repetition costs for a novel combination of three tasks.

#### Appendix A. Model parameter settings

See [Table A1](#).

**Table A1**  
Model parameter settings.

Parameter	Simulation 1	Simulation 2	Simulation 3	Simulation 4
Unit activation max	1.0	1.0	1.0	1.0
Unit activation min	−1.0	−1.0	−1.0	−1.0
Response Threshold	0.15	0.15	0.15	0.15
Step Size	0.0015	0.0015	0.0015	0.0015
Squashing Parameter (task demand units)	0.8	0.8	0.8	0.8
Noise	0.006	0.006	0.006	0.006
Bias (output units)	−6.0	−6.0	−6.0	−6.0
Bias (task demand units)	−4.0	−4.0	−4.0	−4.0
Stimulus Input Strength (A)	3.0	3.0	3.0	3.0
Stimulus Input Strength (B)	3.0	3.0	3.0	1.5–5.0
Stimulus Input Strength (C)	3.0	3.0	3.0	3.0
Topdown Control Strength (A)	12.0	12.0	12.0	12.0
Topdown Control Strength (B)	12.0	12.0	6.0	6.0–20.0
Topdown Control Strength (C)	12.0	12.0	12.0	12.0
Output unit lateral weights (incongruent)	−2.0	−2.0	−2.0	−2.0
Output unit lateral weights (congruent)	2.0	2.0	−2.0	2.0
Output unit - task demand weights (congruent)	1.0	1.0	1.0	1.0
Output unit - task demand weights (incongruent)	−1.0	−1.0	−1.0	−1.0
Task demand - output inhibitory weights	−2.5	−2.5	−2.5	−2.5
Task demand - output excitatory weights	2.5	2.5	2.5	2.5
Squashing Parameter (conflict units)	0.5	0.5	0.5	0.5
Bias (conflict units)	−10.0	−40 to 0	−7.5	−10
Gain (conflict units)	75.0	0–100	39.0	75.0
Conflict - task demand unit weights	−10.0	−30 to 0	−14.0/0	−2.6 to 0

## References

- Allport, A., Styles, E., & Hsieh, S. (1994). Shifting intentional set – Exploring the dynamic control of tasks. In C. Umiltà & M. Moscovitch (Eds.), *Attention and performance xv: Conscious and nonconscious information processing* (Vol. 15, pp. 421–452). Cambridge: MIT Press.
- Allport, A., & Wylie, G. (2000). Task switching, stimulus-response bindings, and negative priming. In *Control of cognitive processes attention and performance XVIII* (pp. 35–70). Cambridge, Mass: MIT Press.
- Altmann, E. M., & Gray, W. D. (2008). An integrated model of cognitive control in task switching. *Psychological Review*, 115(3), 602–639.
- Arbuthnott, K. D. (2005). The influence of cue type on backward inhibition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31(5), 1030–1042.
- Arbuthnott, K. D. (2008). Asymmetric switch cost and backward inhibition: Carryover activation and inhibition in switching between tasks of unequal difficulty. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale*, 62(2), 91–100.
- Arbuthnott, K. D., & Woodward, T. S. (2002). The influence of cue-task association and location on switch cost and alternating-switch cost. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale*, 56(1), 18–29.
- Botvinick, M. M., Braver, T., Barch, D., Carter, C., & Cohen, J. (2001). Conflict monitoring and cognitive control. *Psychological Review*, 108(3), 624–652.
- Brown, J. W., Reynolds, J. R., & Braver, T. S. (2007). A computational model of fractionated conflict-control mechanisms in task-switching. *Cognitive Psychology*, 55(1), 37–85.
- Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: A parallel distributed processing account of the stroop effect. *Psychological Review*, 97(3), 332–361.
- Cohen, J. D., & Huston, T. (1994). Progress in the use of interactive models for understanding attention and performance. In C. Umiltà & M. Moscovitch (Eds.), *Attention and performance XV: Conscious and nonconscious information processing* (pp. 453–476). MIT Press.
- Cooper, R., & Shallice, T. (2000). Contention scheduling and the control of routine activities. *Cognitive Neuropsychology*, 17(4), 297–338.
- Gade, M., & Koch, I. (2005). Linking inhibition to activation in the control of task sequences. *Psychonomic Bulletin & Review*, 12(3), 530–534.
- Gade, M., & Koch, I. (2007a). Cue-task associations in task switching. *Quarterly Journal of Experimental Psychology*, 60(6), 762–769.
- Gade, M., & Koch, I. (2007b). The influence of overlapping response sets on task inhibition. *Memory & Cognition*, 35(4), 603–609.
- Gade, M., Schuch, S., Druey, M., & Koch, I. (2014). Inhibitory control in task switching. In *Task switching and cognitive control* (pp. 137). Oxford University Press.
- Gilbert, S. J., & Shallice, T. (2002). Task switching: A PDP model. *Cognitive Psychology*, 44(3), 297–337.
- Goschke, T. (2000). Intentional reconfiguration and involuntary persistence in task set switching. In S. Monsell & J. Driver (Eds.), *Control of cognitive processes. Attention and performance XVIII* (pp. 331–355). Cambridge, Mass: MIT Press.
- Grange, J. A., & Houghton, G. (2009). Temporal cue-target overlap is not essential for backward inhibition in task switching. *Quarterly Journal of Experimental Psychology*, 62(10), 2068–2079.
- Grange, J. A., & Houghton, G. (2014). Models of cognitive control in task switching. In *Task switching and cognitive control* (pp. 160–199). Oxford University Press.
- Grange, J. A., & Juvina, I. (2015). The effect of practice on n-2 repetition costs in set switching. *Acta Psychologica*, 154, 14–25.
- Grange, J. A., Juvina, I., & Houghton, G. (2013). On costs and benefits of n-2 repetitions in task switching: Towards a behavioural marker of cognitive inhibition. *Psychological Research-Psychologische Forschung*, 77(2), 211–222.
- Houghton, G., Pritchard, R., & Grange, J. A. (2009). The role of cue-target translation in backward inhibition of attentional set. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35(2), 466–476.
- 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(5), 849–874.
- Koch, I., Gade, M., & Philipp, A. M. (2004). Inhibition of response mode in task switching. *Experimental Psychology (formerly Zeitschrift für Experimentelle Psychologie)*, 51(1), 52–58.
- Koch, I., Gade, M., Schuch, S., & Philipp, A. M. (2010). The role of inhibition in task switching: A review. *Psychonomic Bulletin & Review*, 17(1), 1–14.
- 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(3), 575–599.

- Logan, G. D., & Schneider, D. W. (2010). Distinguishing reconfiguration and compound-cue retrieval in task switching. *Psychologica Belgica*, 50(3), 413–433.
- Longman, C. S., Lavric, A., Munteanu, C., & Monsell, S. (2014). Attentional inertia and delayed orienting of spatial attention in task-switching. *Journal of Experimental Psychology: Human Perception and Performance*, 40(4), 1580.
- Mayr, U., & Keele, S. W. (2000). Changing internal constraints on action: The role of backward inhibition. *Journal of Experimental Psychology: General*, 129(1), 4–26.
- McClelland, J. L. (1993). Toward a theory of information processing in graded, random, and interactive networks. In D. E. Meyer & S. Kornblum (Eds.), *Attention and performance 14: Synergies in experimental psychology, artificial intelligence, and cognitive neuroscience* (pp. 655–688). Cambridge, MA, US: The MIT Press.
- Meiran, N., Chorev, Z., & Sapir, A. (2000). Component processes in task switching. *Cognitive Psychology*, 41(3), 211–253.
- Norman, D., & Shallice, T. (1986). Attention to action: Willed and automatic control of behavior. In R. Davidson, R. Schwartz, & R. Shapiro (Eds.), *Consciousness and self-regulation: Advances in research and theory IV*. Plenum Press.
- Peirce, J. W. (2009). Generating stimuli for neuroscience using psychopy. *Frontiers in Neuroinformatics*, 2, 10.
- Philipp, A. M., Jolicoeur, P., Falkenstein, M., & Koch, I. (2007). Response selection and response execution in task switching: Evidence from a go-signal paradigm. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33(6), 1062–1075.
- Philipp, A. M., & Koch, I. (2006). Task inhibition and task repetition in task switching. *European Journal of Cognitive Psychology*, 18(4), 624–639.
- Pitt, M. A., Kim, W., Navarro, D. J., & Myung, J. I. (2006). Global model analysis by parameter space partitioning. *Psychological Review*, 113(1), 57.
- Roberts, S., & Pashler, H. (2000). How persuasive is a good fit? A comment on theory testing. *Psychological Review*, 107(2), 358–367.
- Rogers, R. D., & Monsell, S. (1995). Costs of a predictable switch between simple cognitive tasks. *Journal of Experimental Psychology: General*, 124(2), 207–231.
- Rubinstein, J. S., Meyer, D. E., & Evans, J. E. (2001). Executive control of cognitive processes in task switching. *Journal of Experimental Psychology: Human Perception and Performance*, 27(4), 763–797.
- Schneider, D. W. (2014). Modeling graded response congruency effects in task switching. *Acta Psychologica*, 153, 160–168.
- Schneider, D. W. (2016). Perceptual and conceptual priming of cue encoding in task switching. *Journal of Experimental Psychology: Learning, Memory, and Cognition*.
- Schneider, D. W., & Logan, G. D. (2005). Modeling task switching without switching tasks: A short-term priming account of explicitly cued performance. *Journal of Experimental Psychology: General*, 134(3), 343.
- Schuch, S., & Koch, I. (2003). The role of response selection for inhibition of task sets in task shifting. *Journal of Experimental Psychology: Human Perception and Performance*, 29(1), 92–105.
- Sohn, M. H., & Anderson, J. R. (2001). Task preparation and task repetition: Two-component model of task switching. *Journal of Experimental Psychology: General*, 130(4), 764.
- Steinhauser, M., & Hübner, R. (2009). Distinguishing response conflict and task conflict in the stroop task: Evidence from ex-gaussian distribution analysis. *Journal of Experimental Psychology: Human Perception and Performance*, 35(5), 1398–1412.
- Vandierendonck, A., Liefoghe, B., & Verbruggen, F. (2010). Task switching: Interplay of reconfiguration and interference control. *Psychological Bulletin*, 136(4), 601–626.
- Yeung, N., & Monsell, S. (2003). Switching between tasks of unequal familiarity: The role of stimulus-attribute and response-set selection. *Journal of Experimental Psychology: Human Perception and Performance*, 29(2), 455–469.