

Task Switching: A PDP Model

Sam J. Gilbert and Tim Shallice

Institute of Cognitive Neuroscience, University College London, London, United Kingdom

When subjects switch between a pair of stimulus–response tasks, reaction time is slower on trial N if a different task was performed on trial $N - 1$. We present a parallel distributed processing (PDP) model that simulates this effect when subjects switch between word reading and color naming in response to Stroop stimuli. Reaction time on “switch trials” can be slowed by an extended response selection process which results from (a) persisting, inappropriate states of activation and inhibition of task-controlling representations; and (b) associative learning, which allows stimuli to evoke tasks sets with which they have recently been associated (as proposed by Allport & Wylie, 2000). The model provides a good fit to a large body of empirical data, including findings which have been seen as problematic for this explanation of switch costs, and shows similar behavior when the parameters are set to random values, supporting Allport and Wylie’s proposal. © 2001 Elsevier Science

Key Words: task switching; task set; Stroop effect; parallel distributed processing; executive functions.

Atkinson and Shiffrin (1968) proposed a distinction between relatively permanent cognitive structures, such as short- and long-term memory, and *control processes* which harness those fixed structures in order to attain specific goals. This distinction was elaborated in the following years (e.g., Posner & Snyder, 1975; Shiffrin & Schneider, 1977) and has been generally accepted (Shallice, 1994). Yet research with normal subjects in the decades following Atkinson and Shiffrin’s article paid relatively little attention to the control processes required to select and organize fixed cognitive structures (with some notable exceptions, e.g., Gopher, Weil, & Siegel, 1989; Logan, 1979, 1980, 1985; Spelke, Hirst, & Neisser, 1976).

In recent years, studies have begun to investigate the control processes of normal subjects when they switch between different cognitive tasks on

The research reported in this article was supported by a research studentship awarded to Sam Gilbert by the UK Medical Research Council. We thank Gordon Logan, Nachshon Meiran, Glenn Wylie, Nick Yeung, and an anonymous reviewer for their helpful comments on an earlier version of this article.

Address correspondence and reprint requests to Sam Gilbert, Institute of Cognitive Neuroscience, Alexandra House, 17 Queen Square, London WC1N 3AR, United Kingdom. E-mail: sam.gilbert@ucl.ac.uk. Fax: +44 (0)20 7916 8517.



successive trials (e.g., Allport, Styles, & Hsieh, 1994; Meiran, 1996; Rogers & Monsell, 1995; and contributions to Monsell & Driver, 2000). The task-switching paradigm therefore appears to offer a valuable tool for studying “executive control” (Logan, 1985; Monsell, 1996), both in normal subjects and in patients (e.g., Rogers, Sahakian, Hodges, Polkey, Kennard, & Robbins, 1998). However, there has been disagreement over the interpretation of experiments involving task switching, making it difficult to draw firm conclusions from them about executive control. Here, we seek to clarify this debate by presenting a computational model of task switching.

MAIN EMPIRICAL FINDINGS AND THEORETICAL INTERPRETATIONS

A number of methodologies have been developed for studies of task switching. One approach is simply to compare pure and mixed blocks, i.e., to compare blocks where the subject performs the same task on every trial with blocks in which s/he must alternate between two tasks on successive trials (e.g., Allport et al., 1994; Jersild, 1927; Spector & Biederman, 1976). Other studies have employed the “alternating runs” paradigm (Rogers & Monsell, 1995), where subjects are required to switch tasks predictably every n th trial, where n is at least 2. This has the advantage that “switch” trials (where the task differs from the one performed on the previous trial) can be compared with “nonswitch” (or “repeat”) trials within the same block. In a third methodology for task switching experiments, each trial is preceded by a cue which instructs the subjects which task to perform (e.g., Meiran, 1996; Sudevan & Taylor, 1987). This paradigm makes the requirement to switch tasks unpredictable.

Experiments using each of these paradigms have demonstrated costs, both in reaction time and error rate, for switch compared with nonswitch trials (or mixed compared with pure blocks). These “*switch costs*” have been reported to vary from zero to several hundred milliseconds per item, depending on the experimental conditions (see below). The main test of the model presented in this article will be whether it can predict the effects of various experimental manipulations on switch costs.

Explanations of the Switch Cost

An important concept in studies of task switching is the “task set,” which may be defined loosely as the set of cognitive operations required to effectively perform a task (“To form an effective intention to perform a particular task, regardless of which of the range of task-relevant stimuli will occur, is to adopt a *task-set*”; Rogers & Monsell, 1995, p. 208, emphasis in original). This concept is often used in the same way as the concept of the action or thought “schema” (Norman & Shallice, 1986). The distinction between individual S-R mappings and task sets, composed of all of the individual S-R

mappings which are required to carry out an experimental task, is crucial. As we shall see, an important issue concerns the degree to which task switching should be understood in terms of processes occurring at the level of discrete S-R mappings rather than those which occur at the level of the task set (see also Monsell, Taylor, & Murphy, 2001).

Two main theoretical accounts of the switch cost have been put forward. Allport and colleagues (e.g., Allport et al., 1994; Allport & Wylie, 2000) have suggested that switch costs index an interference effect caused by a carryover of the previous task set into switch trials (this will be called the “task carryover account”). As well as a carryover of the previously active task set into switch trials, inhibition of competing task sets can also persist on switch trials, when the previously competing task is now required. According to this theory, there is no need to postulate differences between the higher level cognitive processes that occur during switch and nonswitch trials in order to explain switch costs; switch trials are simply prolonged by greater competition caused by the carryover effect (i.e., switch costs reflect a form of priming). This account does not deny the involvement of control processes in task switching. Without such processes, the subject could never switch tasks at all. What *is* denied by this account is that these control processes are measured in any direct way by the switch cost.

An opposing explanation of switch costs has been advanced by Monsell and colleagues (Monsell, Yeung, & Azuma, 2000; Rogers & Monsell, 1995). This hypothesis proposes that switch costs *do* reflect the duration of a stage-like executive control process which reconfigures the cognitive system for the upcoming task. However, it is held that this process cannot be completed until the arrival of the first stimulus of the new task (i.e., it is a stimulus-driven “exogenous control process”). An additional component of the switch cost is hypothesized to reflect the operation of an “endogenous control process” which can be executed before the arrival of the first stimulus in the new task. In a similar account, Rubinstein et al. (2001) attribute switch costs to the duration of a “goal shifting” stage, which can be executed before the arrival of a stimulus, and a “rule activation” stage which must await stimulus presentation.

Evidence for the Task Carryover Account

Evidence for the task carryover account of switch costs has come from findings which demonstrate an effect of the performance of an earlier task on subsequent switch costs (Allport et al., 1994; Allport & Wylie, 2000; see also Mayr & Keele, 2000). Two of these findings are of particular theoretical relevance.

Asymmetric switch costs. Allport et al. (1994, Experiment 5) carried out an experiment where subjects switched between word reading and color naming in response to incongruent Stroop stimuli (e.g., “red” written in blue ink; Stroop, 1935). As expected, word reading yielded faster reaction

times than color naming (see MacLeod, 1991). However, this experiment also produced a most unexpected result. Reaction times for word-reading trials were slower when color naming had been performed on the previous trial (i.e., there was a switch cost), but Allport et al. failed to detect a switch cost for color-naming trials that followed performance of the word-reading task. In other words, the switch cost appeared to be confined to the switch from the nondominant into the dominant (i.e., better learned, easier) task.

More recent experiments have replicated this asymmetry in switch costs (Allport & Wylie, 2000; Wylie & Allport, 2000). Allport and Wylie found that there is a cost for a switch into the color-naming task, but it is smaller than the cost of a switch into the word-reading task. Meuter and Allport (1999) reported an analogous finding when subjects switched between digit naming in their dominant and nondominant languages. This paradoxical finding—larger reaction time costs for a switch into a better learned, more dominant task—is difficult to explain if switch costs reflect the time taken to reconfigure the cognitive system for the upcoming task. Why should a switch into an easier, better learned task take longer to complete than a switch into a less familiar task? Allport et al. argue that the result can be explained if the primary determinant of the switch cost is the nature of the previous task. When subjects name the color of an incongruent Stroop stimulus, Allport et al. hypothesize that inhibition of the word-reading task may be required. According to the task carryover account, this inhibition will persist on a switch trial where word-reading is now appropriate, leading to a large switch cost. But in the absence of any requirement to suppress color naming in order to perform the word-reading task, there will be no carryover of inhibition into color naming switch trials, hence the small or absent switch costs.

Item specific costs. Allport and Wylie (2000, Experiment 5) investigated the extent to which the switch cost is caused by the repetition of stimuli between the two tasks. Subjects alternated between short runs of color naming and word reading, but only a subset of the stimuli that appeared in the word-reading task were also presented for color naming. Thus, it was possible to compare two types of stimuli in the word-reading task: “primed” stimuli, which had also appeared recently in the color-naming task, and “unprimed” stimuli, which were only ever seen in the word-reading task. Allport and Wylie found that the reaction time to primed stimuli on switch trials was slower than the reaction time to unprimed stimuli, but there was no reliable difference in reaction time between primed and unprimed stimuli on repeat trials. Thus, there was a greater switch cost when the stimulus on the switch trial was primed. This item specific component of the switch cost provides further evidence for the hypothesis that switch costs involve a carryover effect caused by prior performance of a different task. The finding suggests that stimuli might themselves evoke the task sets with which they

were recently associated, even when this task set is not appropriate. As a result, Allport and Wylie (2000) have updated the “task set inertia” (TSI) theory of Allport et al. (1994). They propose an “associative-TSI” account of switch costs, according to which stimuli are able to evoke recently associated task sets from memory. When a stimulus appears on one trial, associated with task A, and then reappears on a subsequent switch trial, requiring performance of task B, Allport and Wylie hypothesize that task set A may nevertheless be evoked by the presentation of the stimulus (cf. perceptual “trigger conditions” in Norman & Shallice, 1986). The resulting competition between task sets A and B may lead to an extended response selection process, explaining the enhanced switch cost obtained for word reading when the stimulus was recently presented for color naming.

Evidence for the Exogenous Control Process Account

In arguing for the exogenous control process account, Monsell and colleagues (Monsell et al., 2000; Rogers & Monsell, 1995) have demonstrated the robustness of the switch cost even when subjects have long intervals between trials to prepare for the upcoming task. They provided evidence that there is no further reduction in switch costs after the first 600 ms or so of the preparation interval: an asymptotic “residual switch cost” remains (e.g., Rogers & Monsell, 1995). This is hypothesized to correspond with the time taken by the execution of the exogenous control process, which must await stimulus presentation and is therefore insensitive to the preparation interval.

A second line of evidence for the exogenous control process account comes from studies which have investigated task switching using the alternating runs paradigm when there are more than two trials of each task before a switch. Experiments with run lengths of four (Rogers & Monsell, 1995, Experiment 6) and eight (Monsell, Azuma, Eimer, Le Pelley, & Strafford, 1998) trials before each switch have found that the cost of task switching is confined to the first trial of a run (see Fig. 1). Rogers and Monsell (1995) argue that if switch costs result from a carryover effect from the previous task, they should dissipate more gradually, rather than being limited to the first trial of a run. Similarly, Monsell (1996) points out that “although having ample time to prepare before a predictable task switch does not eliminate switch cost, performing the task just once appears to do so” (p. 138). This pattern of results does not seem to be compatible with the idea that switch costs simply reflect a form of priming. It can more easily be explained if switch costs reflect a one-off process of task set reconfiguration at the beginning of switch trials.

An empirical argument against the task carryover account has been made by Monsell et al. (2000). They demonstrated a number of cases in which the cost of a switch from a stronger to a weaker task is greater than the cost of a switch in the reverse direction, i.e., the opposite asymmetry of switch




FIG. 1. Mean reaction time (RT) and error rate in Rogers and Monsell's (1995) Experiment 6, as a function of the position in a run of four trials, averaged over both tasks, adapted from Rogers and Monsell (1995, Fig. 5).

costs to that reported by Allport et al. (1994). Thus, it does not seem to be a general rule that switches from a less dominant to a more dominant task yield greater RT costs than switches in the reverse direction.

Ironically, recent evidence (e.g., Meiran et al., 2000; Salthouse, Fristoe, McGurthy, & Hambrick, 1998) has also shown that it is not a general rule that switch costs are confined to the first trial in a run. Nevertheless, it still appears to be a challenge to the task carryover account that switch costs can be confined to the first trial in a run, at least sometimes. Similarly, the finding of larger costs of a switch into a dominant task, even if it is not universal, appears to challenge the exogenous control process account.

It is of course possible that the switch cost measures a combination of both a carryover of task set (leading to an extended response selection process) *and* an exogenous control process (see Meiran, 2000a, 2000b, for such an account). Indeed, Monsell et al. (2000, p. 254) have pointed out that "it is perfectly possible that an extra 'control' process is required precisely in order to overcome . . . interference, once it arises as the result of the stimulus retrieving a recently activated task-set or the inhibition associated with a recently suppressed task-set." Thus, the task carryover and exogenous control process accounts need not be mutually exclusive. However, before considering a combination of these two accounts, we should first investigate the

abilities of each one to account for the data on its own. To do this, each account needs to be specified more precisely in order to assess its consistency with the available empirical data (of course, a sufficiently vague theory is consistent with *any* empirical data). Second, we need to establish where the crucial differences between the theories lie and which differences are better classified as differences in emphasis. Computational modeling can play a role in both of these steps. First, by implementing a theory in a computational model, the theory can be specified precisely, allowing a clear assessment of its ability to explain the empirical data. Second, if it were possible to produce a model onto which two competing theories may be mapped, this would suggest that the difference between the theories is not a fundamental one.

A crucial area of disagreement is the extent to which involuntary persistence of task sets can explain the dramatic improvement in reaction time from the first to the second trial in a run, following a switch of task. Although Monsell et al. (2000) “certainly accept that there are relatively long-term carry-over effects of the kind that Allport and colleagues have demonstrated,” they are inclined to doubt whether the dramatic improvement in RT from switch trials to immediately successive nonswitch trials can be explained “even in part” by a carryover effect. Thus, the exogenous control process account does not deny that reaction time in some circumstances may be affected by a carryover of task set, but it claims that this is insufficient as an explanation of the drop in reaction time from switch to subsequent nonswitch trials.

In order to test this claim, we have implemented a version of the task carryover account in a computational model. In the remainder of this article, we take an existing computational model of task performance in pure-block conditions and extend it in accordance with the task carryover account of switch costs. We then test its performance in mixed-blocks against the empirical data. Our reasoning is as follows. If an existing computational model of pure-block performance, augmented with the mechanisms required by the task carryover account of task switching, can produce similar behavior to subjects in comparable experimental conditions, then the task carryover account would be strengthened. Furthermore, if the model is successful, this would indicate that, contrary to the arguments of Rogers and Monsell (1995) and Monsell et al. (2000), a task carryover account can offer a sufficient explanation of the relevant set of task-switching phenomena. Thus, the exogenous control process invoked by alternative accounts may be unnecessary to explain the cost of task switching.

The experimental tasks that will be simulated are word reading and color naming in response to Stroop stimuli. This domain was chosen for three reasons. First, a relatively large corpus of data has been accumulated concerning the effects of switching between dominant and nondominant tasks on reaction times in domains such as Stroop word reading/color naming (Allport et al., 1994; Allport & Wylie, 1999, 2000; Monsell et al., 2000),

bilingual language switching (Meuter & Allport, 1999), and pairs of tasks with different S-R compatibility (Monsell et al., 2000). Thus, there is enough data to allow the evaluation of a computational model against the performance of human subjects. A second reason for this choice is the theoretical weight attached to the finding of “paradoxical” asymmetric switch costs (Allport et al., 1994; Monsell et al., 2000), which requires a pair of tasks of different “strengths.” The final reason is that an influential model of the Stroop effect has been developed by Cohen and colleagues (Cohen, Braver, & O’Reilly, 1996; Cohen, Dunbar, & McClelland, 1990; Cohen & Huston, 1994; Cohen & Servan-Schreiber, 1992; see also Phaf, Van der Heijden, & Hudson, 1990; Zhang, Zhang, & Kornblum, 1999). Thus, an existing computational model of pure-block performance is already available.

This strategy is similar to the approach taken by Logan and Gordon (2001), who added control processes to an existing model of visual attention (Bundesen, 1990) in order to simulate situations that require switching between tasks. However, the type of model employed by Logan and Gordon (2001) and the phenomena simulated are rather different from those tackled here, making direct comparison between the two models difficult.

THE MODEL

Although the present model is based on the earlier models of the Stroop task by Cohen, Dunbar, and McClelland (1990) and Cohen and Huston (1994), it also has many modifications. Below, we provide a full description of our model, followed by a brief comparison with the earlier models on which it is based.

We have implemented an interactive activation model (McClelland & Rumelhart, 1981), composed of two separate pathways, for word reading and color naming (see Fig. 2). In each pathway, there are three input units (representing, in the word pathway, the words “red,” “green,” and “blue” and, in the color pathway, the colors red, green, and blue). In addition, each pathway has three output units, representing the responses “red,” “green,” and “blue.” In other words, each possible response is represented twice, once in the word-reading pathway and once in the color-naming pathway. Thus, the model has a total of six input units and six output units. Each input unit has a positive connection with its corresponding output unit. For example, in order to simulate a stimulus of the word “red” written in green ink color, the “red” word input unit and the “green” color input unit would both be activated. This would send activation to the “red” output unit in the word-reading pathway and the “green” output unit in the color-naming pathway.

Processing in the model, i.e., the passing of activation between units along their connections, is iterated for a number of cycles. This allows the simula-




FIG. 2. Architecture of the present model.

tion of reaction time: on each cycle, "evidence" is collected from the activation values of the six output units, two of which represent each possible response, "red," "green," or "blue." When the evidence for one of these three responses passes a fixed threshold, the trial is terminated. In this way, it is possible to compare the number of cycles required for the model to reach its response threshold with the mean reaction time of human subjects.

The connection strengths from the input to the output units are stronger in the word-reading pathway than in the color-naming pathway. This simulates the greater experience that people have of naming written words than colors. As a result, the word-reading output units become more strongly active than the color-naming output units when the model is presented with a Stroop stimulus. The evidence for the response represented in the word-reading pathway is therefore greater than the evidence for the response represented in the color-naming pathway and as a result the model will tend to respond by "reading out" the word that it is presented with. However, since people are able to name the ink color of a color-word, even when the color and word are incongruent, some mechanism is required to prevent the model from always executing the word-reading task. This is provided by the color-naming and word-reading "task demand" units, which send activation to their corresponding pathways. For example, when the color-naming task de-

mand unit is activated, it sends activation to the output units in the color-naming pathway, allowing them to win competition with the output units in the word-reading pathway. As well as sending a positive input to the output units of their corresponding pathway, the task demand units also send a negative (i.e., inhibitory) input to the output units of the other pathway.

The word-reading and color-naming output units also send activation back to the task demand units. This introduces feedback as well as feedforward connectivity into the model, allowing activity in the color and word pathways to modulate activity in the task demand units. Similar connections were introduced into the model of Cohen and Huston (1994) in order to simulate phenomena such as attentional capture, where stimuli are able to “draw attention” to themselves (e.g., Posner, 1980). The task demand units receive an additional “top-down control input,” which specifies which task the model should perform. The word and color output units are interconnected, so that congruent word and color response units (e.g., the two “red” units) have reciprocal positive connections and incongruent pairs of units (e.g., word “red”/color green) have reciprocal negative connections. Finally, there are lateral inhibitory connections between all units within the word output “module” (i.e., set of word output units), color output module and task demand module. This encourages the network to settle into stable states with no more than one unit active in each module.

Comparison with Earlier Models

Although the model is based on the earlier models of Cohen, Dunbar, and McClelland (1990) and Cohen and Huston (1994), its architecture differs from them in two main respects. First, there are three possible words and colors (red, green, and blue) in the present model as opposed to just two in the earlier models. This was chosen because the Cohen et al. (1990) model has been criticized for failing to capture differences between word reading and color naming when the set size is increased beyond two (Kanne, Balota, Spieler, & Faust, 1998; but see Cohen, Usher, & McClelland, 1998, for a reply).

A second difference is that, unlike our model, the earlier models included a “winner-take-all” response layer. In the Cohen et al. (1990) model, the units corresponding to the word and color output units in the present model sent inputs into a pair of response units. Thus the word-reading “red” output unit and the color-naming “red” output unit sent activation to a single “red” response unit. Evidence was collected from these response units in order to determine when each trial should end. These additional units are unnecessary in the present model: since we have interconnected the word and color output units there is a beneficial effect of activating congruent output units and a detrimental effect of activating incongruent output units. This plays a role

in our model similar to the convergent inputs into the response units of the Cohen et al. (1990) and Cohen and Huston (1994) models.

Implementation of the Task Carryover Account

The model was extended in two ways in order to implement the task carry-over account of task switch costs. First, rather than reinitializing the network at the beginning of every trial, the model was modified so that the state of the task demand units can persist into successive trials. As a result, the most recently implemented task set remains active at the beginning of the next trial and the most recently inhibited task set remains inhibited. This can be seen as an implementation of the hypothesis put forward by Allport et al. (1994) that control states persist, involuntarily, from one trial to the next. The most simple way to implement this would be to start each trial with the task demand units in the state they were in at the end of the previous trial. However, this could well lead to “perseverative” behavior, with the model unable to switch from one task into the other. Thus, a “squashing” parameter was introduced so that the activation levels of the task demand units following each response are squashed, i.e., set to some proportion of their activation levels at the end of the previous trial. Such reductions in activation levels between trials are common in models of sequential processes (e.g., Burgess, 1995; Dayan, 1998; O’Reilly & Farah, 1999) and seem to be biologically plausible. The task demand units in our model are thought to reflect activity in prefrontal cortex (see Miller & Cohen, 2001) and single-cell recording studies of monkeys performing cognitive tasks have demonstrated a sharp reduction in the firing rate of cells in prefrontal cortex following the production of a response (see Fuster, 1997, pp. 121–134; for an example in some ways akin to the task switching paradigm see Asaad, Rainer, & Miller, 2000).

An additional requirement for an implementation of the associative-TSI account, as opposed to the earlier theory of Allport et al. (1994), is some mechanism for individual stimuli to evoke task sets with which they have recently been associated. In order to achieve this, we added a connection from each of the stimulus input units to both task demand units. Thus, stimuli are able to evoke task sets by providing an extra input into the task demand units. The weights of these connections between stimulus input and task demand units are determined by Hebbian learning at the end of each trial, so that the weights between coactive units are adjusted in proportion to the product of their activation values. One potential danger with this learning algorithm is that the weights between a pair of repeatedly coactive units will grow without bound, which may lead to such a strong input into the task demand units that the model is unable to switch task. In order to avoid this problem, the weights between stimulus input and task demand units are reset to zero at the end of each trial, before the new weights are calculated, so

that the effects of learning on trial N persist only for trial $N + 1$. This is a simplifying assumption rather than a theoretical position which we adopt. There is strong evidence for such item-specific priming effects lasting for longer than one trial (e.g., Allport & Wylie, 2000, Experiment 5; Waszak, Hommel, & Allport, submitted). However, long-term priming effects, though clearly of great interest, are not addressed in this work.

Multiple Inputs to Task Demand Units

As well as the control input that indicates which task is to be performed, each task demand unit also receives an input from the stimulus input units, from the color and word output units and from the other task demand unit. When we discuss the model's performance we refer especially to two of these inputs: the control input, which indicates which task is currently appropriate (we refer to this as the "top-down input"), and the input that the task demand units receive from the stimulus input units (we refer to this as the "bottom-up input"). One crucial feature of the top-down input is that it is not equal for the two task demand units. We assume that the control mechanism provided by the task demand units is required more for the color than the word task, since the color-naming pathway is weaker (see, e.g., LaBerge & Samuels, 1974; Posner & Snyder, 1975 for similar ideas). Thus, the top-down input received by the color task demand unit, when color naming is the required task, is greater than the top-down input received by the word task demand unit on word-reading trials. As we shall see, this difference in top-down input plays an important role in the model's behavior.

Operation of the Model

The steps taken to simulate a trial are as follows: (1) For trials other than the first, the task demand units are set to a proportion of their activation values at the end of the previous trial. This proportion is set by the squashing parameter discussed above. The activations of the stimulus input and output units are set to zero.¹ (2) The appropriate top-down input is added to the net input of the color or the word task demand unit, depending on which task is required. This input is added to the task demand unit's net input on every cycle. (3) The preparation interval begins. With all of the stimulus input units set to zero, the top-down input is applied to the task demand units for the number of cycles set as the preparation interval. The activation levels of the output units are not updated during this period. (4) After the end of the preparation interval, the appropriate task demand unit is activated as before and either one or two of the stimulus input units have their activation values

¹ A more sophisticated model might incorporate persisting activation of the output units. This could be used in an attempt to simulate trial-to-trial effects of response repetition, which has been shown to interact with task switching (e.g., Rogers & Monsell, 1995). However, such effects are beyond the scope of the present model.

set to the maximum value, until the end of the trial. In order to simulate neutral trials (affording a response in only one of the tasks, e.g., “xxxxx” written in blue ink or “green” written black ink), just one of the stimulus input units is activated in this way. Congruent (e.g., “red” in red ink) and incongruent (e.g., “red” in blue ink) trials are simulated by activating two stimulus input units, one in the word-reading pathway and one in the color-naming pathway. (5) Activation is allowed to propagate until a response threshold is reached (see below). The number of cycles since stimulus presentation is recorded as the “reaction time.”

Activation Level and Weight Update

Activation levels are determined by the standard interactive activation equations (McClelland & Rumelhart, 1981). On every cycle, the net input for each unit is calculated by summing the activation values of every unit from which it receives a connection, multiplied in each case by the appropriate connection weight. In addition, units in the task demand, word and color modules each have a bias, a constant which is added to their net inputs on every cycle as well as the inputs received from other units. Each unit's activation value is then updated according to the following equations:

$$\text{If the net input is positive: } \Delta \text{act} = \text{step} \times \text{net} \times (\text{max} - \text{act}) \quad (1)$$

$$\text{If the net input is negative: } \Delta \text{act} = \text{step} \times \text{net} \times (\text{act} - \text{min}) \quad (2)$$

Where act = current activation, step = step size, net = net input, max = maximum activation value, and min = minimum activation value. The step size parameter determines the magnitude of the change in activation on each cycle, setting the speed of processing. When the activation values have been updated for each unit, the net inputs are calculated again and a new cycle begins. On each cycle, a random noise term is also added to the activation values of each unit. This term is drawn from a Gaussian distribution, with a mean of zero; the standard deviation of this distribution determines how much disruption is caused by noise on each cycle. After noise has been added, the activation levels of any units outside the maximum and minimum values are reset to the relevant extreme.

At the end of each trial, the weights between the stimulus input and task demand units are set according to the following equation:

$$w_{ij} = \text{lrate} \times a_j \times a_i \quad (3)$$

In this equation, w_{ij} is the weight of the connection from unit j to unit i , lrate sets the learning rate (i.e., the magnitude of the change in weights for each trial), and a_j and a_i are the activation levels of units j and i respectively. Note that this equation does not take into account the previous weight of the

connection between the two relevant units. Thus, these weights are calculated anew at the end of each trial, i.e., the weights derived from the activation levels of the units at the end of trial N only affect the model's behavior for trial $N + 1$.

Response Threshold

The purpose of the response threshold is to provide a way to determine the moment when enough evidence has been accumulated from the word and color output units for a response to be emitted. This is implemented as follows. At the end of each cycle, the word or color output unit with the greatest level of activation is compared with each other word and color output unit, except for the unit which commands the same response in the other stimulus dimension (e.g., the red color-naming output unit if the red word-reading output unit has the highest activation level). The difference in activation between the most active unit and the next most active unit is calculated. When this difference passes a fixed threshold the trial is terminated. Thus, a response is simulated as occurring when the amount of evidence for that response exceeds the evidence for any other response by a fixed amount. This is essentially a "random walk" model (see Logan & Gordon, 2001; Nosofsky & Palmeri, 1997; Palmeri, 1997). A similar response threshold, based on the amount of evidence for the different responses, was used by Cohen et al. (1990).

PERFORMANCE OF THE MODEL

The parameters, including the connection weights, were initially set by hand and remain fixed for the first set of simulations described below, except where stated (see the Appendix for parameter values). We show later that the model's behavior generally remains similar when the parameters are set to random values within a certain range. Thus, the properties of the model result from its general processing characteristics rather than the specific parameter values that were chosen.

Whenever incongruent stimuli are used in the simulations reported below, the same stimulus is presented on every trial, whichever the task (except where stated). A small set of stimuli is typically used for both tasks in experimental studies of task switching. Since the effects of learning only last for one trial in the model, it was necessary in the basic simulations to use the same stimulus on every trial, to simulate the possible contribution of item-specific effects which are assumed to have occurred in previous studies of task switching. This is of course a major simplification, which is addressed further in a later section.

In each of the following simulations, mean reaction times are based on 50,000 simulated reaction times. Trials where the model produced an incorrect response, or where reaction time was greater than 400 cycles (corre-

sponding to around 2500 ms according to linear regressions of simulated on empirical data) were counted as errors and excluded from the reaction time analyses. For every simulation below where the performance of the model is being compared with a specific set of empirical data, we have performed a linear regression of the simulation data (in cycles) on the empirical data (in milliseconds) and reported the model's simulations in milliseconds. Where the model is not being compared with any specific data set, we have used the regression equation derived from the model's simulation of the Stroop interference and facilitation data (see below).

Stroop Interference and Facilitation

The first aspect of the model's performance to be investigated is whether it can accurately predict the reaction times of subjects in pure blocks of word reading and color naming with neutral, incongruent, and congruent Stroop stimuli (see MacLeod, 1991, for a review). Cohen et al. (1990) demonstrated that their model was able to simulate three important findings in pure blocks of the Stroop tasks: word-reading reaction times are faster than color-naming reaction times; the word-reading task interferes more with the color-naming task than vice versa (i.e., there is a greater difference between neutral, incongruent, and congruent trials in the color-naming task); and the cost of incongruent trials, in comparison with neutral trials, is greater than the benefit of congruent trials. Pure blocks of word-reading and color-naming were simulated, with no preparation interval. As shown in Fig. 3, the model accurately

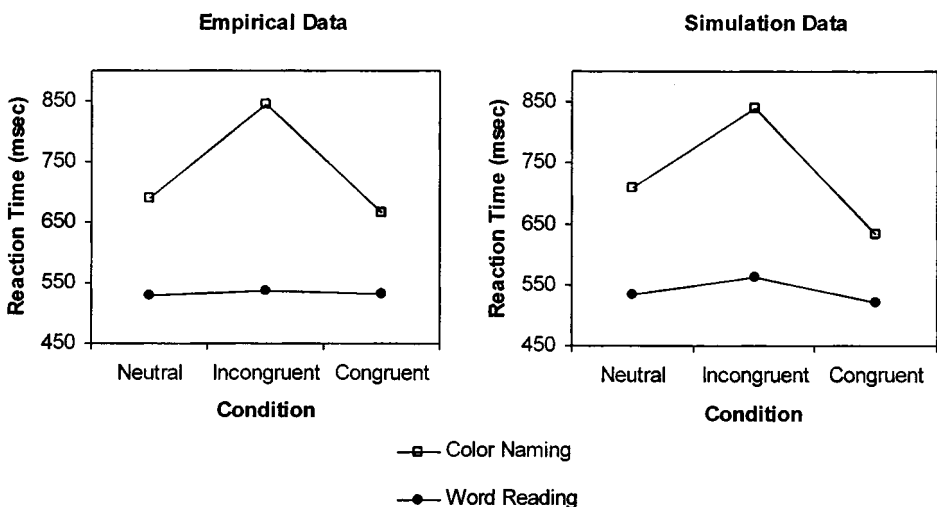


FIG. 3. Performance of the model in pure blocks of the standard Stroop task, with equivalent empirical data, based on Dunbar and MacLeod (1984, Experiment 1B). Simulated RTs = $5.8 * \text{cycles} + 318 \text{ ms}$.

reproduces all three aspects of the data which were discussed in relation to the Cohen et al. model. Errors were infrequent ($<0.5\%$ for every trial type).

Each of these phenomena can be readily explained in terms of the architecture and processing of the model. Word-reading responses are faster than color-naming responses because of the stronger connection strengths in the word pathway which lead to stronger activation of the output units and therefore a faster rate of evidence accumulation. The color-naming task suffers much greater interference from the word-reading task than vice versa because the connections in the word pathway are sufficiently strong to cause strong activation in the output units, even when the task demand units favor the color-naming task. The greater cost of incongruent trials, in comparison with the benefit of congruent trials, can be explained in terms of the activation function used [Eqs. (1) and (2)]. The effect of a positive net input into a unit is reduced, the higher the activation value of that unit (the same is true of the somewhat different activation function used by Cohen et al., 1990, and Cohen & Huston, 1994). Thus, the greater the positive input to a unit, the less will be the contribution of additional positive inputs. Consequently, adding an extra positive input to an active output unit from a congruent stimulus dimension will have less effect than adding a positive input to a less active (incorrect) output unit from an incongruent stimulus.


The simulation of the Stroop interference and facilitation data is not a demanding test of the present model which, despite several differences, is based on an existent model of the Stroop task (Cohen et al., 1990; Cohen & Huston, 1994). A more interesting question concerns the ability of the model to generalize beyond the mean reaction times found in pure blocks and predict reaction times when subjects switch between the two tasks within the same block.

Task Switching in Mixed Blocks

To test the ability of the model to predict reaction times in mixed blocks of word reading and color naming, the alternating runs paradigm (Rogers & Monsell, 1995) was simulated with a run length of four before each switch of task. If switch costs are found, this would then allow one to investigate whether they are confined to the first trial of a run, as the empirical data would suggest (Rogers & Monsell, 1995, Experiment 6). Again, there was no preparation interval in this simulation.

The mean reaction times of the model when it performs four trials of word reading, followed by four trials of color naming, and then a further four trials of word reading are illustrated in Figs. 4a (with item repetition allowed between the two tasks) and 4b (without item repetition). The model was reset after each 12-trial run, so that the mean reaction time on the first word-reading trial was not affected by a previous color-naming trial. Four features of the model's performance are noteworthy. First, the model does produce switch costs; that is, switch trials have longer reaction times than nonswitch

(a) with item repetition



(b) without item repetition




FIG. 4. Performance of the model in a block of 12 trials (a) using the same stimulus throughout and (b) without item repetition. Simulated RTs were calculated with the regression equation used in Fig. 3.

trials. Second, this switch cost appears to be confined to the first trial, as found by Rogers and Monsell (1995). Inspection of Fig. 4a shows that there is a very small increase in RT on the first repeat trial compared with the second: 19.7 ms for the word task and 6.3 ms for the color task. However, this effect is negligible in comparison with the RT drop from the first to the second trial: 273 ms for the word task and 123 ms for the color task.

The third noteworthy feature of Fig. 4, especially Fig. 4a, is that switch costs are markedly asymmetrical: much larger for the word than the color task. This is the pattern reported by Allport et al. (1994) and Allport and Wylie (2000). The switch cost asymmetry is less extreme in Fig. 4b, in line with some recent empirical evidence (Waszak et al., submitted); we return to this below. Finally, the model produces a slightly increased RT on the first trial of the run, even though this does not follow the performance of any other tasks. Such "restart costs" have been observed by Allport and Wylie (2000) and Gopher, Armony, and Greenshpan (2000).

Errors were relatively infrequent (<5% in all conditions), but were more common in the color task than in the word task and on switch in comparison with nonswitch trials. However, since reaction times have received greater attention than errors in studies of task switching, error rates are not discussed in the following simulations.

Explanations of the Observed Phenomena

Switch/restart costs. In order to understand why the model produces switch costs, we need to look at the role of the task demand units in its behavior. Recall that the task demand units bias processing in the word and color pathways so that the intended task is facilitated (i.e., receives positive connections) and the unintended task is inhibited (i.e., receives negative connections). Thus, the task demand units help to resolve competition between the two tasks, leading to one of the output units in the correct response dimension (color or word) becoming most active. The model was run without noise in order to produce Fig. 5, which illustrates the activation levels of the two task demand units over the course of a switch and a nonswitch trial, both for the color and the word task. The vertical lines indicate the moment of response execution, i.e., the moment when the difference between the activation levels of the most active and the second most active output unit passed the threshold. Although the activation levels after the end of the trial are shown in Fig. 5, for illustrative purposes, the next trial immediately follows response execution when the model is actually run.

On switch trials, the task demand units take longer to reach the activation levels required to facilitate the intended task and inhibit the unintended one. Thus, competition between the responses is extended, in comparison with nonswitch trials, because the task demand units are less effective at biasing the network toward the correct task. This is for two reasons. First, the activation values of the task demand units have the wrong sign at the beginning

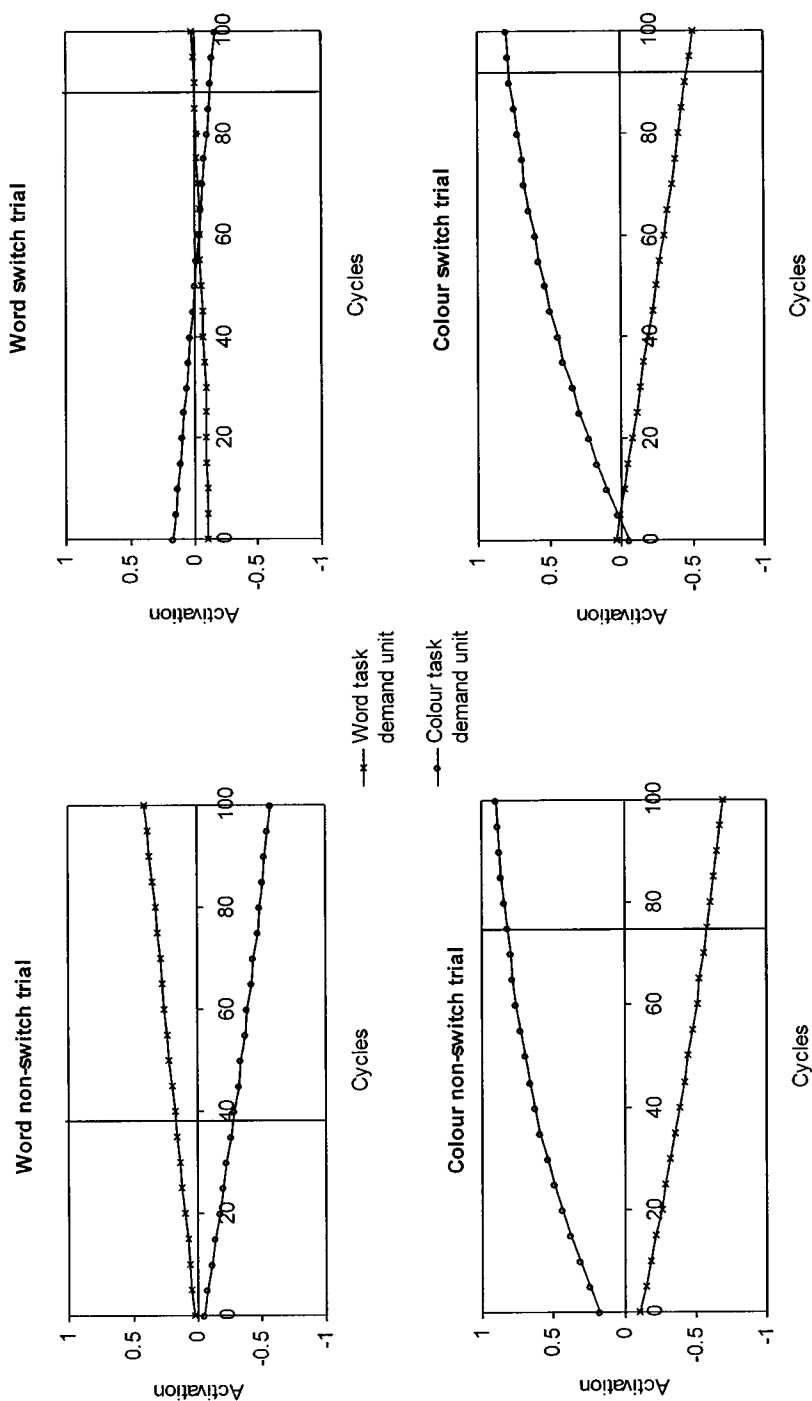


FIG. 5. Activation levels of the task demand units during the course of switch and nonswitch trials, for both the word and the color task. Vertical lines indicate the moment of response execution.

of the trial and so take longer to reach activation levels which would facilitate the intended task and inhibit the unintended one. This is caused by the carryover of the task demand units' activation values from the previous trial. A second reason for the relative ineffectiveness of the task demand units on switch trials results from the network's learning mechanism and the repetition of the same item for both tasks. This item-specific effect will be investigated in more detail below.

Restart costs arise because, on the first trial of a run, the network lacks the facilitation caused by a carryover of task set from switch to subsequent nonswitch trials. Since they only reflect the absence of a facilitatory effect, rather than being caused by an inhibitory effect combined with the carryover of a competing task set, restart costs are more modest than switch costs.

Switch costs confined to the first trial in a run. A determinant of the size of switch costs in the model is the state of the task demand units at the end of the trial preceding the switch. This determines the magnitude of the carryover of task set: if the task demand units are very active at the end of one trial, this will cause a large carryover into the subsequent trial. As Fig. 5 shows, although the task demand units have inappropriate activation levels at the start of switch trials, their activation levels become appropriate for the required task by the end of such trials. The stimulus input units become associated with these new values by Hebbian learning. Consequently, the appropriate task demand activations are carried over into the immediately successive nonswitch trial. In other words, the original cause of the switch cost (incorrect task demand unit activations) is dissipated by the end of the switch trial. It therefore does not contribute to successive nonswitch trials in the run. (At the end of a switch trial, the task demand units generally have a weaker activation than they do at the end of a nonswitch trial, since they started with inappropriate activations at the beginning of the trial. As a result, the first nonswitch trial after a switch may receive a weaker carryover of the appropriate task than subsequent nonswitch trials. This accounts for the very slight improvement which is observed between the first and second repeat trial.)

One might ask whether this account would be tenable in a model where there was long-term learning between the stimulus input and task demand units (to account for long-term item-specific priming effects). Long-term connections between stimulus input and task demand units may contribute to the switch cost, and these connections will not have dissipated by the end of the switch trial. However, it is assumed that the long-term connection weights will have a slow enough decay that the contribution of these connections will not be noticeably different from trial to trial. Thus, the effect of long-term connection weights between stimulus input and task demand units will effectively be constant over successive trials in a run and will not cause a gradual reduction in RT from trial to trial. Note that switch costs confined to the first trial in a run are not dependent on the model's learning rule, the effects of which last for just one trial. The same pattern of data is seen in

Fig. 4b, where there is no item repetition. Rogers and Monsell (1995) are therefore incorrect to argue that the task carryover account of switch costs is incompatible with such costs being limited to the first trial of a run. As Fig. 4 shows, a model in which the only source of switch costs is a carryover of task set is capable of producing this pattern of data.

Asymmetric switch costs. Examination of Fig. 5 also reveals a reason for asymmetric switch costs. Comparison of the word and color switch trials shows that, on the color switch trial, the task demand units quickly become strongly activated in the appropriate manner for the color-naming task but on the word switch trial the task demand units continue to favor the color task until just before the moment of response execution. This is due to the difference in the strength of top-down inputs for the two task demand units. Recall that a larger value is added to the net input of the color task demand unit on color trials than is added to the net input of the word task demand unit on word trials. Thus, on switch trials, reconfiguration of the task demand units for the intended task is slower for the word-reading task than the color-naming task, hence the paradoxical asymmetry in switch costs.

It should be noted that this explanation of asymmetric switch costs differs from the explanation proposed by Allport et al. (1994). Allport and colleagues propose that the color task requires strong suppression of word reading, but not vice versa. As a result, the word task is strongly suppressed following a switch from color naming. This generates a large switch cost in the word task, but there is a smaller cost for switches into the color task since it was not previously suppressed or was suppressed to a lesser degree. We agree that differences in “competitor suppression” between the two tasks are plausible (equal excitatory and inhibitory weights from the color and word task demand units were implemented as a simplifying assumption). However, the model shows that a paradoxical asymmetry in switch costs can result from differences in top-down inputs for the two tasks, regardless of their possible differences in the requirement for competitor suppression.

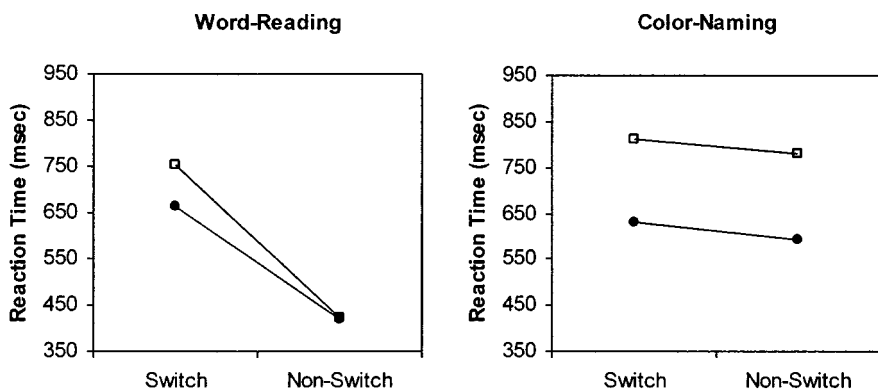
It would be problematic if the model rigidly predicted asymmetric switch costs, with a larger cost of switches into the dominant task. As discussed above, a number of experiments have been carried out which have yielded an asymmetry of switch costs in the reverse direction (i.e., greater switch costs for switches into the weaker task; see Monsell et al., 2000). However, we will see later that the asymmetry in switch costs observed in Fig. 4a and (to a lesser degree) Fig. 4b is sensitive to certain parameter manipulations. Thus, it is possible for the model to produce asymmetric switch costs with a larger cost for switches into the weaker task. The circumstances that give rise to this, both experimentally and in the model, are discussed below.

Reverse Stroop Interference

One of the pieces of evidence for the task carryover account of switch costs is the “reverse Stroop” interference seen on word switch trials (e.g., Allport et al., 1994; Allport & Wylie, 2000). On switch trials, word reading

is substantially slower for incongruent than neutral stimuli but the difference is generally negligible on nonswitch trials, suggesting a carryover of the color task set into word switch trials. This contrasts with color naming, which shows a large difference between incongruent and neutral stimuli for both switch and nonswitch trials. As can be seen from Fig. 6, the model also shows this effect. On word nonswitch trials, the color pathway is too weak to interfere noticeably with activity in the word pathway. As a result there is little difference in reaction time to neutral and incongruent stimuli on word

Empirical Data



Simulation Data

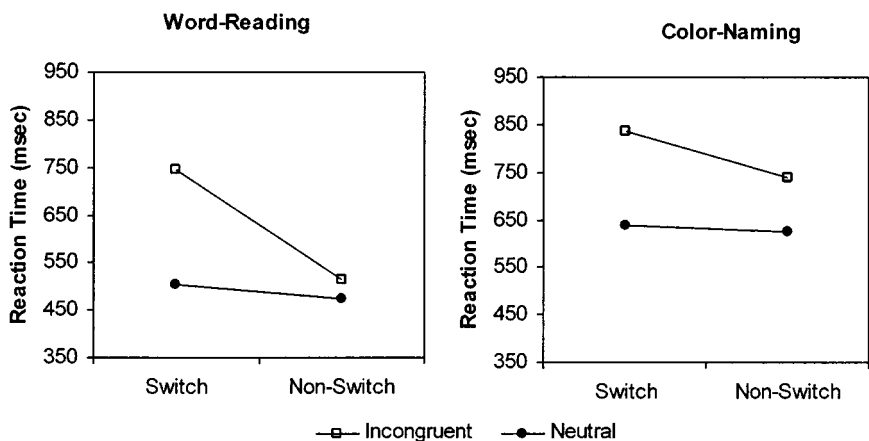


FIG. 6. Simulated and experimental effects of incongruent versus neutral stimuli on switch and nonswitch trials. Empirical data for word reading is based on Allport and Wylie (2000, Fig. 5). Data for color naming is based on Allport and Wylie (2000, Fig. 1). Simulated RTs = $4.9 * \text{cycles} + 288 \text{ ms}$.

nonswitch trials. But on word switch trials, the carryover of task set facilitates the color pathway and inhibits the word pathway. This causes the color pathway to be strong enough, and the word pathway to be weak enough, for interference to occur, hence the difference in reaction time between incongruent and neutral stimuli on switch trials. When the task is color naming, the word pathway is sufficiently strong to cause interference even on nonswitch trials when the color and word are incongruent. Thus, there is a large difference between reaction times to incongruent and neutral stimuli both on switch and nonswitch trials.

Item-Specific Switch Costs

In an experiment which was discussed in the introduction, Allport and Wylie (2000, Experiment 5) investigated the extent to which switch costs might be contingent on repetition of the same stimulus items between the two tasks. They found that the cost of a switch into the word task was substantially higher if the stimulus on the switch trial had earlier appeared in a color-naming trial. However, there was no significant difference in reaction time between these “primed” and “unprimed” stimuli on nonswitch trials. Thus, there appears to be a strongly item-specific component of the switch cost. It is clear that priming from the color task cannot affect subsequent word nonswitch reaction times in the model for the simple reason that learning between the stimulus input and task demand units only lasts for one trial. As a result, even if there is a greater reaction time on word switch trials for primed stimuli (i.e., stimuli just seen in the color-naming task), this priming will have decayed before the following word nonswitch trial.

This is not a satisfactory account of the interaction between stimulus repetition and switch/nonswitch trials because we have already seen evidence for priming effects that last for more than one trial (Allport & Wylie, 2000, Experiment 5; Waszak et al., submitted). The simplifying assumption that learning effects last for only one trial does not allow us to test the model’s simulation of the interaction between item repetition and switch/nonswitch trials. Thus, to provide a stronger and more principled test of the model, we allowed learning to last beyond a single trial for this simulation only.

Equation (3) was modified so that the weight at the end of the last trial was added to the new weight, i.e., there was no decay in the connection strengths between the stimulus input and task demand units. First, the model was presented with a sequence of four nonrepeating incongruent stimuli (referred to as “ABCD”), the first two for color naming and the last two for word reading, in order to obtain baseline reaction times for unprimed stimuli.² Then, after reinitializing the network, an “ABBC” sequence was pre-

² In order to obtain these unprimed baseline RTs, Hebbian weights were only allowed to persist for one trial (as in the earlier simulations), and different stimuli were presented on consecutive trials. This was necessary because there are only three possible incongruent stimuli with nonoverlapping stimulus representations.

sented, again with color naming for the first two trials and word reading for the last two. Finally, an "ABCB" sequence was presented. (As before, each of these simulations was run 50,000 times). By comparing the mean reaction time to the second B stimulus in the ABBC sequence with the C stimulus in the ABCD sequence, it is possible to assess the effect of item repetition on word switch trials. Likewise, comparing the mean reaction time to the second B stimulus in the ABCB sequence with the D stimulus in the ABCD sequence should reveal the effect of item repetition on word nonswitch trials. Since there was no decay in the connection strengths between the stimulus input and task demand units, the distance between prime and probe trials had no effect on the model's behavior. Figure 7 shows the results of these simulations.

The model does an excellent job of simulating the interaction between item repetition and trial type (switch/nonswitch) for word-reading reaction times. Not only does item repetition between the color and word tasks lead to slower response times in the word-reading task, this effect is essentially confined to switch trials. The explanation is as follows. When a stimulus is presented for color naming, the Hebbian learning at the end of the trial leads to a positive connection weights from the stimulus input units to the color task demand unit, and negative connection weights from the stimulus input units to the word task demand unit. If the same stimulus is presented on a subsequent wordreading trial, it sends some activation to the color task demand unit and inhibition to the word task demand unit. This means that the




FIG. 7. Simulated and experimental effects of item repetition across both tasks on word-reading switch and nonswitch trials. Empirical data are based on Allport and Wylie (2000, Fig. 6). Simulated RTs = $5.6 * \text{cycles} + 221 \text{ ms}$.

task demand units take longer to reach appropriate activation values for the word task, since the top-down activation they receive is countered, to some extent, by the inappropriate bottom-up (i.e., stimulus-driven) activation. On nonswitch trials, the task demand units start off favoring the word task and inhibiting the color task. Since the word pathway is much stronger than the color pathway, the precise values of the task demand units have very little effect on word-reading trials when they already favor the word task. Thus, an incorrect bottom-up activation of the task demand units on a word nonswitch trial has little effect on the model's behavior.

Similar simulations were run to assess the effects of priming from word-reading to color-naming trials; these revealed two differences between the two tasks. First, although the mean reaction time was slower for primed than unprimed color-naming switch trials, this effect was smaller than in the word-reading task. This is because the color-naming top-down input is stronger, allowing it to oppose the inappropriate bottom-up activation more effectively. Second, unlike in the word-reading task where the effect of item repetition is very small on nonswitch in comparison with switch trials, the relative size of the effect on color-naming nonswitch trials, in comparison with switch trials, was greater than half. This is because the color-naming pathway, unlike the word-reading pathway, is heavily dependent on the task demand units even on nonswitch trials. Support for these two predictions was recently obtained, using a picture-word Stroop task, by Waszak et al. (submitted), who found that the effect of item repetition on reaction times in the less dominant picture-naming task was smaller than the effect on word-reading switch trials, and it was just as large on switch and nonswitch trials. Note that the model's simulation of item repetition increases the switch cost substantially in the word-reading task, but very little in the color-naming task, where switch and nonswitch trials are affected similarly. This accounts for the greater asymmetry in switch costs between the two tasks when item repetition is allowed (compare Figs. 4a and 4b), as found by Waszak et al. (submitted).

Now that we have seen the effect of item repetition on word switch and nonswitch trials, another explanation of reverse Stroop interference is possible. Recall that incongruent word switch trials are slower than neutral trials, but there is no difference in reaction time to incongruent and neutral stimuli on word nonswitch trials. By definition, a neutral stimulus in one task cannot appear as a stimulus in the other task. Thus, for neutral trials, there was no item repetition between the two tasks. However, in the above simulation, the same incongruent stimulus was used in both tasks. In other words, the incongruent stimulus on the word switch trial was primed, but the neutral stimulus was not. Reinterpreted in this way, the reverse Stroop data shows the same pattern of results as the itemspecific data: a difference between primed and unprimed stimuli on switch, but not nonswitch trials (at least in the word task). In order to assess the contribution of item-specific priming

effects to reverse Stroop interference, we reran the simulations illustrated in Fig. 6, this time using a different incongruent stimulus for the word and color tasks. Essentially the same pattern of results was found, but with a smaller difference between incongruent and neutral stimuli on the word switch trial. Thus, item-specific priming effects can play a role in the model's simulation of reverse Stroop interference, but such interference is seen even in the absence of any item-specific effects.

Effects of Neutral Trials

Many experiments (e.g., Allport & Wylie, 2000, Experiment 4; Rogers & Monsell, 1995, Experiment 1) have found that switch costs are reduced, but not eliminated, when neutral stimuli are used. Inspection of Fig. 6 shows that the model also produces reduced switch costs when the stimulus is neutral, both in the word and color tasks. This can be explained in terms of the earlier explanations of reverse Stroop interference. Incongruent switch trials are slower than neutral switch trials, but there is a smaller difference in reaction time between incongruent and neutral nonswitch trials. Thus, the difference in reaction time between switch and nonswitch trials (i.e., the switch cost) is larger for incongruent than neutral stimuli. This is because response selection in incongruent trials is particularly slow following a switch of task.

A second example of the effect of neutral trials on switch costs is reported by Allport and Wylie (2000, Experiment 1; also reported in Wylie & Allport, 2000, Experiment 1). In this experiment, the focus was the effect of neutral trials in the previous task on subsequent switch costs rather than the effect of neutral stimuli on switch trials. Allport and Wylie found that the cost of a switch into the word-reading task was greatly increased when the stimuli in the previous color-naming task were changed from neutral to incongruent, even though the word-reading stimuli were always incongruent. Similarly, the model's mean word-reading switch cost increases from approximately 151 to 273 ms when the stimulus for the previous color-naming trial is changed from neutral to incongruent, using incongruent stimuli for word reading in both cases.

The main cause of this effect is the absence of item repetition between the two tasks when a neutral stimulus is used, as discussed above, as a possible explanation of reverse Stroop interference. Another contributing factor to the reduced switch cost following the appearance of neutral stimuli in previous trials is that, on neutral trials, a response can be selected before the task demand units have become strongly active. This means that the task demand units have smaller activation values at the end of neutral than incongruent trials, leading to a greater carryover of task set into switch trials following incongruent stimuli in the other task. Thus, even when items are not repeated between the two tasks, the model produces a larger switch cost following incongruent stimuli.

Preparation

Preparation is simulated by activating the task demand units in advance of the stimulus input units, without activating the output units. As Fig. 8 shows, the preparation interval has a large facilitatory effect on the reaction times for switch trials, but little effect on nonswitch RTs: the preparation interval is therefore able to reduce switch costs, as found by Meiran (1996) and Rogers and Monsell (1995, Experiment 3). In this respect, the model's simulation of the empirical data is successful. However, one aspect of the model's behavior does not fit the experimental evidence. In the model, the switch cost is entirely eliminated at the longest preparation interval of 150 cycles; this corresponds with an interval of approximately 1200 ms according to the regression equation derived from the model's simulation of the Stroop interference and facilitation data. In contrast, some experimental investigations into task switching have found residual switch costs even at preparation intervals as long as 8 s (Kimberg, Aguirre, & D'Esposito, 2000). This forms a central plank in the evidence for the exogenous control process account of switch costs (Rogers & Monsell, 1995).

However, other studies have found that the switch cost is virtually eliminated by the preparation interval (De Jong et al., submitted; Meiran, 1996; Meiran et al., 2000; Tornay & Milán, 2001). On the basis of this and other evidence, De Jong (2000) has argued that the residual switch cost does not result from some inexorable limitation in subjects' ability to prepare for a switch of tasks, but rather from their occasional failure to make effective use of the preparation interval. This is perfectly compatible with the model presented here, although it is not necessary to assume, as does De Jong, that subjects are either fully prepared or not at all prepared for the new task at

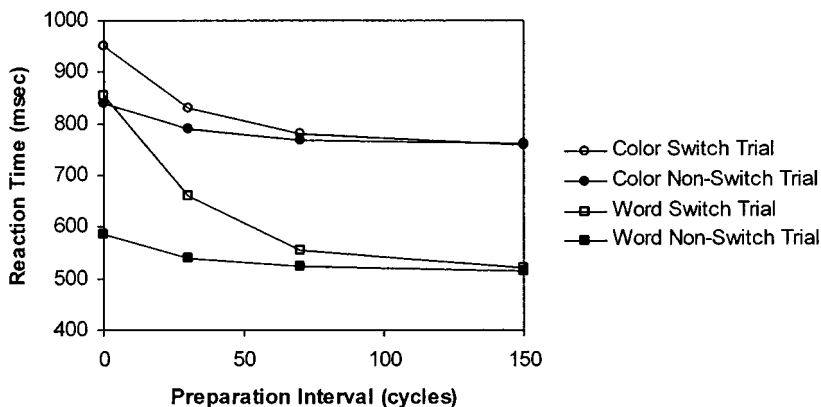


FIG. 8. Simulated effects of the preparation interval on reaction time. Simulated RTs were calculated with the regression equation used in Fig. 3.

the moment of stimulus onset. Thus, the model's failure to produce a residual switch cost may result from the way that the preparation interval is simulated rather than reflecting a more serious problem.

Another potential shortcoming of the model's simulation of preparation is that some authors have reported that preparation, while reducing switch costs, does not seem to reduce the amount of Stroop-like interference between tasks (e.g., De Jong, 1995; Meiran, 1996). The model would predict some reduction in Stroop interference with preparation. However, this reduction in Stroop-like interference with preparation has been found by other researchers (Goschke, 2000; Sudevan & Taylor, 1987), so it is not clear whether this should be considered a problem for the model.

EFFECTS OF PARAMETER SETTINGS

The model produces a good fit to a large body of empirical data from task-switching experiments. However, it has a number of free parameters which were set by hand. It is possible that the model owes its successes to the specific parameter settings chosen. If this is the case, then the model cannot be said to provide strong support for the task carryover account of switch costs, since it might be able to fit any set of data with the appropriate parameter settings (see Roberts & Pashler, 2000). It would be more impressive if it could be shown that the properties of the model result from the general principles which govern its operation rather than its specific parameters. We investigated this question by evaluating the effects of setting the parameters to random values (within certain limits).

The model was run 1000 times with a random value selected for each weight and each other free parameter on every run, with the exception of minimum and maximum activation values and step size (see the Appendix for a list of parameters and their possible values). Each run consisted of 500 simulated reaction times in each condition to provide mean reaction times. In keeping with a central assumption of the model, the connections from stimulus input to output units had a greater strength in the word pathway than in the color pathway, reflecting the greater strength of the word task (Cohen et al., 1990). On each run, two random values were chosen for the connection strengths from the stimulus input units to the word and color output units. If the value for the color units was greater than the value for the word units, another pair of values was chosen until the word output units received a stronger connection than the color output units. Similarly, the level of top-down control input for the color task demand unit was greater than the level of top-down control input for the word task demand unit.

Before these simulations could be run, it was necessary to establish limits for the random values chosen for each parameter. This was done by taking the model with the standard parameter settings and adjusting each parameter in turn, to establish its maximum and minimum values, outside of which the

model produced unacceptably high error rates. We defined “unacceptably high” as an error rate of 20% or more for one or more trial types in a repeating run of word-reading and color-naming trials using the same incongruent stimulus on every trial.

The level of top-down control input to the color task demand unit had no upper limit above which the model’s performance deteriorated. It was therefore given a maximum value which made the standard parameter setting exactly midway between the maximum and the minimum values. The noise parameter was allowed to vary between zero and its standard setting. We assumed that, since the maximum and minimum parameter values were chosen with the noise parameter at its standard setting, making this the maximum noise setting should ensure that the majority of simulations do not produce an excessive number of errors.

Results

Models that produced an error rate of greater than 10% on any trial type were excluded from the following analyses (as before, reaction times of greater than 400 cycles were counted as errors). This led to 366 simulations of 1000 being excluded (see Table 1).

TABLE 1

Thirteen Properties of the Model and the Percentage of Models Showing Those Properties When Their Parameters Are Set to Random Values

	Property	Percentage of models
A1.	Word reading is faster than color naming in pure blocks	88.5
A2.	In pure blocks there is a greater effect of incongruent and congruent stimuli in the color task than in the word task	76.3
A3.	In pure blocks, interference from incongruent stimuli is greater than facilitation from congruent stimuli	66.9
B1.	Switch trials are slower than repeat trials	98.9
B2.	Switch costs are confined to the first trial in a run	72.7
B3.	There is a greater cost of switches into the word task than switches into the color task	77.1
B4.	Reverse Stroop interference occurs on word switch trials	100.0
B5.	Less reverse Stroop interference occurs on word nonswitch trials	99.5
B6.	Reaction time is slower on word switch trials when the stimulus was seen on the previous color-naming trial	98.3
B7.	There is a smaller effect of item repetition between the two tasks, or no effect, on word nonswitch trials (for those models with property B6)	93.9
B8.	Smaller switch costs occur when the subject is presented with neutral stimuli than when incongruent stimuli are used	99.2
B9.	Switch costs in the word task, using incongruent stimuli, are greater following color naming with incongruent stimuli than following color naming with neutral stimuli	99.8
B10.	A reduction in switch costs occurs with the preparation interval	99.7

We have taken 13 empirical phenomena simulated above and categorized them into those that relate to performance in pure blocks of the Stroop tasks (A1–A3) and those that relate to task-switching in mixed blocks (B1–B10). Table 1 shows the percentage of simulations with each of these properties. Operationally, property B2 (switch costs confined to the first trial in a run) was defined as a reduction in RT between the first and the second nonswitch trial that amounted to less than 6.5% of the reduction in RT between the switch trial and the second nonswitch trial, averaged over both tasks (this is approximately the reduction shown by the model with its standard parameter settings). Property B10 was satisfied if switch costs were reduced by a 30-cycle preparation interval.

Binomial tests showed that the models with random parameter settings were more likely have each of these properties than not ($p < .0001$ for every phenomenon). Thus, these phenomena are simulated by the model because its architecture and general processing principles are predisposed toward them rather than as a result of manipulation of the parameter settings.³

Only two of the task-switching phenomena are shown by less than 90% of the models—switch costs confined to the first trial of a run (property B2, shown by 73% of the models) and paradoxical asymmetric switch costs (property B3, shown by 77% of the models). Interestingly, both of these phenomena have not always been reported in the empirical literature (see Meiran et al., 2000, and Salthouse et al., 1998, for evidence that switch costs need not be confined to the first trial of a run; and Monsell et al., 2000, for evidence that switch costs need not be “paradoxically” asymmetrical). Thus, the phenomena that the models with random parameter settings are least likely to produce are also those which have not always been observed in the empirical data. In order to further investigate the model’s simulation of switch costs being confined to the first trial in a run, we generated a frequency histogram of the drop in reaction time from the first to the second nonswitch trial, as a percentage of the drop in reaction time from the switch trial to the second nonswitch trial (Fig. 9). As the figure shows, the majority of simulations were in the bars representing 0–3% and 3–6%, indicating that there was generally little or no further drop in reaction time from the first to the

³ Only a subset of the parameter space was tested, in accordance with the assumption that the bottom-up connections should be stronger in the word than the color pathway and the top-down input stronger for the color than the word task. It is also important to show that the model fails when these constraints are not met. This can be seen when one considers that the only difference between the word and color pathways (i.e., the reason for one being called “word” and not the other) is the strength of the bottom-up and top-down connections. Thus, if the constraints concerning the strength of bottom-up and top-down connections were reversed, this would be equivalent to renaming the word task “color” and the color task “word” in the present simulations. In this case, the three phenomena that reflect a direct comparison between the two tasks (A1, A2, and B3) would not be shown by the majority of simulations currently showing them.

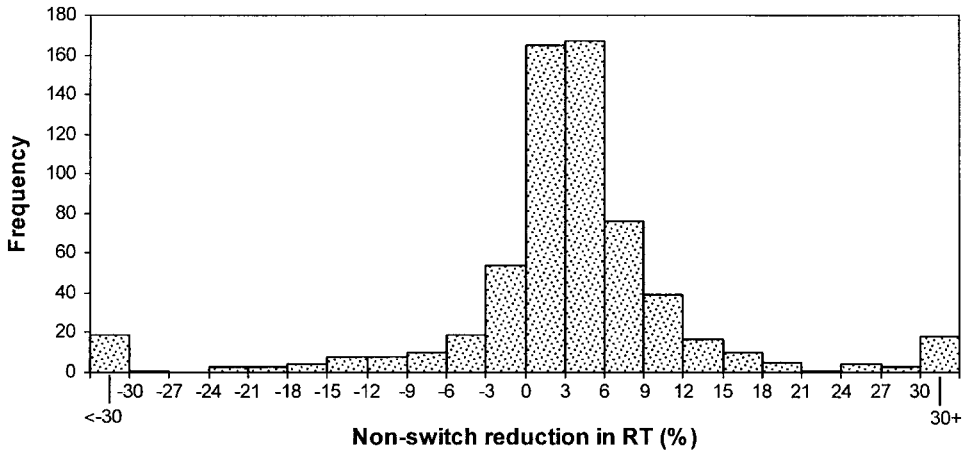


FIG. 9. Frequency histogram of the reduction in reaction time from the first to the second nonswitch trial, as a percentage of the reduction from the switch trial to the second nonswitch trial.

second nonswitch trial (as found by Rogers & Monsell, 1995). However, there were some simulations that showed a larger drop in reaction time between nonswitch trials, and a few that showed an *increase* in reaction time from the first to the second nonswitch trial.

Parameter Settings and Asymmetric Switch Costs

Although we have seen that the model does not rigidly predict a paradoxical asymmetry of switch costs when the parameters are set to random values, we have not yet seen which parameter manipulations are responsible for this variability in behavior. Only four of the model's parameters have effects which are specific to one of the two tasks: color and word pathway connection strengths and color and word level of top-down control input. Thus, any differences between the model's performance in the two tasks are due to differences between the settings of these parameters for the color and word tasks. As an important assumption of the model, the word pathway has greater connection strengths than the color pathway and the color top-down control input is greater than the word top-down control input. We investigated the effects of these differences on asymmetric switch costs. Figure 10 shows the mean *difference* between the pathway strengths (favoring the word task) and between the top-down control inputs (favoring the color task) separately for models which produce paradoxical and nonparadoxical asymmetric switch costs. The "paradoxical models" show a greater difference in top-down level of task control than the "nonparadoxical models," but the nonparadoxical models show a greater difference in pathway strength than the paradoxical models. Thus, nonparadoxical asymmetric switch costs seem to

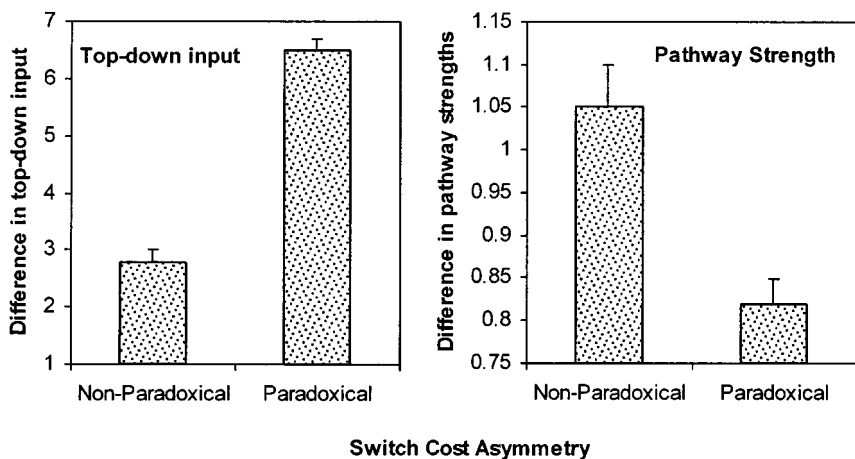


FIG. 10. Differences between the color and word tasks in top-down control input and pathway strength for models showing a paradoxical and a nonparadoxical asymmetry of switch costs. Error bars represent standard error of the mean.

be produced by the combination of a relatively small difference in top-down control input for the two tasks and a relatively large difference in pathway strength.

The effect of top-down control input on asymmetric switch costs makes sense in terms of the earlier explanation of asymmetric switch costs. Since switch trials are slowed by the need to reconfigure the state of the task demand units, a smaller top-down control input will lead to slower reconfiguration of the task demand units and therefore larger switch costs. Thus, a large difference in the level of top-down control input for the two tasks, favoring the color task, will lead to larger switch costs in the word than the color task. But this is clearly not the full explanation of asymmetric switch costs. If the only reason for asymmetric switch costs were the difference in top-down control inputs for the two tasks, we would expect the paradoxical asymmetry every time, since the top-down control input is *always* larger for the color task.

When the top-down control input for the two tasks is equal, but the pathway strengths are different, the model still produces asymmetric switch costs. This time, however, the switch cost is larger for switches into the nondominant task. Thus, two opposing effects seem to be occurring. When the level of top-down control input for two tasks is different, all else being equal, larger switch costs are produced in the dominant task (i.e., the task with the smaller top-down control input). But different pathway strengths (all else being equal) will produce larger switch costs in the nondominant task (i.e., the task with the weaker pathway strength). This is because the nondominant task is more dependent on the task demand units than the dominant task.

Thus, the effect of an inappropriate state of activation in the task demand units (the cause of switch costs in the model) is greater for the nondominant than the dominant task. It appears that the balance between these two forms of asymmetry in task strength is largely responsible for the nature of the asymmetric switch costs.

An experiment by Yeung (1999; also described in Monsell et al., 2000) has provided results in line with the model's prediction that nonparadoxical asymmetric switch costs should be associated with a small difference in top-down control input between the two tasks. In this experiment, subjects again switched between word-reading and color-naming tasks. A black color-word was displayed on a colored rectangle, following the procedure of Glaser and Glaser (1982). When the word and colored background were presented simultaneously on each trial, a paradoxical asymmetry was obtained, with larger switch costs for the word task. But when the word was presented 320 ms after the appearance of the colored background on each trial, the asymmetry in switch costs was reversed. Yeung (1999) suggests that this procedure reduces the difference in top-down control input between the two tasks, leading to a nonparadoxical asymmetry of switch costs. This is exactly in line with the predictions of the model.

DISCUSSION

We have demonstrated that an implementation of one theory of the causes of switch costs—the task carryover account—is able to reproduce a large number of empirical phenomena that have been reported in the literature. The model simulates the facilitation and interference effects found in Stroop color naming and word reading, with asymmetric interference between the two tasks and greater interference than facilitation, the existence of switch costs, the asymmetry of those switch costs between dominant and nondominant tasks, the confinement of those switch costs to the first trial in a run, reverse Stroop interference and its confinement to switch trials, the effects of neutral stimuli on switch costs when they appear in the task being switched from and when they appear in the task being switched into, item-specific priming effects, and the interaction of item repetition with switch versus nonswitch trials in the word task. The model also simulates the reduction of switch costs by the preparation interval, although it does not produce residual switch costs at long preparation intervals. Two of these effects have been claimed to be problematic for theories such as Allport and Wylie's to explain, viz., switch costs confined to the first trial of a run (see Rogers & Monsell, 1995, Experiment 6) and nonparadoxical asymmetric switch costs (see Monsell et al., 2000). The model has these properties as a result of its architecture and general processing principles rather than as a result of the specific parameter values that were chosen.

Comparison with Other Models

Norman and Shallice (1986). Norman and Shallice (1986) have also produced a model of the processes involved in the selection of routine and non-routine actions (see also Cooper & Shallice, 2000; Shallice, 1982, 1988, 1994). Although it was not designed with the task-switching paradigm in mind, the Norman–Shallice model has a great deal in common with the model presented in this article. According to the Norman–Shallice model, behavior is controlled by “schemata,” which select and coordinate the elementary processes involved in simple cognitive tasks. Individual schemata can become selected in two ways: They can be triggered by environmental cues, following which a process called “contention scheduling” selects the most active schema for the control of action. This is appropriate in routine circumstances where environmental cues are sufficient to determine behavioral responses. But in circumstances requiring the overriding of prepotent responses to environmental cues, or in novel situations where there are no learned associations between environmental cues and behavioral responses available, the “supervisory attentional system” (SAS) biases the contention scheduling process by providing a top-down input favoring the appropriate schema (and/or inhibiting unwanted schemata). The elements of this theory can easily be mapped onto elements of the model presented in this article: Schemata correspond to the task demand units, the environmental triggering of schemata corresponds with the connections that form between the stimulus input units and the task demand units, and the SAS corresponds with the top-down control input into the task demand units.

However, one apparent discrepancy is that, whereas in the standard version of the theory the SAS is only operative when nonroutine behavior is required, the top-down control input into the task demand units is always present. Thus, the model presented in this article does not entail a strict division between routine behavior mediated by contention scheduling and non-routine behavior which also requires the SAS. The model is equally consistent with other models where there is a single action-controlling hierarchy (e.g., Schwartz, Reed, Montgomery, Palmer, & Mayer, 1991)—the top-down control input into the task demand units may reflect an input from higher level to lower level schemata.

Extra process accounts of switch costs. As discussed in the introduction, some researchers have explained switch costs in terms of the duration of a stage-like executive control process that operates only on switch trials. Versions of this account have been proposed by De Jong (2000), Monsell et al. (1998, 2000), Rogers and Monsell (1995), Rubinstein et al. (2001), and others. We concur with Rogers and Monsell (1995; Monsell, 1996) that stimulus-driven processes play an important role in determining switch costs. However, the present model indicates that the assumption of an extra process only present on switch trials is not necessary to account for the findings

reported so far. The crucial difference between the model presented in this article and the “extra process” models discussed above is that, in the present model, precisely the same computational processes occur on switch and non-switch trials. The only difference between the two types of trial is in the initial state of the network. Thus, the switch cost does not measure the duration of any distinct process.

Note that it is quite possible that, at least on some occasions, additional processes may take place on switch trials. For example, a time-consuming process of retrieving the now-appropriate task from memory may occur more often on switch trials in some circumstances (Mayr & Kliegl, 2000). Indeed, one could examine the behavior of the model when such additional assumptions are made. However, we have demonstrated above that such extra processes on switch trials are not necessary to explain the switch cost.

Yeung (1999). Yeung (1999) has recently put forward a task carryover model of switch costs. Although very different in implementation, this model is very similar to the present one at a deeper theoretical level.⁴ As in the current model, Yeung’s model simulates switch costs as an extended response selection process resulting from a carryover of task set, although Yeung’s model incorporates separate response activation and response resolution stages rather than the single response selection stage that takes place in our model. Yeung’s model, like ours, explains paradoxical asymmetric switch costs in terms of a greater top-down control input for the nondominant task. Unlike our model, however, the carryover of task set is identical every time the subject switches task (it is a constant added to the “readiness” of the most recently performed task). One further difference is that there is no mechanism for stimulus-driven retrieval of task set in this model, although it could presumably be accommodated by a small extension of the model. Yeung’s model is implemented as a set of mathematical equations relating reaction time to theoretical variables such as “task readiness.” These equations are intended to be “descriptive rather than explanatory with respect to the nature of the response selection process.” By implementing our model in a PDP network, we aim to make the computational principles that may underlie the behavioral effects of task-switching more explicit.

Concluding Comments

The model presented in this article implements the idea that switch costs reflect an interference effect caused by a carryover of task set. We have shown that this account can provide explanations for a large body of empirical data. In general, supporters of alternative explanations of switch costs have not provided quantitative predictions for this empirical data and so it is difficult to evaluate the ability of other theories to explain the effects simulated by the model. One exception to this is the model put forward by Kieras,

⁴ The development of the present model and Yeung’s model occurred independently

Meyer, Ballas, and Lauber (2000). Kieras et al. have implemented the extra process account proposed by Rubinstein et al. (2001), using the EPIC architecture (see, e.g., Meyer & Kieras, 1997, 1999). However, this model has so far only been used to simulate a small set of data, from an unpublished study by Lauber (1995), making direct comparison with our model difficult. Meiran (2000a) has also put forward an explicit computational model of task switching. Again, the nature of the model and the data it is used to simulate make it difficult to compare our models directly.

It might seem that the task switching paradigm has less relevance to executive control if the switch cost simply reflects an interference effect caused by a carryover of task set. However, we have two suggestions for the role that the task switching paradigm can play in furthering our understanding of executive control. First, even if switch costs result from an automatic process of response conflict resolution, according to the present framework they can still reflect states of executive control, i.e., the top-down control inputs that both tasks receive. Thus, it should be possible to study these top-down control inputs indirectly, using the model to infer top-down input settings from behavioral data. We have recently utilized this methodology (Gilbert & Shallice, unpublished data; see Yeung, 1999 for a similar approach): We reasoned that if subjects were to perform a working memory task at the same time as switching between previously learned tasks of Stroop word reading and color naming, this should interfere more strongly with the (larger) color-naming top-down input than with the word-reading top-down input. In order to simulate the effect of concurrent performance of a working memory task, we therefore reduced the color-naming top-down input to a value closer to the weaker word-reading top-down input. When the color-naming top-down input is strongly reduced, the model produces a reversed asymmetry of switch costs (greater for the color-naming task), and this is also what we found in our subjects. A second role for the task switching paradigm in the study of executive control could occur once the effects of a carryover of task set and stimulus-driven retrieval of task set are understood. It would then be possible to attribute new findings in more complex situations to the duration of executive processes involved in task switching. In this respect, the current model could be used in the same way as it itself uses the earlier model of Cohen et al. (1990).

APPENDIX

Parameters of the Model, Standard Settings, and Permitted Range When Set to Random Values

Parameter	Standard setting	Possible range for random values
Weights		
Stimulus input units to word output units	3.5	1.9–4.2
Stimulus input units to color output units	1.9	1.5–4.1
Task demand units to word/color output units ^a	2.5	1.6–5.7
Word/color output units to task demand units	1.0	0.0–6.7
Connections between color and word output units ^b	2.0	0.0–2.8
Within module inhibitory connections ^c	-2.0	-5.0–0.0
Other parameters		
Bias for word/color output units	-6.0	-11.3–-3.5
Bias for task demand units	-4.0	-15.2–0.0
Top-down control input into color task demand unit	15.0	9.0–21.0
Top-down control input into word task demand unit	6.0	3.2–21.0
Learning rate between stimulus input and task demand units	1.0	0.0–1.6
Response threshold	0.15	0.09–0.32
Minimum activation value	-1.0	-1.0
Maximum activation value	1.0	1.0
Step size [see Eq. (1)]	0.0015	0.0015
“Squashing” of task demand unit activations between trials	80%	63–100%
Noise (i.e., standard deviation of the distribution used to provide noise terms added to activation values)	0.006	0–0.006

^a This parameter determines both the positive connection strength from (e.g.) the word task demand unit to each of the word output units and the negative connection strength from (e.g.) the word task demand unit to each of the color output units. Likewise, the positive and negative connection strengths for the backprojections from the color and word output units to the task demand units are determined by a single parameter.

^b This parameter determines both the positive connection strengths between congruent words and colors and the negative connection strengths between incongruent words and colors.

^c This parameter determines the negative connection strengths between each member of the task demand, word, and color modules.

REFERENCES

- Allport, D. A., Styles, E. A., & Hsieh, S. (1994). Shifting intentional set: Exploring the dynamic control of tasks. In C. Umiltà & M. Moscovitch (Eds.), *Attention and performance XV* (pp. 421–452). Cambridge, MA: MIT Press.
- Allport, A., & Wylie, G. (2000). ‘Task-switching’, stimulus-response bindings, and negative priming. In S. Monsell & J. S. Driver (Eds.), *Attention and performance XVIII: Control of cognitive processes* (pp. 35–70). Cambridge, MA: MIT Press.
- Asaad, W. F., Rainer, G., & Miller, E. K. (2000). Task-specific neural activity in the primate prefrontal cortex. *Journal of Neurophysiology*, **84**, 451–459.

- Atkinson, R. C., & Shiffrin, R. M. (1968). Human memory: A proposed system and its control processes. In K. W. Spence & J. T. Spence (Eds.), *The psychology of learning and motivation* (Vol. 2). New York: Academic Press.
- Bundesen, C. (1990). A theory of visual attention. *Psychological Review*, **97**, 523–547.
- Burgess, N. (1995). A solvable connectionist model of immediate recall of ordered lists. In D. S. Touretzky, G. Tesauro, & T. K. Leen (Eds.), *Advances in neural information processing systems* (pp. 51–58). Cambridge, MA: MIT Press.
- Cohen, J. D., Braver, T. S., & O'Reilly, R. C. (1996). A computational approach to prefrontal cortex, cognitive control and schizophrenia: Recent development and current challenges. *Philosophical Transactions of the Royal Society B*, **351**, 1515–1527.
- 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**, 332–361.
- Cohen, J. D., & Huston, T. A. (1994). Progress in the use of interactive models for understanding attention and performance. In C. Umiltà & M. Moscovitch (Eds.), *Attention and performance XV* (pp. 453–476). Cambridge, MA: MIT Press.
- Cohen, J. D., & Servan-Schreiber, D. (1992). Context, cortex and dopamine: A connectionist approach to behavior and biology in schizophrenia. *Psychological Review*, **99**, 45–77.
- Cohen, J. D., Usher, M., & McClelland, J. L. (1998). A PDP approach to set size effects within the Stroop task: Reply to Kanne, Balota, Spieler, and Faust (1998). *Psychological Review*, **105**, 188–194.
- Cooper, R., & Shallice, T. (2000). Contention scheduling and the control of routine activities. *Cognitive Neuropsychology*, **17**, 297–338.
- Dayan, P. (1998). A hierarchical model of binocular rivalry. *Neural Computation*, **10**, 1119–1136.
- De Jong, R. (1995). Strategic determinants of compatibility effects with task uncertainty. *Acta Psychologica*, **88**, 187–207.
- De Jong, R. (2000). An intention-activation account of residual switch costs. In S. Monsell & J. S. Driver (Eds.), *Attention and performance XVIII: Control of cognitive processes* (pp. 357–376). Cambridge, MA: MIT Press.
- De Jong, R., Emans, B., Eenshuistra, R., & Wagenmakers, E.-J. (submitted). Structural and strategic determinants of intentional task control.
- Dunbar, K., & MacLeod, C. (1984). A horse race of a different color: Stroop interference patterns with transformed words. *Journal of Experimental Psychology: Human Perception and Performance*, **10**, 622–639.
- Fuster, J. M. (1997). *The prefrontal cortex*, 3rd ed. Philadelphia: Lippincott–Raven.
- Glaser, M. O., & Glaser, W. R. (1982). Time course analysis of the Stroop phenomenon. *Journal of Experimental Psychology: Human Perception and Performance*, **8**, 875–894.
- Gopher, D., Armony, L., & Greenspan, Y. (2000). Switching tasks and attention policies. *Journal of Experimental Psychology: General*, **129**, 308–339.
- Gopher, D., Weil, M., & Siegel, D. (1989). Practice under changing priorities: An approach to training of complex skills. *Acta Psychologica*, **71**, 147–179.
- Goschke, T. (2000). Intentional reconfiguration and involuntary persistence in task-set switching. In S. Monsell & J. S. Driver (Eds.), *Attention and performance XVIII: Control of cognitive processes* (pp. 331–355). Cambridge, MA: MIT Press.
- Jersild, A. T. (1927). Mental set and shift. *Archives of Psychology, Whole No. 89*.
- Kanne, S. M., Balota, D. A., Spieler, D. H., & Faust, M. E. (1998). Explorations of Cohen,

- Dunbar, and McClelland's (1990) Connectionist Model of Stroop Performance. *Psychological Review*, **105**, 174–187.
- Kieras, D. E., Meyer, D. E., Ballas, J. A., & Lauber, E. J. (2000). Modern computational perspectives on executive mental processes and cognitive control: Where to from here? In S. Monsell & J. Driver (Eds.), *Attention and Performance XVIII: Control of cognitive processes*. Cambridge, MA: MIT Press.
- Kimberg, D. Y., Aguirre, G. K., & D'Esposito, M. (2000). *Neural activity associated with task-switching: An fMRI study*. Poster presented at meeting of the Cognitive Neuroscience Society, San Francisco, April 2000.
- Lauber, E. J. (1995). *Executive control of task switching operations*. Unpublished doctoral dissertation, University of Michigan, Ann Arbor, MI.
- Logan, G. D. (1979). On the use of a concurrent memory load to measure attention and automaticity. *Journal of Experimental Psychology: Human Perception and Performance*, **5**, 189–207.
- Logan, G. D. (1980). Attention and automaticity in Stroop and priming tasks: Theory and data. *Cognitive Psychology*, **12**, 523–553.
- Logan, G. D. (1985). Executive control of thought and action. *Acta Psychologica*, **60**, 193–210.
- Logan, G. D., & Gordon, R. D. (2001). Executive control of visual attention in dual-task situations. *Psychological Review*, **108**, 393–434.
- MacLeod, C. M. (1991). Half a century of research on the Stroop effect: An integrative review. *Psychological Bulletin*, **109**, 163–203.
- Mayr, U., & Keele, S. W. (2000). Changing internal constraints on action: The role of backward inhibition. *Journal of Experimental Psychology: General*, **129**, 4–26.
- Mayr, U., & Kliegl, R. (2000). Task-set switching and long-term memory retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **26**, 1124–1140.
- McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review*, **88**, 159–188.
- Meiran, N. (1996). Reconfiguration of processing mode prior to task performance. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **22**, 1423–1442.
- Meiran, N. (2000a). Modeling cognitive control in task-switching. *Psychological Research*, **63**, 234–249.
- Meiran, N. (2000b). Reconfiguration of stimulus task sets and response task sets during task switching. In S. Monsell & J. S. Driver (Eds.), *Attention and performance XVIII: Control of cognitive processes* (pp. 377–399). Cambridge, MA: MIT Press.
- Meiran, N., Chorev, Z., & Sapir, A. (2000). Component processes in task switching. *Cognitive Psychology*, **41**, 211–253.
- Meuter, R. F. I., & Allport, D. A. (1999). Bilingual language switching in naming: Asymmetrical costs of language selection. *Journal of Memory and Language*, **40**, 25–40.
- Meyer, D. E., & Kieras, D. E. (1997). EPIC—A computational theory of executive cognitive processes and multiple-task performance: Part 1. Basic mechanisms. *Psychological Review*, **104**, 3–65.
- Meyer, D. E., & Kieras, D. E. (1999). Précis to a practical unified theory of cognition and action: Some lessons from computational modeling of human multiple-task performance. In D. Gopher & A. Koriati (Eds.), *Attention and performance XVII* (pp. 15–88). Cambridge, MA: MIT Press.
- Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, **24**, 167, 202.

- Monsell, S. (1996). Control of mental processes. In V. Bruce (Ed.), *Unsolved mysteries of the mind* (pp. 93–148). Hove, UK: Erlbaum/Taylor & Francis.
- Monsell, S., Azuma, R., Eimer, M., Le Pelley, M., & Strafford, S. (1998). *Does a prepared task switch require an extra (control) process between stimulus onset and response selection?* Poster presented at the 18th International Symposium on Attention and Performance, Windsor, England.
- Monsell, S., & Driver, J. S. (Eds.) (2000). *Attention and performance XVIII: Control of cognitive processes*. Cambridge, MA: MIT Press.
- Monsell, S., Taylor, T. J., & Murphy, K. (2001). Naming the color of a word: Is it responses or task sets that compete? *Memory & Cognition*, **29**, 137–151.
- Monsell, S., Yeung, N. P., & Azuma, R. (2000). Reconfiguration of task-set: Is it easier to switch to the weaker task? *Psychological Research*, **63**, 250–264.
- Norman, D. A., & Shallice, T. (1986). Attention to action: Willed and automatic control of behaviour. In R. J. Davidson, G. E. Schwartz, & D. Shapiro (Eds.), *Consciousness and self-regulation* (Vol. 4, pp. 1–18). New York: Plenum.
- Nosofsky, R. M., & Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review*, **104**, 266–300.
- O'Reilly, R. C., & Farah, M. J. (1999). Simulation and explanation in neuropsychology and beyond. *Cognitive Neuropsychology*, **16**, 49–72.
- Palmeri, T. J. (1997). Exemplar similarity and the development of automaticity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **23**, 324–354.
- Phaf, R. H., Van der Heijden, A. H. C., & Hudson, P. T. W. (1990). SLAM: A connectionist model for attention in visual selection tasks. *Cognitive Psychology*, **22**, 273–341.
- Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, **32**, 3–25.
- Posner, M. I., & Snyder, C. R. R. (1975). Attention and cognitive control. In R. L. Solso (Ed.), *Information processing and cognition: The Loyola Symposium*. Hillsdale, NJ: Erlbaum.
- Roberts, S., & Pashler, H. (2000). How persuasive is a good fit? A comment on theory testing. *Psychological Review*, **107**, 358–367.
- Rogers, R. D., & Monsell, S. (1995). Costs of a predictable switch between simple cognitive tasks. *Journal of Experimental Psychology: General*, **124**, 207–231.
- Rogers, R. D., Sahakian, B. J., Hodges, J. R., Polkey, C. E., Kennard, C., & Robbins, T. W. (1998). Dissociating executive mechanisms of task control following frontal lobe damage and Parkinson's disease. *Brain*, **121**, 815–842.
- Rubinstein, J., Meyer, D. E., & Evans, J. E. (2001). Executive control of cognitive processes in task switching. *Journal of Experimental Psychology: Human Perception and Performance*, **27**, 763–797.
- Salthouse, T. A., Fristoe, N., McGurthy, K. E., & Hambrick, D. Z. (1998). Relation of task switching to speed, age, and fluid intelligence. *Psychology and Aging*, **13**, 445–461.
- Schwartz, M. F., Reed, E. S., Montgomery, M., Palmer, C., & Mayer, N. H. (1991). The quantitative description of action disorganisation after brain damage: A case study. *Cognitive Neuropsychology*, **8**, 381–414.
- Shallice, T. (1982). Specific impairments of planning. *Philosophical Transactions of the Royal Society of London B*, **298**, 199–209.
- Shallice, T. (1988). *From neuropsychology to mental structure*. Cambridge, UK: Cambridge Univ. Press.
- Shallice, T. (1994). Multiple levels of control processes. In C. Umiltà & M. Moscovitch (Eds.), *Attention and performance XV* (pp. 395–420). Cambridge, MA: MIT Press.

- Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review*, **84**, 127–190.
- Spector, A., & Biederman, I. (1976). Mental set and shift revisited. *American Journal of Psychology*, **89**, 669–679.
- Spelke, E., Hirst, W., & Neisser, U. (1976). Skills of divided attention. *Cognition*, *4*, 205–230.
- Stroop, J. R. (1935). Studies in interference in serial verbal reactions. *Journal of Experimental Psychology*, **18**, 643–662.
- Sudevan, P., & Taylor, D. A. (1987). The cuing and priming of cognitive operations. *Journal of Experimental Psychology: Human Perception and Performance*, **13**, 89–103.
- Tornay, F. J., & Milán, E. G. (2001). A more complete task-set reconfiguration in random than in predictable task switch. *Quarterly Journal of Experimental Psychology*, **54A**, 785–803.
- Waszak, F., Hommel, B., & Allport, A. (submitted). Task-switching and long-term priming: Role of episodic S-R bindings in task-shift costs.
- Wylie, G., & Allport, A. (2000). Task switching and the measurement of “switch costs.” *Psychological Research*, **63**, 212–233.
- Yeung (1999). *Switching between simple cognitive tasks: interactions between executive control and task properties*. Unpublished doctoral dissertation, University of Cambridge, Cambridge, England.
- Zhang, H., Zhang, J., & Kornblum, S. (1999). A parallel distributed processing model of stimulus–stimulus and stimulus–response compatibility. *Cognitive Psychology*, **38**, 386–432.