



The Parallel Episodic Processing (PEP) model 2.0: A single computational model of stimulus-response binding, contingency learning, power curves, and mixing costs

James R. Schmidt^{a,*}, Jan De Houwer^a, Klaus Rothermund^b

^a Department of Experimental Clinical and Health Psychology, Ghent University, Belgium

^b Department of Psychology, Friedrich-Schiller-Universität Jena, Germany

ARTICLE INFO

Article history:

Accepted 24 October 2016

Available online 7 November 2016

Keywords:

Computational modelling

Episodic memory

Contingency learning

Practice

Mixing costs

Binding

ABSTRACT

The current paper presents an extension of the Parallel Episodic Processing model. The model is developed for simulating behaviour in performance (i.e., speeded response time) tasks and learns to anticipate both how and when to respond based on retrieval of memories of previous trials. With one fixed parameter set, the model is shown to successfully simulate a wide range of different findings. These include: practice curves in the Stroop paradigm, contingency learning effects, learning acquisition curves, stimulus-response binding effects, mixing costs, and various findings from the attentional control domain. The results demonstrate several important points. First, the same retrieval mechanism parsimoniously explains stimulus-response binding, contingency learning, and practice effects. Second, as performance improves with practice, any effects will shrink with it. Third, a model of simple learning processes is sufficient to explain phenomena that are typically (but perhaps incorrectly) interpreted in terms of higher-order control processes. More generally, we argue that computational models with a fixed parameter set and wider breadth should be preferred over those that are restricted to a narrow set of phenomena.

© 2016 Elsevier Inc. All rights reserved.

1. Introduction

As is inevitably the case in most disciplines, research in cognitive psychology tends to be highly segregated. Research is often focused on individual phenomena, rather than on broader-picture processes (Hommel & Colzato, 2015), which can frequently lead to “reinvention of the wheel” in differing subfields. Although a single “big theory” of the full complexities of human perception, cognition, timing, memory, etc. would be impossible, cross-paradigm comparisons are often highly informative (e.g., Hommel, Proctor, & Vu, 2004; Schmidt, Cheesman, & Besner, 2013) and larger-scale cognitive frameworks are highly useful (e.g., Anderson, 2007; Anderson et al., 2004; Laird, 2012; Laird, Newell, & Rosenbloom, 1987). Of course, some ideas have been more widely applied to various subfields, but not necessarily within one integrated framework. In the current manuscript, we present Version 2.0 of the Parallel Episodic Processing (PEP) model, and show how a small set of assumptions about memory storage and retrieval can have wide explanatory power across a range of performance (i.e., response time) paradigms in the learning, timing, binding, and attentional control domains. More generally, we argue that there is much to gain by developing fixed-parameter models of broadly-applicable processes.

* Corresponding author at: Ghent University, Henri Dunantlaan 2, B-9000 Ghent, Belgium.

E-mail address: james.schmidt@ugent.be (J.R. Schmidt).

The PEP model that we present belongs to a general class of models, often labelled as *episodic*, *instance*, *exemplar*, or *event file* memory models. Such models, often formally specified (Hintzman, 1984, 1986, 1988; Logan, 1988; Medin & Schaffer, 1978; Nosofsky, 1988a, 1988b), differ in some respects from one instantiation to another, but also have a lot in common. In particular, such models make two key assumptions. First, representations of individual events (e.g., trials in an experiment) are recorded into individual memories in episodic memory. These episodes might contain information about the stimuli experienced and the response that was executed (Hommel, 1998), in addition to timing information such as response times (Kinoshita, Forster, & Mozer, 2008; Kinoshita, Mozer, & Forster, 2011; Mozer, Kinoshita, & Davis, 2004; Schmidt, 2013c, 2014; Schmidt, Lemerrier, & De Houwer, 2014; Schmidt & Weissman, 2016), and perhaps even more complex information such as task sets (Waszak, Hommel, & Allport, 2003) or cognitive control settings (Egner, 2014). Second, on subsequent events episodic memories are retrieved to aid in performance. For instance, if a participant has already encoded several episodes in which a blue stimulus was presented and J-key keyboard response was made, then presentation of a blue stimulus again will lead to the retrieval of these “blue” episodes, which will in turn automatically facilitate a J-key response. In this way, simple storage and retrieval of episodes can explain the performance benefits that come with practice (Logan, 1988).

To begin, we note that the scope of the current work is on performance paradigms, for which response times are the typical dependent measure. Episodic memory frameworks have also been instantiated extensively for simulating memory effects proper, such as recognition (e.g., Shiffrin & Steyvers, 1997) and recall (e.g., Howard & Kahana, 2002) of studied items. Much work has further aimed to tie episodic encoding and retrieval into biologically plausible models of the brain (for a review, see O’Reilly & Rudy, 2001). Though the reach of the present manuscript is broad, the discussion will be focused on the impact of episodic memories on speeded response times.

In addition to practice curves, episodic memory models have been proposed for a broad range of other performance phenomena, including contingency learning (Schmidt, Crump, Cheesman, & Besner, 2007; Schmidt & De Houwer, 2012a, 2012c, 2012d, 2016a; Schmidt, De Houwer, & Besner, 2010), feature binding effects (Frings, Rothermund, & Wentura, 2007; Hommel, 1998), negative priming (Rothermund, Wentura, & De Houwer, 2005), proportion congruent effects (Schmidt, 2013a, 2013b), congruency sequence effects (Hommel et al., 2004; Mayr, Awh, & Laurey, 2003; Mordkoff, 2012; Schmidt & De Houwer, 2011), evaluative conditioning (Schmidt & De Houwer, 2012b), task switch costs (Logan & Bundesen, 2003, 2004; Logan & Schneider, 2006a, 2006b; Logan, Schneider, & Bundesen, 2007; Schmidt & Liefoghe, 2016; Schneider & Logan, 2007, 2009), rhythmic responding (Kinoshita et al., 2008; Kinoshita et al., 2011; Mozer et al., 2004; Schmidt, 2013c, 2014; Schmidt, Lemerrier et al., 2014; Schmidt & Weissman, 2016), and various other phenomena.

Though similar accounts have been forwarded for a diverse range of phenomena, what is lacking is a general purpose computational modelling framework to explain these phenomena. Most of the above-described phenomena, for instance, have only been verbally described in terms of episodic binding. Many other findings have been formally modelled, of course, but typically with different models (or different adaptations of models) from one phenomenon to the next. One of the greatest strengths of computational modelling is that it forces a researcher to clearly specify the dynamics of a theoretical account, which eliminates any ambiguity in whether a given effect should or should not be predicted. When models are adjusted on a simulation-by-simulation basis (e.g., by changing learning rates or other major parameters in the model) or when altogether different models are constructed for different phenomena, most of this benefit is lost. Any given model can probably be “forced” to fit an individual dataset with enough overfitting. More impressive is a model that can fit multiple datasets from multiple diverse domains (or predict novel findings) without any simulation-by-simulation adjustments. For instance, if one version of a model is constructed to simulate Result A and another version is constructed to simulate Result B, this does not necessarily provide good support for the model. For example, it could be that the way the model is parameterized to produce Result A is fundamentally incompatible with Result B, and vice versa.

The present manuscript presents some initial steps in the development of a broader-focus episodic memory computational model aimed to simulate a range of performance phenomena within a single processing framework. Of course, one paper is insufficient for simulating all relevant phenomena at once, but we aim to model a range of findings from differing fields: practice effects, contingency learning, stimulus-response (S–R) binding effects, and mixing costs. Below, we first describe in general terms how these various findings might be related to the same encoding and retrieval processes. Following this, we explain the PEP 2.0 framework in detail and present a series of simulations of data reported in the literature. We also examine the backward compatibility of PEP 2.0 by verifying whether PEP 2.0 can model the (cognitive control) phenomena that the original PEP model could account for (see Appendix B).

1.1. Power law of practice

Over time, performance improves with practice in most tasks that involve repeated behaviours (Newell & Rosenbloom, 1981). Indeed, the speedup in performance follows a highly predictable function: a power function in participant-averaged, blocked data, or an exponential function with individual-participant, trial-by-trial analyses (Heathcote, Brown, & Mewhort, 2000; Myung, Kim, & Pitt, 2000). In less technical terms, speed of responding rapidly improves early on in a task, then continues to improve at an ever-diminishing rate. This highly-regular phenomenon is often referred to as the *power law of practice*. Logan (1988) demonstrated how episodic storage and retrieval is sufficient to simulate the power law of practice. Though the math of his “winner takes all” model of retrieval is slightly different than the mechanism in the PEP model (i.e., where multiple episodes bias responding in unison), the same principle applies to all episodic learning models. Each time a

model makes a response to a stimulus, a new episode coding for the stimulus-response binding is recorded. Over practice, these episodes accumulate such that there are more and more memories to retrieve to automatically bias the correct response on presentation of the stimulus.

As Schmidt (2016a) briefly pointed out (but did not demonstrate directly), most effects observed in any paradigm should generally decrease with practice. For instance, consider the Stroop congruency effect (Stroop, 1935). Participants identify the print colours of colour words slower and less accurately when the word and colour are incongruent (e.g., the word “red” printed in green) relative to congruent (e.g., “red” in red). Along with a generalized decrease in response speed, this congruency effect decreases with practice (Dulaney & Rogers, 1994; Ellis & Dulaney, 1991; MacLeod, 1998; Simon, Craft, & Webster, 1973; Stroop, 1935). Of course, there can be many reasons why an effect will shrink with time. However, one (full or partial) reason for the decrease in the congruency effect is simply that the faster participants can respond to the target colour, the less time there is for the word to conflict with this decision. In different words, response time is initially slow but then speeds toward an asymptotic ideal over time. Initially slow incongruent trials stand to gain more as responding speeds up (i.e., the difference between starting and asymptotic RT is large) than the already-fast congruent trials. Thus, as we will demonstrate for the first time, an episodic learning model predicts both a general speedup in responding and decreasing congruency effects with practice. Simulation 1 investigates the ability of the PEP model to simulate these observations.

1.2. Contingency learning

Learning via episodic retrieval is not exclusive to regularities between relevant (target) stimuli and responses (i.e., repeated responding to targets, as is the case in practice effects). Participants can also learn regularities between irrelevant (distracter) stimuli and responses. Consider the colour-word contingency learning paradigm (Schmidt et al., 2007). Participants respond to the print colour of colour-unrelated neutral words and each word is presented most often in one colour (e.g., “move” most often in blue, “sent” most often in green, etc.). Similar to related learning paradigms (e.g., Miller, 1987), participants respond faster and more accurately to *high contingency* stimuli (e.g., “move” in blue) than to *low contingency* stimuli (e.g., “move” in green). This learning effect can also be explained by episodic storage and retrieval (Schmidt et al., 2010). For instance, because “move” is presented most often in blue, most “move” episodes will be linked to a blue response. As such, on subsequent presentations of the word “move,” the blue response will be facilitated via episodic retrieval. Thus, high contingency trials will be speeded, because the predicted (blue) response is also the correct response. On low contingency trials, a different response must be made (e.g., green), so the benefit is lost. This contingency learning benefit has only been indirectly simulated in PEP simulations of the proportion congruent effect (Schmidt, 2013a, 2016a). In Simulations 2–4 of the current paper, we model colour-word contingency learning benefits directly. We also simulate acquisition curves and two recent dissociation procedures aimed to distinguish possible mechanisms underlying contingency learning (Schmidt & De Houwer, 2016a).

1.3. Stimulus-response binding

So far, we have considered how episodic encoding and retrieval can produce longer-term learning effects. That is, both target-response practice benefits and distracter-response contingency learning effects result from the accrual of multiple episodes linking a stimulus to a response. Work in the binding domain, however, has investigated how recently-encoded events can also affect behaviour. For instance, in the binding paradigms of Hommel (1998) participants are presented with pairs of trials. The stimulus on Trial 1 either repeats or alternates (i.e., changes) on Trial 2. Orthogonal to this, the response on Trial 1 either repeats or alternates on Trial 2. On *complete repetition* trials, the same response is made to the same stimulus on both trials. On *complete alternation* trials, a *different* response is made to a *different* stimulus. Both complete repetition and complete alternation trials are responded to relatively quickly. Performance is much slower on *partial repetition* trials, where either (a) the same response is made to a different stimulus, or (b) a different response is made to the same stimulus. According to the feature integration account of these data, stimuli and responses are bound together in *event files*. On partial repetition trials, the repeated stimulus (or repeated response) has to be “unbound” from the response (or stimulus) that it was previously bound to in order to bind to the new response (or stimulus). This is said to account for the slowing. In contrast, on complete repetition trials the stimulus and response do not have to be unbound (i.e., because they are already correctly bound in the event file). Unbinding is similarly not necessary on complete alternation trials, because a different stimulus and response were just encoded.

Similar binding effects have been investigated between distracting stimuli and responses (Frings et al., 2007; Rothermund et al., 2005). For instance, Frings and colleagues had participants respond to target letters while ignoring distracting “flanker” letters (e.g., DFDFD, where F is the target and D the distracter). The target letter response could either repeat or alternate. Orthogonal to this, the distracting letters could either repeat or alternate. Responses were faster when both the response and distracter repeated (complete repetition; e.g., DFDFD → DFDFD) than when just the response repeated (partial repetition; e.g., DFDFD → JFJFJ). Similarly, responses were faster when both the response and distracter alternated (complete alternation; e.g., DFDFD → JKJKJ) than when just the distracter repeated (partial repetition; e.g., DFDFD → DKDKD). Analogous to effects in the paradigm of Hommel (1998), this has been argued to be due to the distracter retrieving the response that it was

just linked with, which benefits a complete repetition (where the same response has to be made), but impairs a partial repetition (where a different response has to be made).

In these sorts of binding paradigms, it is the immediately preceding (or at least recent; Pösse, Waszak, & Hommel, 2006) trial that is of interest. However, the same mechanism that produces practice benefits and contingency learning effects across many trials might also be responsible for S–R binding effects. The only reasonable assumption that has to be made is that the most recently encoded events have a larger effect on retrieval than more distantly encoded events. In other words, memory retrieval may have some cumulative effect over time, but new memories are the most accessible. This makes sense intuitively. If you have just seen the distracting flanker “D” for the hundredth time, for instance, you would certainly expect that you would be more primed by this 100th occurrence if you see “D” again than you would be by, say, the 20th time you saw the letter “D”. This is already the case in the PEP model: older memories become progressively less accessible as new memories are created. As will be discussed extensively in the Model Description section, each time a stimulus is encountered, older memories of that stimulus are weakened while a new memory is being written. This is akin to (partially) forgetting where you have previously left your keys after you place them somewhere new: you want to remember where you have put your keys *recently*, not where you put them days, weeks, or months ago.

The episodic storage and retrieval mechanisms described earlier should correctly model binding effects. On trials where the response repeats, repetition of the distracting (or target) stimulus leads to better retrieval of the just-encoded episode (i.e., where the exact same stimulus–response pairing was encoded). Thus, a complete repetition will be strongly facilitated, whereas a partial repetition will not. When the response alternates, repetition of the distracter will boost retrieval (i.e., activation) of the just-encoded episode, which will in turn bias retrieval in favour of the just-encoded, but now *incorrect*, response. Thus, partial repetitions will be impaired relative to complete alternations. Though this retrieval account does not, strictly speaking, involve any “unbinding,” it shares obvious similarities with previous accounts. Simulation 5 will evaluate whether the learning mechanism in the PEP model is sufficient to replicate Frings et al. (2007).

1.4. Mixing costs

The preceding sections discuss how an episodic learning model can bias one response over others. That is, the model learns to anticipate *what* response to emit given the stimuli it is presented. Episodic models can also learn to anticipate *when* to respond. In any repetitive performance task, responding becomes highly rhythmic (Grosjean, Rosenbaum, & Elsinger, 2001). That is, participants begin responding at a similar speed from one trial to the next. Indeed, this phenomenon is obvious enough that it can be heard in the keystrokes of participants performing essentially any speeded response time task. Rhythmic responding is even observed when response–stimulus intervals are randomly varied (Schmidt, 2016b), suggesting that participants learn regularities in timing between stimulus onset and the response.

One way to study the effect of task pace on performance is with mixing cost paradigms. For instance, Los (1999a; see also, Los, 1994, 1996, 1999b; Lupker, Brown, & Colombo, 1997; Van Duren & Sanders, 1988) presented participants with digits that were either normal or degraded in some way (e.g., missing segments or added noise). Of course, degraded stimuli are responded to more slowly than normally-presented stimuli. Critically, the two item types were either presented in separate *pure blocks* or intermixed together in *mixed blocks*. Performance on both the normal and degraded item types are generally slower in mixed blocks than in pure blocks, but this is especially the case for normal items.

There are many differing accounts of how mixing costs arise (Grice, 1968; Kohfeld, 1968; Ollman & Billington, 1972; Van Duren & Sanders, 1988), but many of these accounts assume that criterion adjustments are made for each list type (i.e., pure easy, pure hard, and mixed). This is also the case in the PEP model. Like practice, contingency, and binding effects, mixing costs can result from episodic storage and retrieval. Timing information is stored in episodes (see Dignath, Pfister, Eder, Kiesel, & Kunde, 2014; Matzel, Held, & Miller, 1988), and this timing information can subsequently be retrieved to anticipate when to respond on following trials. That is, retrieval of memories about how long it took to respond in the past provides the cognitive system with information about the point in time in a trial when a response is likely (e.g., x ms after stimulus onset), which results in heightened preparedness to respond at that moment. This produces rhythmic responding. Mixing costs result from this rhythmic responding because a fixed rhythm can easily be maintained when all of the items are of similar difficulty (pure lists), but this is less effective when harder (i.e., slow) and easier (i.e., fast) items are intermixed with one another (mixed lists). For instance, when all trials are easy, the time to response can be easily anticipated, because most responses fall in the same (early) response window. Similarly, most responses fall in the same (later) response window in the pure hard list. However, the expected time to respond is less certain when easy and hard trials are randomly intermixed. In Simulation 6, we will assess the ability of the PEP model to simulate mixing costs.

2. Model description

A full explanation of the math underlying the PEP model can be found in Appendix A. Here, we present a briefer conceptual overview. A pictorial representation of the PEP model is presented in Fig. 1. The PEP model contains five different types of nodes. Two of these are Input nodes, one set for target stimuli (e.g., the print colours of Stroop stimuli) and another set for distracters (e.g., the words of Stroop stimuli). As in most models, all Input nodes receive some random noise activation. This noise adds in some realistic variability in response times and the occasional error. In addition, the model is “presented” a

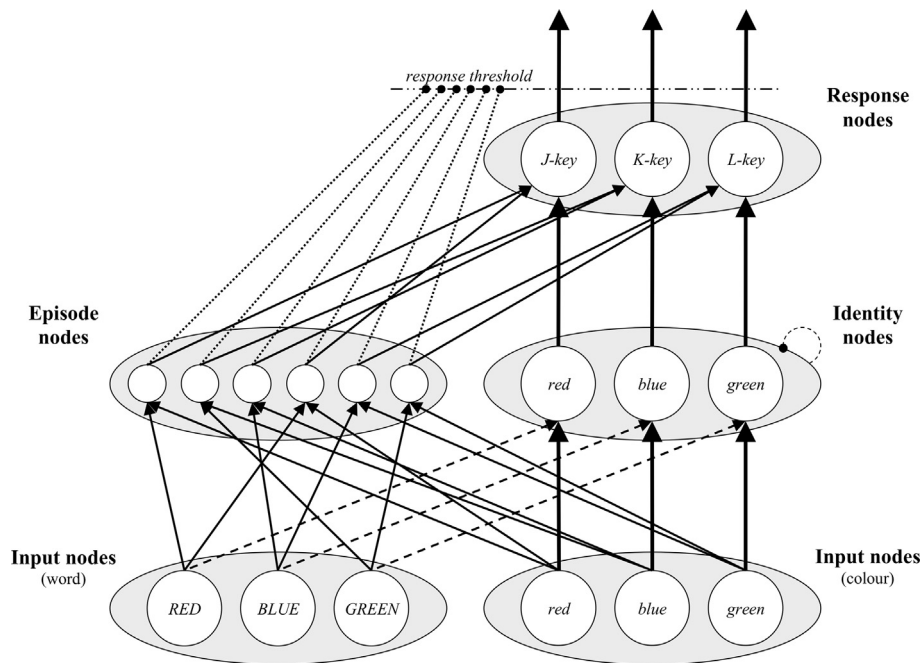


Fig. 1. A representation of the Parallel Episodic Processing (PEP) model. Input nodes feed activation through Identity nodes (where conflict occurs) and Response nodes in an algorithmic route. Learning occurs via Episode nodes, which encode the stimuli and response for each trial. Anticipation of when to respond is also achieved via episodic biasing of the response threshold. Dashed lines indicate connections between distracter Input nodes and Identity nodes that only exist when distracters are related to targets (e.g., Stroop-like stimuli). Dotted lines indicate inhibitory connections.

stimulus by adding extra activation (i.e., signal) to the Input node for the stimulus that is “shown” on the trial. With this input, activation accrues (logistically) on each processing cycle (i.e., simulated millisecond). Input nodes feed activation forward to Identity nodes, one for each target concept. Within Identity nodes, competition occurs if more than one candidate response is activated (note that neutral distracting stimuli do not connect to Identity nodes, because such Identity nodes would have no valid Response node to connect to, anyway). Identity nodes further feed activation on to Response nodes. After sufficient evidence for one of the responses accrues (i.e., when the activation for a response exceeds the response threshold), the model responds. This part of the model, which might be called the *algorithmic route*, is similar to most models developed for Stroop-type tasks (e.g., Cohen, Dunbar, & McClelland, 1990; Cohen & Huston, 1994). Other than subtle adjustments to some parameters to compensate for added changes elsewhere (described below), this part of the model remains unchanged from previous versions.

The more crucial part of the model is the episodic store and it is there where all the key effects of interest emerge from. On each trial, a new Episode node is created. Each episode encodes both the target and distracter presented, and the response that is made (and also the response time, which will be described later). Note that in previous versions of the model, the target was not encoded into the episodes, but this was merely for computational simplicity. This represents the only major change to the model in the current series of simulations. Also, because the target now activates responses via both the algorithmic and episodic routes, some rebalancing of other parameters was of course necessary (see Appendix A), but no fundamental changes were made. As in the previous versions, stimuli and responses are encoded in episodes in connection weights that are strengthened via Hebbian learning (i.e., nodes that are strongly active become linked to the episode). Thus, stimuli that were presented on the current trial are linked to the new episode in a “binding” (Hommel, 1998; Treisman & Gelade, 1980) or “chunking” (Wickelgren, 1979) process and these stimuli will therefore later serve as effective retrieval cues of the episode.

On subsequent trials, both the target and distracter activate Episode nodes via these connection weights. For instance, the colour red will activate nodes in which the colour red has been encoded. Similarly, the word “red” will activate nodes in which the word “red” is encoded. Each activated Episode node will facilitate the Response node corresponding to the encoded response. For instance, an Episode node coding a trial in which the J-key response was made will facilitate a J-key response. Note that each activated Input node activates all Episode nodes that it is connected to. Each of these activated Episode nodes activates the response that it is connected to, such that many episodes cumulatively bias responses. To the extent that more activated episodes point to one response or another, each response will be proportionally activated. That is, each activated episode increases the activation of its connected Response node by some amount. This contrasts slightly with some models in which episodes “race” each other and the “winner” determines the response (e.g., Logan, 1988). In the Discussion, we will expand on this distinction and indicate one strength of the PEP formulation relative to a race model of retrieval.

Importantly for the PEP model, it is assumed that recently-encoded episodes have a much stronger impact on behaviour than older ones. This is important for allowing the model to be adaptable to, for instance, a change in task contingencies: new relationships can be learned rapidly only if the newly-encoded information is not diluted in a sea of older information. For instance, it would be hard to remember where you parked your car today if every single place you have parked it before simultaneously came to mind at equal strength. Forgetting is therefore a part of learning in the model. In particular, Episode nodes become less accessible (via weakened connection strengths) each time they are retrieved (i.e., activated). As a result, more recently encoded episodes have more effect on behaviour than older ones.

Note that weakening the accessibility of retrieved episodes might initially seem counterintuitive, particularly because the episodes *most like the current trial* will be those that are weakened the most. However, it is important to realize that the *newly-encoded episode* is added to memory and thus serves to strengthen the link between the presented stimuli and the response. Thus, the model will have *better* memory for a just-presented stimulus, even though older episodes of this stimulus are weakened. That is, retrieval does *not* cause forgetting in the model (typically it *strengthens* a memory), *unless* one tries to retrieve a memory but identifies something else. Note, too, that this exponential weighting of episodes is akin to simple association formation via Hebbian learning with decay (but distributed across episodes), backpropagation, or virtually any other adaptive learning computation in computational models: if the just-encountered event updates an association strength, then it is by definition the case that the influence of previous events is decreased. Indeed, weakening some connections while strengthening others is a *necessary component* of an adequately functioning learning mechanism (Gerstner & Kistler, 2002). This memory consolidation is important for simulating binding effects, as previously discussed (i.e., as the effect of the just-encoded episode would be far too “diluted” if older episodes were not weakened), and also ensures a rapid learning rate. Furthermore, *recency* of an event is known to effect retrievability of memories (Ebbinghaus, 1913), and the current mechanism is one way to achieve this. Similarly, this mechanism is further consistent with retrieval-induced forgetting phenomena (Anderson, Bjork, & Bjork, 1994), only without center-surround inhibition: episodes that are partially retrieved (i.e., weakly activated) but also partially inconsistent with the current event are weakened (thereby weakening these old memories); episodes that are identical to the current event are also weakened, but a new episode re-encodes the same information (thereby strengthening this memory). Note that the above aspects regarding retrieval decay were always true in the PEP model, but there was some slight adjustment of *rate* with which current episodes have more influence than older ones in the current work (discussed in detail later).

In addition to coding the stimuli that were presented and the response that was made, the PEP model also encodes response time information into episodes. This response time information is later retrieved in order to anticipate *when* to respond. In particular, the response times from recently encoded episodes are retrieved on each trial and this retrieval leads to dynamic adjustment of the response threshold (i.e., amount of evidence needed for a response). Like the contingency learning mechanism, more recent events have a larger impact on the threshold adjustment than older ones. The response threshold decreases (temporarily) at the time corresponding to recently encoded response times. That is, the model expects to be able to respond at a similar speed as it has just been responding. If there is enough evidence for one of the responses at the anticipated time, then responding is facilitated. That is, less response activation is needed to cross the threshold while it is (temporarily) reduced, so as long as the activation for a response *does* cross this reduced threshold (maintaining the “pace”) responding is fast. If the anticipated timing is missed (i.e., activation does not cross the reduced threshold for any of the responses), then the threshold for responding goes back up, the pace is broken, and the rhythmic benefit is lost. This causes the model to respond rhythmically, that is, to respond at a similar speed from one trial to the next. As long as the model has enough evidence to respond at the anticipated time (i.e., when the response threshold is decreased), responding will be speeded (i.e., because less evidence is needed to cross the reduced threshold). If the rhythm is broken (e.g., not enough evidence for a response when the threshold is temporarily lowered), then responding is slower. In previous work, it was demonstrated that this mechanism is important for simulating several findings in the cognitive control literature (see Schmidt, 2013c; Schmidt & Weissman, 2016) and also proves important for mixing costs. No changes were made to this mechanism in the present work. Note also that the model simulates rhythmic timing, but does not currently simulate time perception. The model is simply given the simulated RT to store, rather than modelling an internal clock or some other time perception mechanism. This abstraction of the time perception process is for simplicity only, and future versions of the model might aim to model time perception directly, for instance, by borrowing from one of the many time perception models (e.g., Church, 1984; French, Addyman, Mareschal, & Thomas, 2014; Gibbon, Church, & Meck, 1984; Miall, 1989).

3. Modelling principles

3.1. Open-access and documented source code

The Java code for Version 2.0 of the PEP model is freely available on the web page of the lead author (users.ugent.be/~jaschmid/PEP/), along with the code for all previous versions of the model. We made the program available because, as Addyman and French (2012) rightly point out, it can be difficult to replicate a model from descriptions in a paper, even for an experienced programmer. We have also made efforts to clearly document the model, provide a relatively user-friendly interface, and label variables intuitively.

3.2. Single parameter set

As we have emphasized elsewhere in this manuscript, one fixed parameter set was used for all simulations. That is, nothing in the model was adjusted from one simulation to the next. All that changed was the stimuli presented. Of course, ideal parameters were determined through initial testing, but the final reported simulations were with one fixed model. On the plus side, this reduces concerns that the model is being forced to “fit” a single dataset with no generality to other datasets. On the down side, the decision to fix the parameter set is so conservative that not all fits to the data will be as impressive. Indeed, there are meaningfully-large differences between the various experiments we model (e.g., verbal Stroop, keypress Stroop, prime-probe, flanker, digit identification, etc.) that produce real differences in human data, so changes in parameters might be justified in some cases. However, we feel that a good model should either be robust to these task differences or be formally specified to explain these differences (e.g., Kornblum, Stevens, Whipple, & Requin, 1999).

3.3. Large simulation sizes

For all simulations, 500 simulated “participants” were run. Obviously, this means that standard error in effect estimates will be very small in the model. Because of this, we do not report statistics. However, all discussed patterns in the data were statistically significant (i.e., $p < 0.05$), typically by a large margin. Indeed, during initial testing of the model the findings reported throughout this paper were robust, at least qualitatively, to reasonable adjustments in parameters (to be discussed in further detail in the Discussion section). We also report only cycle times (i.e., simulated response times). This is in part for brevity. Additionally, the model produced quite low error rates in the absence of conflict (often at or near zero). This is similar to what can be seen in many behavioural investigations (particularly verbal naming tasks), though errors were even less frequent in the model. With conflict, errors were comparable to behavioural results (albeit again lower). Further refinements might be made to produce a better balance between cycle times and errors, but we do note that (a) none of the error data produced inconsistent evidence to that reported for cycle times, and (b) the model was not producing good RT data at the expense of inflated errors (as is sometimes a problem in other models). The full dataset and condition means are available for download on the website of the lead author (with the model code download).

3.4. Backwards compatibility

Checks for backwards compatibility are sometimes made (e.g., Perry, Ziegler, & Zorzi, 2010), but this is not always the case. Typically entirely new models or variations of old models are used for each simulated dataset. Related to our fixed-model philosophy, we also ensured that the current version of the model is able to simulate the findings that the older versions of the model were able to simulate, namely item-specific proportion congruent effects (Schmidt, 2013a), list-level proportion congruent effects (Schmidt, 2013c), congruency sequence (or Gratton) effects (Schmidt & Weissman, 2016), and asymmetric list shifting effects (Schmidt, 2016a). For brevity, these results are presented in Appendix B, but further details can be obtained in the original papers or from the lead author on request. These “cognitive control” paradigms are also discussed in the Discussion section.

4. Simulation 1: Stroop and power curves

In Simulation 1, we model a simple Stroop task, with three target colours (blue, red, green), three distracting colour words (“blue,” “red,” “green”), and two neutral words (“make,” “find”). Each of the five words were presented equally often in all colours, making for three congruent pairings, six incongruent pairings, and six neutral pairings. Each of these was presented to the model twice per randomized block, with 10 blocks total (300 trials). Two main predictions are assessed. First, in line with the power law of practice, we should see a generalized speedup in responding over time (i.e., a main effect of block). In the PEP model, this is both the result of the accumulation of target-response pairings in memory and the gradual strengthening of the connections between Identity and Response nodes (see Appendix A). Second, the congruency effect should decrease across blocks. In the model, this is due to the fact that the target can more rapidly retrieve the correct response from memory with practice, thereby speeding responding and leaving the distracting word less time to influence colour naming. Worded differently, accumulation of episodes leads the model to more quickly determine the response on the basis of the colour, speeding performance over time toward asymptote (i.e., toward the fastest response time possible). The initially-slow incongruent trials can gain more from this general speeding than the already-fast congruent trials, because the responding on congruent trials is already close to the fastest response time possible (i.e., a floor effect). As a result of the differential impact of speeding performance on incongruent compared to congruent trials, the congruency effect will be reduced.

The results of Simulation 1 are presented in Fig. 2. As can be observed, a general speedup in performance was observed across blocks, which roughly obeys a power function (trend lines on figure). Most interestingly, we also observe that the congruency effect diminishes with practice, from 170 cycles in Block 1 down to 56 cycles in Block 10. Note that the rate with which the distracting word interferes with target processing does not actually change in the model. This decrease in the congruency effect is observed simply because the colour more quickly retrieves the correct response the more practice has pro-




Fig. 2. Simulation 1 cycle times of congruent, neutral, and incongruent trials across training blocks with fitted power curves.

gressed, so the word has less time to impair colour naming. As a minor note, we also observe an asymmetry in the congruency effect, such that incongruent trials are more impaired relative to neutral trials (54 ms) than congruent trials are speeded (31 ms). This asymmetry is a hallmark of the Stroop effect (for a review, see MacLeod, 1991). This is produced by the logistic activation function used in various computational models (e.g., Cohen & Huston, 1994; Cohen et al., 1990), newly added to the PEP model (see Appendix A). In less technical terms, activation of the correct response benefits little from redundant activation from a congruent word, but can be impaired by competition with an incongruent response.

5. Simulation 2: Contingency learning acquisition curves

Contingency effects in the colour-word contingency learning paradigm are acquired incredibly quickly. For instance, the effect was already present in the very first block of 18 trials in Schmidt et al. (2010). Such rapid learning is not atypical for performance-based learning paradigms (e.g., Lewicki, 1985, 1986; Lewicki, Hill, & Czyzewska, 1992; Nissen & Bullemer, 1987). Interestingly, the effect does not seem to increase or decrease noticeably over time. On the one hand, this might seem to suggest that contingency knowledge is maximally acquired almost immediately and does not increase afterwards. On the other hand, the previous simulation demonstrated that effect magnitudes tend to decrease with practice. This should presumably also be the case for contingency learning: as performance improves, the contingency should have less time to influence behaviour, and the contingency effect should *decrease* with practice. Possibly, there are two competing forces: (a) acquisition is partially cumulative, but (b) as the encoded contingency is strengthened, there is also less time for the contingency to effect behaviour (i.e., due to practice). Simulation 2 tests how the contingency effect evolves over time in the PEP model.

In particular, we simulate the study from Schmidt and De Houwer (2016b), which is similar to other acquisition studies in our lab (e.g., Schmidt et al., 2007; Schmidt et al., 2010). The stimulus frequencies are presented in Table 1. In each of ten blocks, each of three words is presented most often (8 of 10 times) in one colour and once each in the remaining two colours (300 trials). In addition, the learning blocks are preceded by three practice blocks, in which participants respond to the stimulus “@@@” in each of the three colours 10 times each (90 trials). Two main findings are of interest. First, performance should speed up across blocks, particularly in the initial practice phase. Second, the contingency effect should either: (a) remain relatively stable, or (b) increase slightly with practice.

The results of Simulation 2 are presented in Fig. 3. First, note that performance improves with practice, particularly in the initial practice phase where mean cycle times were 413, 383, and 374 cycles. This replicates our observations. Secondly, the contingency effect increases slightly early on in the learning phase, with contingency effects of 6, 24, and 32 cycles in the first three blocks, but is relatively stable otherwise (peaking at 44 ms in Block 8). This is again consistent with our observations. As a minor note, performance slowed between the end of the practice phase and the beginning of the learning phase. This was also observed in the original dataset, probably due to the increased number and complexity of the distracting stimuli.

6. Simulation 3: Proportional retrieval

Schmidt and De Houwer (2016a) conducted two experiments to evaluate possible mechanisms producing contingency effects. Here we consider their Experiment 1. In Simulation 4, we simulate their Experiment 2. The design of Experiment

Table 1
Simulation 2 stimulus frequencies.

Colours	Words		
	Search	Choose	Drive
Purple	8	1	1
Orange	1	8	1
Grey	1	1	8




Fig. 3. Simulation 2 cycle times of initial practice trials (left) and the following high and low contingency trials across (right) across training blocks, with original data (Schmidt & De Houwer, 2016b, Experiment 2).

Table 2
Simulation 3 stimulus frequencies.

Colours	Words		
	Give	Hear	Make
Purple	9	1	5
Orange	1	9	5
Grey	5	5	5

1 is presented in Table 2. Two words were presented 9 of every 15 presentations in one colour, 5 of 15 presentations in another colour, and 1 of 15 presentations in a third colour. The remaining word was presented 5 times each per 15 presentations in all three colours. With these contingencies, 400 trials were selected randomly with replacement. The manipulation produces five trial types. The word-colour combinations on *high contingency* trials occur frequently (e.g., “hear” in orange), and infrequently on *low contingency* trials (e.g., “hear” in purple). The remaining three trial types are all *medium contingency* (i.e., medium frequency pairings). Though performance on the three medium contingency trial types did not differ, the conditions do have subtle differences. Specifically, on *biased word* trials (e.g., “hear” in grey) the distracting word is normally predictive of a specific colour (e.g., “hear” is most predictive of an orange response). On *biased colour* trials (e.g., “make” in orange), the *colour* is predictive of the likely word (e.g., orange is normally presented with the word “hear”). On *unbiased* trials (e.g., “make” in grey), neither the colour nor the word are predictive.

In the original study, all three medium contingency trials produced similar performance, whereas high contingency trials were faster and low contingency trials slower. This was taken as evidence that episodic retrieval biases each response proportionally to the frequency with which a presented word co-occurs with the response. For instance, presentation of “hear” will lead to strong facilitation of the orange response, weak facilitation of the purple response, and intermediate facilitation of the grey response. In contrast, presentation of the word “make” will lead to intermediate facilitation of all three responses. The results of Simulation 3 are presented in Fig. 4. As can be seen, the model replicates the predicted pattern of means: high contingency trials are the fastest (372 cycles), low contingency the slowest (402 cycles), and the three medium contingency conditions are intermediate and roughly equivalent (385 cycles).

7. Simulation 4: Frequent, infrequent, and novel items

In a second experiment, Schmidt and De Houwer (2016a) further explored the idea that episodic retrieval biases might be inherently facilitative in nature. That is, not only does correct response prediction facilitate responding, but *incorrect* response prediction does not impair performance. For instance, if “think” is presented highly frequently in purple and infrequently in orange, then presentation of “think” will not impair an orange response. In fact, it might facilitate orange slightly




Fig. 4. Simulation 3 cycle times of high contingency, low contingency, and three types of medium contingency trials, with original data (Schmidt & De Houwer, 2016a, Experiment 1).

Table 3
Simulation 4 learning phase stimulus frequencies.

Colours	Words		
	Think	Find	Search
Purple	40	5	5
Orange	5	40	5
Grey	5	5	40

(purple much more, of course). In support of this, they found that low contingency trials were responded to (slightly) faster than once-presented novel-word trials, not slower. The design of this experiment is presented Table 3. In the first learning block, 3 of 33 words were presented 40 of 50 times in one colour, and 5 of 50 times in each of the remaining two colours (150 trials). In the second test block, the same stimuli were presented again in addition to novel word trials. Each of the remaining 30 words was presented only once in one of the three colours (180 trials).

The results of Simulation 4 are presented in Fig. 5. As in the original study, a standard contingency learning effect (i.e., low contingency – high contingency) is observed in both blocks. More critical are the novel item trials. Unsurprisingly, cycle times were faster on high contingency trials (381 cycles) than on novel trials (405 cycles). As predicted, low contingency trials (402 cycles) were *also* responded to faster than neutral trials, though not by much. This confirms that the most probable (but incorrect) response does not compete with the correct response on low contingency trials in the model, as observed in the participant data.

8. Simulation 5: Stimulus-response binding

In Simulation 5, we model Experiment 1 of Frings et al. (2007). In this experiment, the letters D, F, J, and K were presented as both target and distracter stimuli. Targets were presented in the second and fourth position of a five letter string (e.g., K in FKFKF), with the remaining three letters as distracting flankers. The flankers and targets were never congruent. Trials were presented in pairs. First, a prime trial was presented, but is not of interest. Next, a probe trial was presented. The four dif-




Fig. 5. Simulation 4 cycle times of high and low contingency trials during training (left) and high contingency, low contingency, and novel word trials during test (right), with original data (Schmidt & De Houwer, 2016a, Experiment 2).

Table 4
Simulation 5 conditions with example stimuli.

Trial type	Prime	Probe
<i>Response repetition</i>		
Distracter repetition	DFDFD	DFDFD
Distracter alternation	DFDFD	JFJFJ
<i>Response alternation</i>		
Distracter repetition	DFDFD	DKDKD
Distracter alternation	DFDFD	JKJKJ

ferent types of trials this produces are presented in Table 4. When the target (and therefore response) repeated from prime to probe, the distracter could either repeat (complete repetition) or alternate (partial repetition). Repetition of the distracter benefitted performance in this case. When the target (and response) alternated (i.e., changed) from prime to probe, the distracter could again repeat (partial repetition) or alternate (complete alternation). In this case, repetition of the distracter *impaired* performance. In this experiment, the prime distracter never matched the probe target and the prime target never matched the probe distracter. Similar to the original report, each simulated “participant” was presented 160 prime-probe trial pairs (i.e., 320 individual trials). Approximately 40 of each trial type were presented to each participant, selected randomly with replacement from all the possible stimulus combinations.

Note that, as in the original report, to have an equal number of each trial type, complete repetitions had to be presented more often than would be expected by chance. For instance, if DFDFD is the prime, then there is only one possible stimulus that can serve as a complete repetition probe (i.e., DFDFD), whereas there were always two possibilities for the remaining conditions. Note, however, that: (a) the same binding effects are observed without this transitional frequency confound (e.g., Giesen, Frings, & Rothermund, 2012), and (b) the PEP model is (currently) unable to learn regularities *across* trials (see Discussion). Thus, this potential bias is not problematic for the original finding or the current simulation.

In the PEP model, both the facilitative effect of distracter repetitions on response repetitions and the interference effect on response alternations should replicate. Retrieval of the just-encoded stimulus-response pair in the prime should benefit performance when the same stimulus-response pair is again presented in the probe. When the response changes, however, repetition of the distracter will lead to retrieval of the just-encoded stimulus-response pair. This will bias retrieval in favour of just-made response, which is not the correct response on the current trial. Note that there is no response competition between the retrieved response and other responses in the PEP model. The impairment for partial repetitions is caused by proportional retrieval (see Simulation 3): to the extent that a distracter repetition biases retrieval of the just-encoded response, retrieval for other responses (including the correct response) will be decreased. This might reasonably be called “retrieval interference” rather than “response competition.”

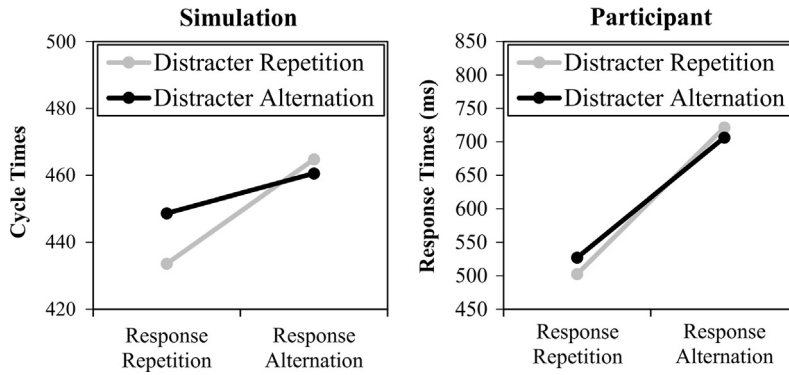


Fig. 6. Simulation 5 cycle times for the orthogonal combination of distracter and response repetitions, with original data (Frings et al., 2007).

The results of Simulation 5 are presented in Fig. 6. As can be seen, the critical effects from the original report are replicated. First, repetition of the distracter sped performance on response repetitions (15 cycles). Second, repetition of the distracter *slowed* performance on response alternations (4 cycles). The interference effect was smaller than the facilitative effect, consistent with the participant data. It might be noted that the repetition cost is quite small. In participant data, this is also generally the case, with the repetition benefit consistently large and the repetition cost often small and non-significant. Also notable, overall performance was faster on response repetitions (441 cycles) than response alternations (463 cycles), consistent with the participant data. The overall main effect of response repetitions was a fair amount smaller in the simulated data, however, possibly because PEP model cycle times are generally on a smaller scale than participant response times.

9. Simulation 6: Mixing costs

For our last simulation, we turn to mixing costs. In particular, we simulate Experiment 1 of Los (1999a). Stimuli consisted only of target stimuli (i.e., no distracters), which were the digits 2, 3, 4, and 5. Letters were either normal or degraded. Degraded stimuli had missing segments in one group of participants and added noise for other participants. However, performance in these two conditions was similar and for modelling purposes this distinction is less important. Thus, we consider the simple comparison between normal and degraded stimuli. There were three types of blocks in the experiment: pure easy (i.e., all normal digits), pure hard (i.e., all degraded digits), and mixed (50% of each item type). Each of these blocks contained 64 trials, with 16 presentations of each digit. In the participant data, performance was overall slower in the mixed block, but this was especially the case for the normal (easy) items.

In the model, degraded stimuli were modelled by simply reducing the *signal* value from 1 to 0.65 (see Appendix A). This means that the model will detect degraded stimuli easily, but only more slowly. Purely for simplicity, the three list types were each simulated separately. This obviates the need to program, run, and compile together multiple counterbalancing orders of blocks. It is also similar to some investigations of mixing with between group manipulations (e.g., Strayer & Kramer, 1994a, 1994b) and the model by definition uses only the most recently-occurring trials, anyway, so contamination across blocks cannot be a problem in the model. The results of Simulation 6 are presented in Fig. 7. As can be observed, a mixing cost was observed: performance was overall slower in the mixed block (464 cycles) than in the pure blocks (458 cycles). Like the participant data, this was especially the case for normal (easy) items (13 cycle effect), and not evident for degraded (hard) items (<1 cycle), producing an interaction between block type (pure vs. mixed) and item type (normal vs. degraded).

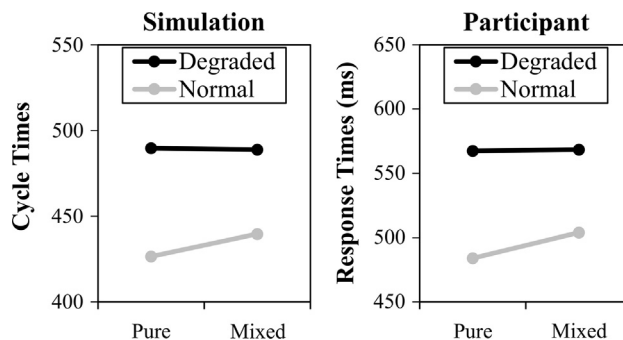


Fig. 7. Simulation 6 cycle times easy and hard items in pure and mixed blocks, with original data (Los, 1999a).

10. Discussion

The present series of simulations aimed to demonstrate how one well-defined, fixed-parameter model can be successfully applied to a range of different phenomena from different literatures. We were able to simulate power curves, the Stroop asymmetry, and the decrease in the congruency effect with practice in Simulation 1. Simulation 2 further simulated the (relatively flat) acquisition curve in the colour-word contingency learning paradigm. Simulations 3 and 4 successfully simulated the relationships between high contingency, medium contingency, low contingency, and novel word trials. Simulation 5 successfully simulated S–R binding effects from the distracter-response binding paradigm. Finally, Simulation 6 successfully simulated a mixing cost. As previously reported (and replicated in the current version of the model; see [Appendix B](#)), the PEP can also simulate a number of important phenomena from the cognitive control literature. Interestingly, as we discuss in more detail in the Conflict Monitoring and Attentional Adaptation section, in the PEP model these latter effects are the incidental result of the very learning processes that we have been discussing (and simulating) in the present manuscript (i.e., rather than an attentional control device, as normally assumed).

Of course, we certainly do not suggest that we have provided a model of cognition that is in any way “complete.” Indeed, the Model Limitations section below discusses further refinements that might be needed to further generalize the model. Importantly, however, we demonstrate that a single fixed model of a small number of processes can have wide applicability to different performance paradigms. The usefulness of this approach is that findings from one domain can often inform another. Indeed, while it might seem inherently more difficult to develop a single model that can be applied to multiple diverse tasks, we observed through testing that many of the parameterisation decisions worked “cooperatively” across tasks. One example of this we have already discussed in Simulation 1: if you model practice-induced benefits in performance, then you get the decrease in the congruency effect across blocks for free. In the following section, we discuss the robustness of the model to parameter changes and the impact of certain parameter decisions on some of the key effects.

10.1. Model robustness

Most of the parameters in the model are not involved in producing the main effects of interest in the paper (see [Table A1](#) in [Appendix A](#)). For instance, the calculations of bias for the Input nodes, the activation function for nodes, the various sources of noise in the model, and the connection weightings and thresholds throughout the algorithmic route of the model are not responsible for the direction of any of the observed effects. Indeed, these parameters have no way of differentiating between conditions (e.g., high vs. low contingency) and therefore are not capable of producing the critical effects of interest. These parameters were set with much simpler goals in mind: cycle times that were of a rough response time-like magnitude, a realistic response time distribution (e.g., see [Schmidt & Weissman, 2016](#)), good accuracy with occasional errors, etc. The actual “work” of model results from the retrieval mechanism. Contingency effects, binding effects, and the various other simulated phenomena result from episodic retrieval. Here, too, there is little room for playing with parameters to observe a given effect. A contingency effect, for instance, is observed with any reasonable parameterization and is always in the same direction. The only parameterizations that fail to produce a contingency effect are those that clearly “break” the mechanism (e.g., no retrieval activation of responses was occurring at all).

The one area where parameter choice does matter more is with the learning rate (i.e., the rate with which older episodes are weakened over time). For instance, stimulus-response binding effects require a high learning rate. That is, if the most-recently encoded event does not have a notably larger effect than older events on retrieval, then the binding effect is lost (or very small). However, this learning rate could not simply be set to explain a single phenomenon, as this one parameter affects multiple simulations. What was particularly interesting is that the relationship between practice curves, binding effects, and contingency acquisition curves worked cooperatively with this parameter choice. After recalibrating the retrieval mechanism to produce larger binding effects (Simulation 5), the effects in the contingency paradigms (Simulations 3–4) emerged more strongly. Additionally, an early issue with the acquisition curves (Simulation 2) was incidentally corrected. In earlier versions of the model, the contingency effect was very large in early blocks, then diminished rapidly (e.g., similar to the Stroop data in Simulation 1). This was corrected after adjusting the retrieval mechanism to better simulate the findings of [Schmidt and De Houwer \(2016a\)](#) and [Frings et al. \(2007\)](#). That is, the same high learning rate needed to simulate binding effects is also needed to simulate a proper acquisition curve. The same is true of the practice curve: if the learning rate is *too* high (such that older events simply do not matter) performance will not continue to improve, and if it is too low performance improvements will be too drastic.

These same changes did not impair the ability of the model to simulate previously reported model results ([Schmidt, 2013a, 2016a](#)), and even boosted the model fit in many cases (esp., the asymmetric list shifting simulations). Similarly, the temporal learning mechanism initially designed to simulate list-level proportion congruent ([Schmidt, 2013c](#)) and congruency sequence effects ([Schmidt & Weissman, 2016](#)) allowed the model to simulate the mixing cost (Simulation 6) without alterations. We believe that these examples are not just lucky coincidences. The closer we get to specifying a cognitive process correctly, the more additional findings should emerge for “free.” In contrast, if a model change made for one effect impairs the ability of a model to explain another effect, then there is probably something fundamentally wrong in the model structure.

It is finally worth noting that the model includes two different prediction mechanisms: one to predict what response to make (contingency learning) and another to predict when to respond. As illustrated in Table 5, each mechanism is responsible for producing different effects (indicated by checkmarks). The same mechanism that produces practice curves, contingency effects, and S–R binding produces item-specific proportion congruent effects (because words are predictive of the congruent response in the mostly congruent condition and of the incongruent response in the mostly incongruent condition), and asymmetric list shifting effects (due to a complicated relationship between task contingencies and the practice benefit across blocks; see Appendix B). The model cannot produce any of these five phenomena if the contingency learning mechanism is lesioned. The model is also able to produce list-level proportion congruent effects, where the critical items are contingency-unbiased, because of the same temporal learning mechanism that produces the mixing cost. The model is also able to produce a congruency sequence effect, even in the version of the paradigm where contingencies and S–R bindings are controlled for by design, due to the same rhythmic timing mechanism acting more locally. The model cannot produce any of these three phenomena if the temporal learning mechanism is lesioned.

10.2. Model comparison

Many of the present series of simulations were highly unique in focus. For instance, Simulation 5 was, to our knowledge, the first simulation of distracter-response binding effects. Similarly, Simulations 2–4 were the first simulations of colour-word contingency learning effects, or any similar “within-trial” learning performance paradigm. Direct comparisons to other models are therefore challenging. We also note that it is not the intent of this paper to argue that our particular episodic memory model is superior to other episodic memory models. Rather, our aim is to demonstrate conceptually how episodic memory biases can have a major impact on a range of performance paradigms. Nevertheless, we here consider some similarities and differences between our model and some of the most related models.

First, we consider the instance model of Logan (1988), which was also used to simulate practice curves. Like all episodic models, this model assumes that each event (e.g., trial) is encoded into a new episode. The primary difference in the model of Logan, however, is in the retrieval mechanism. In the PEP model, multiple episodes are retrieved in parallel and collectively bias responses. This is similar to models of choice behaviour, where each episode “votes” for the response to make (e.g., Howard & Kahana, 2002). This is also similar to how diffusion models (for a review, see Ratcliff, Smith, Brown, & McKoon, 2016) simulate a practice curve with (in addition to other factors) increased accrual of evidence for the likely response with training (Dutilh, Vandekerckhove, Tuerlinckx, & Wagenmakers, 2009; Ratcliff, Thapar, & McKoon, 2006). This is unlike the model of Logan where episodes “race” each other for retrieval and the “winner” determines the response. Practice curves result from this race, because the more episodes there are linking a target to a response the higher the probability that one of the episodes will have been retrieved at time t (i.e., the longer practice has progressed). However, it seems unlikely that this race model of retrieval could simulate binding effects. Each episode in the race is equally weighted, so the just-encoded episode will presumably have no meaningfully large effect on retrieval. Increasing the “running speed” of recently-encoded episodes might be possible, of course, but this would presumably undermine the ability of the model to simulate practice curves: the just-encoded episode would almost always win, meaning that there would be little added gain from older episodes participating in the race.

In contrast, in the exemplar-based random walk model (Nosofsky & Palmeri, 1997) episodes race, but multiple times within each trial. Each race increments evidence for a response (i.e., in contrast to the Logan model, where a single race determines the executed response). “Running rate” can be manipulated in this random walk model (Nosofsky, Little, Donkin, & Fific, 2011), such that the most-recently encoded trial will tend to win most races, but not necessarily every one. Such a random walk model might therefore have similar success as the PEP model in simulating both transient effects (e.g., stimulus-response binding) and longer term learning effects (e.g., colour-word contingency learning effects). Of course, this random walk model is highly similar to the PEP model, with the primary difference being that activation of responses is determined by graded activations from multiple episodes in parallel on each processing cycle of the PEP, whereas multiple episodes contribute to response activation that is graded *across* cycles in the random walk model.

Table 5
Model mechanisms as they relate to simulated effects.

Effect	Contingency	Temporal	Notes
Practice curve	✓		Target-response contingency learning
Contingency effects	✓		Distracter-response contingency learning
S–R binding	✓		Recent distracter-response binding
Mixing cost		✓	Consistent list pace produces pure list advantage
Item-specific PC	✓		Contingency confound produces the effect
List-level PC		✓	Critical items are contingency-unbiased
Congruency sequence		✓	The paradigm used eliminates binding and contingency confounds by design
Asymmetric list shifting	✓		Contingencies and practice produce the effect

Note: PC = proportion congruent.

Another model that shares a lot in common with the PEP model is the Adaptation to the Statistics of the Environment (ASE) model (Jones, Mozer, & Kinoshita, 2009; Mozer et al., 2004; Kinoshita et al., 2008; Kinoshita et al., 2011). Similar to the PEP model, the ASE model learns to anticipate when to respond on the basis of the difficulty of previous trials. As already discussed, in the PEP model this results in rhythmic responding: the model anticipates responding at a similar time as it responded on previous trials. In the ASE model, response time will also be autocorrelated in this manner, but for a more incidental reason: the model aims to initiate responding as quickly as it can without inflating errors. This is related to the notion of response caution (van Maanen et al., 2011). Of course, the PEP model learns both *when* and *what* to respond, giving it wider breadth, but the differences in the temporal learning mechanisms in the two models are interesting. To what extent autocorrelated response times in performance paradigms are due, simply, to our natural tendency to respond rhythmically or to a balancing act between our desire to respond both quickly and accurately is an interesting question for future research.

10.3. Conflict monitoring and attentional adaptation

The PEP model was initially developed as a model of certain findings in the attentional control literature. The proportion congruent effect (Lowe & Mitterer, 1982) is the observation that the congruency effect is reduced when trials are *mostly incongruent* (e.g., 75% incongruent, 25% congruent) relative to when trials are *mostly congruent* (e.g., 25% incongruent, 75% congruent). Typically, this effect is interpreted as resulting from attentional adaptation (Jacoby, Lindsay, & Toth, 1992; Lindsay & Jacoby, 1994). For instance, the highly influential conflict monitoring account argues that detection of frequent conflict in the mostly incongruent condition leads to the adjustment of attention away from the distracting word and/or toward the target colour (Botvinick, Braver, Barch, Carter, & Cohen, 2001).

Another related phenomenon is the congruency sequence (or Gratton) effect (Gratton, Coles, & Donchin, 1992), which is the observation that the congruency effect is reduced when the immediately preceding trial was incongruent rather than congruent. This finding, too, is generally interpreted in terms of conflict adaptation, only of a more transient sort: experiencing conflict on an incongruent trial leads to an adjustment of attention on the following trial. Both effects have been successfully simulated with computational models of conflict monitoring (e.g., Blais, Robidoux, Risko, & Besner, 2007; Botvinick et al., 2001; Verguts & Notebaert, 2008).

However, both effects have been argued to be systematically biased by exactly the sorts of learning biases discussed in the current report (see Schmidt, 2013b; Schmidt, Notebaert, & Van den Bussche, 2015). For instance, proportion congruent effects are often biased by contingencies (Hazeltine & Mordkoff, 2014; Schmidt, 2013a; Schmidt & Besner, 2008) and/or rhythmic response biases (Kinoshita et al., 2011; Schmidt, 2013c). Similarly, congruency sequence effects are typically biased by feature binding biases (Hommel et al., 2004; Mayr et al., 2003; Mordkoff, 2012; Schmidt & De Houwer, 2011) and rhythmic biases (Schmidt & Weissman, 2016). The PEP model is able to simulate both proportion congruent and congruency sequence effects, much like the conflict monitoring model. Unlike the conflict monitoring model, however, the mechanisms responsible for these so-named “attentional control” phenomena are the simple episodic storage and retrieval mechanisms discussed in the present report.

The parsimony of the episodic account is that the only mechanisms we must assume to explain proportion congruent and congruency sequence effects are the very same mechanisms that we must assume anyway to explain various learning and binding phenomena (as confirmed in the present report for the first time). The conflict monitoring account, in contrast, does not share this parsimony: adaptations to conflict in these narrow sets of conflict paradigms can be explained by such models, but explaining the types of (non-conflict) basic learning, binding, and timing phenomena explored in the present report would presumably require appeals to the very learning processes that the PEP model instantiates. Future work might aim to explore whether our simple learning account is sufficient, or whether attentional adaptation to conflict plays an additional role in these phenomena. This is already a heated debate in the attentional control literature (Atalay & Misirlisoy, 2012; Bugg, Jacoby, & Chanani, 2011; Crump & Milliken, 2009; Hazeltine & Mordkoff, 2014; Levin & Tzelgov, 2016; Notebaert & Verguts, 2007; Schmidt, De Schryver, & Weissman, 2014; Schmidt et al., 2015) and the PEP model might serve as a useful reference for how far one can go with episodic learning alone.

It is additionally important to stress that the above discussion is about whether certain phenomena (proportion congruent and congruency sequence effects) exist due to higher-order attentional control processes or lower-level learning biases. The question is not, in contrast, whether attention can be strategically controlled or whether, more broadly, control processes are involved in human cognition. It is obviously the case, for instance, that control processes play an important role in human cognition, including in the paradigms we have simulated here. Researchers do not simply present participants stimuli and “see what they do.” Rather, we instruct participants on how to perform the task, and this will determine how they approach the task (e.g., “do not read the word, name the colour”) and prepare to respond (Cohen-Kadosh & Meiran, 2009). This entails early cognitive control to prepare for the task (Rabbitt, 1997) and is even a fundamental part of the PEP model (i.e., the algorithmic route). Our suggestion is merely that many of the phenomena observed in cognitive paradigms may result from more automated learning processes in operation during task execution (which may, in turn, be instantiated by an initial control process when preparing for the task). Similarly, it is also clearly the case that attention is strategically controllable. Indeed, the ability of participants to perform a Stroop task with reasonably high accuracy is proof enough of this: participants *can* ignore the word and attend to the colour (mostly) successfully, which is undeniable evidence of attentional control. Again, however, this is different than proposing that attention is dynamically adjusted in response to experienced conflict during a Stroop task, as the conflict monitoring account proposes. Thus, our claim is not that cognitive

control does not play an important role in behaviour, but that several important phenomena may not necessarily index conflict adaptation as many assume.

10.4. Model limitations

Further adaptation of the PEP model might be made in future work to expand the breadth of the model. For instance, the present version of the PEP model always responds to target stimuli with a one-to-one stimulus-to-response mapping (e.g., the response to the colour blue is always the blue response). This hard coding is for computational simplicity alone. However, there are many interesting phenomena in which stimulus-response mappings are inconsistent from trial to trial. For instance, in the binding paradigm of [Hommel \(1998\)](#) discussed earlier the response to the first target in each trial pair is cued and unrelated to the stimulus itself. Similarly, in the binding paradigm of [Rothermund et al. \(2005\)](#) participants categorized the colour of words on prime trials, then the word meaning on probe trials. In a similar vein, in task switching experiments participants are asked to respond to different stimulus dimensions on different trials ([Jersild, 1927](#); for reviews, see [Kiesel et al., 2010](#); [Monsell, 2003](#); [Vandierendonck, Liefoghe, & Verbruggen, 2010](#)). Binding effects can also play a major role in task switching experiments ([Goschke, 2000](#); [Logan & Bundesen, 2003](#); [Schmidt & Liefoghe, 2016](#)), so task switching experiments are also relevant for the PEP model. All of these paradigms, however, require a model that allows for stimulus-response mappings that are not one-to-one.

Currently, the model learns the contingencies (or bindings) between stimuli within a given trial (e.g., that “search” tends to be presented in purple). However, learning can also occur across trials. For example, in the serial response time task ([Destrebecqz & Cleeremans, 2001](#); [Nissen & Bullemer, 1987](#)) it is found that participants respond more quickly when responses follow a repeated series than when responses are randomly ordered. Other computational models have been forwarded that can simulate sequential learning and related phenomena (e.g., artificial grammar learning; [Reber, 1967](#)). For instance, the Serial Response Network (SRN; [Cleeremans & McClelland, 1991](#); [Elman, 1990](#)) takes both the current and previous trial (via recurrence) as input on every simulated trial to learn sequences, and the TRACX model ([French, Addyman, & Mareschal, 2011](#)) extracts “chunks” of information from a temporary store of episodes to learn sequential regularities. Though none of the paradigms modelled thus far with the PEP model contained any sequential regularities, presumably some adaptation like this might prove highly effective in explaining both “within-trial” learning (e.g., in the colour-word contingency learning paradigm) and sequential learning with the same encoding and retrieval mechanisms.

In the present version of the model, stimulus, timing, and response information are all stored in episodes. Presentation of a stimulus leads to episodic retrieval of the contingent responses. However, presentation of one stimulus (e.g., distracting word) does not lead to retrieval of other associated stimuli (e.g., target colour). Whether this should change in future instantiations of the model is uncertain. For instance, the results of some dissociation procedures suggest that contingency learning in performance tasks might be exclusively due to stimulus-response learning ([Schmidt & De Houwer, 2012b](#); [Schmidt et al., 2007](#)). Some findings in the binding literature also seem compatible with this (see especially, [Moeller and Frings \(2014\)](#); see also, [Frings et al., 2007](#)). In some reports, however, it has been suggested that distracter-target bindings are retrieved ([Colzato, Raffone, & Hommel, 2006](#); [Giesen & Rothermund, 2014](#); [Hommel & Colzato, 2009](#)). More experimental research might be needed to clarify this issue further, but adaptations might be made to the PEP model to simulate distracter-target retrieval effects if deemed appropriate.

As a further caveat, the current version of the PEP does not link semantically-related targets and distracters in memory. For instance, the distracting word “blue” will not retrieve episodes coding for the target colour blue (or vice versa). These sort of bindings have been argued to have important effects on behaviour, for instance, in negative priming ([Rothermund et al., 2005](#)) and congruency sequence effects ([Mordkoff, 2012](#); [Schmidt & De Houwer, 2011](#)). One way to adapt the PEP model for such bindings might be to link Identity nodes to episodes. Via these nodes, semantically-related distracters and targets would (partially) activate the same episodes.

10.5. Model strengths

In the preceding section, we have elaborated on some of the limitations of the current instantiation of the PEP model. In the current section we discuss the strengths of the model. The greatest strength of the model is in the demonstration that a small number of assumptions about memory storage and retrieval are sufficient to explain a broad range of observations. Consider first Simulations 1–5. Here, we showed that practice curves, diminishing Stroop effects across blocks, contingency learning effects, and S–R binding effects could all be simulated with episodic storage and retrieval. More boldly, the current model suggests that all of these individual phenomena result from the same underlying process, only in different ways. For instance, consider practice curves and the diminished Stroop effect across blocks. The latter is merely a consequence of the former in the current model: both congruent and incongruent response time improve exponentially toward asymptote (practice), which results in smaller and smaller differences between the two trial types across blocks. Similarly, while practice effects may, more conceptually, be due to “target-response” bindings and contingency effects are due to “distracter-response” bindings, there is actually no difference in how the model retrieves these two different “types” of information. Activated episodes facilitate the response that they are connected to in either case. Indeed, the episodes do not even “know” whether they are predicting the response on the basis of the colour or on the basis of the word (or both). Similarly, S–R binding effects are a mere consequence of the learning process. In order to have a high learning rate, recent events must have a

larger impact than more distant events, and this alone is enough to explain why recent stimulus-response co-occurrences produce S–R binding effects. This contrasts with some views that propose that, though related, learning and binding effects are fundamentally different (e.g., Colzato et al., 2006). It is possible that the view espoused in the present manuscript is a little *too* parsimonious, and future research might aim to test to what extent learning, binding, and practice curves are due to the same process and to what extent they differ.

The integration of temporal information to an episodic binding model is a further strength of the PEP model. As we have outlined, rhythmic responding biases can also be explained by episodic storage and retrieval. As fundamental as learning *what* response to make in a given context is the knowledge of *when* to make it. A musician cannot learn how to play a song, for instance, by learning the series of notes alone; the duration and temporal spacing of said notes is equally critical. The same is true for any complex behaviour involving a sequence of actions. As another example, when solving a Rubik's cube a speedcuber will make rapid bursts of moves in quick succession (e.g., right face clockwise, up face clockwise, right face counter-clockwise, then up face counter-clockwise). If one of the moves in the series is made too soon or too late, the cube will “catch” (i.e., get jammed) or, more catastrophically, “pop” (i.e., pieces “explode” out of the cube). It is precision in both the actions and their timing that allows a speedcuber to perform the series of four moves mentioned above in a few hundred milliseconds. This same type of temporal learning can explain proportion congruent effects (Schmidt, 2013c), congruency sequence effects (Schmidt & Weissman, 2016), and mixing costs (Simulation 6).

The conceptual core of the model is its reliance on episodic storage and memory retrieval processes that are applied to a wide range of phenomena, some of which have traditionally been explained by reference to other processes (e.g., selective attention, cognitive control, speed/accuracy trade-offs). Still, the potential of the model is far from being fully exploited. Another broad class of effects that could easily be modelled by an extension of the PEP model are what have been labelled “action effects” in the literature (Hommel, Musseler, Aschersleben, & Prinz, 2001). That is, the consequence (C) of an action will typically also become part of an S–R episode, transforming it into an S–R–C episode. By extending the range of action consequences to affective and emotional action effects (Eder, Rothermund, De Houwer, & Hommel, 2015), the model can also be used to provide explanations of what has traditionally been treated in the domain of learning proper (i.e., operant conditioning).

Besides being able to simulate a wide range of phenomena from the literature, the episodic model also generates novel predictions that can be evaluated in future research. Essentially, the PEP model predicts retrieval effects that are based on a repetition of stimulus features that then leads to a retrieval of previous episodes. A straightforward test of the validity of the model thus consists in comparing repetition trials with a neutral baseline without stimulus repetitions. Although this principle has already been applied in some of the paradigms that were mentioned (e.g., S–R binding; see also, Simulation 4), the inherent potential has not yet been fully utilized. Further applications of this idea are possible, for instance, in investigating mixing costs.

Similarly, seemingly more technical features of the model can be used to derive novel and possibly surprising empirical predictions. For instance, consider the general notion that effects should reduce with continued practice, which is explained in the model with a strengthening of connections between input nodes and episodes. This notion could be tested further in an experiment by intermixing probe trials for which no previous episodes have been created (i.e., using new stimuli). As an example, what would happen in a Stroop task if, after extensive training with the words “red,” “blue,” and “green” printed in red, blue, and green, new colours and colour words were added, such as “yellow,” “orange,” and “purple” in yellow, orange, and purple? It might be predicted that for these new stimuli the Stroop effect would increase back up to the original effect size observed at the start of the experiment. After all, participants have extensive practice in identifying red, blue, and green stimuli, but have no practice in identifying yellow, orange, and purple stimuli (i.e., within the context of the experiment). Such work could inform us to what extent Stroop (or other) effects decrease with time due to item-specific episodic learning, and to what extent these decreases might be due to more general processes (e.g., general preparedness, familiarity with the task structure, or practice in ignoring words).

The retrieval of temporal information to learn when to respond and the temporary lowering of the response threshold at the anticipated interval for responding is a further interesting feature of the PEP model that has a lot of interesting and (mostly) untested implications. As currently programmed, the mechanism is quite simple, sampling only a few (five) of the most recently-encoded episodes, but further application of the retrieval principals for *response* retrieval in the model to the temporal expectancy mechanism are possible. For instance, if the expected time to respond is determined, instead, by the very same episodes that are retrieved to anticipate the likely response, then we might anticipate that temporal expectancies will be largely determined by which stimuli are presented. If the model is presented with the colour blue and the last response to a blue stimulus was made quickly, then the model might anticipate another quick response (i.e., relative to the case where the last blue stimulus was responded to slowly). There is already some (albeit marginal) evidence for partial item-specificity in rhythmic responding in Schmidt (2014), and further evidence that rhythmic responding can be context-specific (e.g., two different rhythms for two different display locations) in Schmidt, Lemerrier et al. (2014). Further experimental and computational modelling work in this domain could prove highly informative.

10.6. Concluding remarks

The present work aimed to achieve two things. More narrowly, we aimed to show how one model of episodic storage and retrieval can be highly effective in explaining behaviour in a range of performance paradigms. Including the backwards

compatibility checks (Appendix B), twelve different experiments were simulated with one model parameterisation, spanning across the learning, timing, practice, binding, and cognitive control domains. More globally, we aimed to stress the utility of modelling work that focuses on broader-picture processes (Hommel & Colzato, 2015), rather than narrowly focuses on a small set of observations within a particular subfield. No model is perfect, of course (e.g., we are sure that the reader could easily imagine some data for which the current version of the PEP model will likely fail to simulate), but models with explanatory power over a range of datasets and subfields might serve to defragment the splintered subfields in cognition and promote a focus on the bigger picture.

Author notes

This research was supported by a postdoctoral researcher Grant (121184N) of the Fonds Wetenschappelijk Onderzoek (FWO – Vlaanderen) to James R. Schmidt, Grant BOF16/MET_V/002 of Ghent University to Jan De Houwer, the Interuniversity Attraction Poles Program initiated by the Belgian Science Policy Office (IUAPVII/33), and Grant DFG RO 1272/6-2 from the Deutsche Forschungsgemeinschaft to Klaus Rothermund.

Appendix A. Model description

Since the original publication of the model (Schmidt, 2013a), various minor adjustments have been made to the PEP framework (see Schmidt, 2013c, 2016a; Schmidt & Weissman, 2016). In the current work, however, major adjustments to the model are made (while maintaining “reverse compatibility” with previous simulations) in order to simulate a wider range of phenomena. Below, we describe the math of the model.

A.1. Basic formulas

On every cycle (i.e., simulated millisecond) of the model, the *activation* state of each node i is updated with Formula (1),

$$activation_i = activation_i(1 - decay) + input_i(decay) \quad (1)$$

where $input_i$ is the incoming activation to the node, and $decay$ is 0.01. Note that in this formula the $activation_i$ value on the right hand of the formula is the value before applying the formula, and the value on the left hand of the formula is the value after applying the formula. The same general rule applies to similar formulas below. With this formula, the *activation* state of the node will slowly increase or decrease toward the *input* value. The *input* value is determined by first summing all the positive and negative *incoming* inputs to the node, using Formula (2),

$$incoming_i = \sum_{j=1}^n (output_j(weight_{ji})) + noise \quad (2)$$

where $output_j$ is the amount of activation being sent from node j and $weight_{ji}$ is the connection strength between j and i . The *noise* parameter is calculated from a Gaussian distribution with a mean of 0.5 and a standard deviation of 0.2, restricted between 0 and 1, then scaled by 0.5 for Input nodes, and 0.01 for Identity and Response nodes. Purely for performance (i.e., simulation time) purposes, no noise was added to Episode nodes. Note that the model still functions the same with noise in the Episode nodes, but it greatly slows down simulation times to include such noise. The result of Formula (2) is then inserted in a logistic transform, as computed by Formula (3),

$$input_i = \frac{1}{1 + e^{-(incoming_i)(8)-4}} \quad (3)$$

Thus, Formula (3) returns a value between 0 and 1 that grows in a sigmoid (S-shaped) function. That is, *input* increases minimally with smaller *incoming* values, rapidly with intermediate values, and only slightly more rapidly with larger values. This logistic transform (often referred to as the “softmax” function in neural net research) has previously been demonstrated to be important for simulating the asymmetry in the Stroop congruency effect (see Simulation 1; Cohen et al., 1990), but is new to the PEP model. Note that this function is identical to that in *Parallel Distributed Processing (PDP) models*, only on a different scale (i.e., minimum, middle, and maximum values of 0, 0.5, and 1 in the PEP and –4, 0, and 4 in the PDP).

A.2. Input nodes

There are two sets of Input nodes, one for the target stimuli and another for distracting stimuli. The *incoming* activation to Input nodes is computed with Formula (4),

$$incoming_i = bias_i + noise \quad (4)$$

where $bias_i$ is the current cycle bias toward the stimulus and *noise* is a random number added to each cycle. At the start of each trial, *bias* is set for each node from a Gaussian distribution with a mean of 0.4 and standard deviation of 0.16, restricted between 0 and 0.8. On each cycle, *noise* is computed as a Gaussian number with a mean of 0.25 and standard deviation of 0.1,

restricted between 0 and 0.5. The above math for *bias* setting is slightly different than in previous models, but practically little different. On each cycle, *bias* is adjusted toward a *signal* value, computed with Formula (5),

$$bias_i = bias_i(1 - momentum) + signal_i(momentum) \quad (5)$$

where $signal_i$ is 1 for presented items (except the degraded stimuli in Simulation 6, where this value is 0.65) and 0 for unrepresented items and $momentum$ is 0.004. The purpose of this calculation (e.g., rather than just setting input to 1 for presented items and 0 for unrepresented items) is to add realistic noise to the model. With these formulas, the model can begin a trial with a bias for the wrong stimulus, but will eventually discover the correct stimulus given enough time. This makes for a more appropriate response time distribution (ex-Gaussian) and the occasional error. Activation exceeding a *threshold* of 0.5 is passed on to Identity and Episode nodes, described later.

For distracter nodes, the *signal* value is further modified by the “attentional wandering” mechanism described in Schmidt and Weissman (2016). In particular, the *signal* value of a presented distracter is not fixed at 1, but determined on each trial with Formula (6),

$$signal_i = signal_i(0.9) + signal_{base}(0.1) + deflection_n \quad (6)$$

where $signal_{base}$ is 1. The *deflection* for trial n is computed with Formula (7),

$$deflection_n = deflection_{n-1}(0.95) + random_n(0.05) \quad (7)$$

where $random_n$ is a roughly normal number between -1 and 1 computed as one random number subtracted from another. Note that this is simply “slow wave” noise. That is, the strength of the word varies randomly, but this random variation is correlated across trials. This is only useful for more advanced sequential analyses investigating correlations between previous trial and current trial RT (see original report) and is unrelated to any of the mean RT analyses investigated in the current report.

A.3. Identity nodes

Identity nodes receive input from target Input nodes with a connection *weight* of 3.2. Distracter Input nodes also send input to Identity nodes, at a rate of 1.8, but only if these correspond to potential responses (e.g., in the Stroop task). Identity nodes are also inhibited by the other Identity nodes with a connection *weight* of 0.5. This is where congruency effects emerge in the model. Activation exceeding a *threshold* of 0.5 is passed on to Response nodes, using Formula (6),

$$output_i = (activation_i - threshold)(weight_{ij}) \quad (8)$$

This is actually the same formula used for all node types. The only thing unique to Identity nodes (i.e., relative to other node types) is that, as previously described (Schmidt, 2016a), the connection *weight* between Identity and Response nodes strengthens over time. In the current version of the model, this computation is simplified by being computed only once per trial at the time of responding (rather than continually on every processing cycle), which is also probably more plausible. Specifically, the *weight* for each connection between Identity node i and (correct) Response node j starts at 1.12, then is strengthened each time the correct response j is made with Formula (7),

$$weight_{ij} = weight_{ij}(1 - change) + 1.33(change) \quad (9)$$

where *change* is 0.005. As such, $weight_{ij}$ approaches 1.33 over time. This models a general strengthening of the algorithmic route over time. Along with episodic retrieval, this adjustment helps explain practice curves.

A.4. Response nodes

Response nodes receive input from both Identity nodes (described above) and Episode nodes (described below). Retrieval from episodic memories was (slightly) simplified. Specifically, *retrieval* (input from episodes) simply equaled the sum of all inputs minus 0.1 (subtracted to wipe out low level activations, thereby increasing the “spread” in the amount of retrieval for high versus low contingency responses) with the following constraints: (a) if the total amount of retrieval activation for all responses exceeded 1, then the *retrieval* input for each individual response was first divided by the total amount of retrieval activation for all responses (i.e., proportional retrieval), and (b) the resulting *retrieval* value was constrained between 0 (i.e., to prevent retrieval inhibition) and 0.15 (to prevent the model from responding prematurely). The model responds when *activation* exceeds the response *threshold* for one of the responses. The default *threshold* is 0.5. However, this value can be decreased by the “temporal expectancy” mechanism first described by Schmidt (2013c), then adjusted by Schmidt and Weissman (2016), as will be explained in the following section.

A.5. Episode nodes

Most important to the model are Episode nodes, and these received the largest overhaul in Version 2.0 of the model. The most novel change is that both the distracting and target stimuli are recorded into the episode. Specifically, on each trial a

new Episode node is created. The targets, distracters, and responses j are recorded into the episode i on each cycle using Formula (8),

$$weight_{ij} = weight_{ij}(1 - write) + incoming_i(write) \quad (10)$$

where *write* is 0.01 and *incoming_i* is the amount of activation being received by the Input node or the recording rate of the Response node. Note that *weight_{ij}* begins at 0 at the start of the trial. For Input nodes, *incoming* activation is the *activation* exceeding *threshold* multiplied by 5 for distracters and 2 for targets. For Response nodes, *incoming* activation is the *activation* of the Response node multiplied by 1. Note that the threshold is not subtracted from response activation, because (as currently programmed) response activation stops growing once the threshold is reached. For Input nodes, subtracting the threshold wipes out the noisy *bias* activation for unrepresented stimuli. Presumably, this could be reprogrammed to use the same formula for Input and Response nodes (e.g., by allowing response activation to continue growing after response selection, or by simply recording the highest-active stimulus/response) and might be adapted in the future. This detail seemed of relatively minor importance to us, however. The key point is that the identified stimuli and executed response get encoded. Thus, any sufficiently activated node will be recorded in the Episode, but typically only the presented stimuli and the response made will be strongly connected. In addition to input from stimuli, Episode nodes receive a constant input of 0.05 during the trial in which it is being created (to ensure it is active).

On subsequent trials, *activation* exceeding a *threshold* of 0.02 activates connected Response nodes at a *strength* of 1. This *strength* value is reset to 0.02 for episodes in which an error was made. In the previous versions of the model, connection *weights* were reduced on every cycle proportional to the amount of retrieval. In the current version of the model, this was simplified slightly by reducing the *weights* between an Input or Response node j and the Episode node i once at response, using Formula (9),

$$weight_{ij} = weight_{ij}(1 - activation_i(loss)) \quad (11)$$

where *activation_i* is the *activation* of the Episode node and *loss* is 0.25. That is, the connections an episode has to other nodes are weakened proportionally to how active the Episode node was. This is needed for a high learning rate, as discussed in the Introduction. Also, failing to weaken some connections as you strengthen others is fatal in any learning model, and, as Gerstner and Kistler (2002) point out, a good learning mechanism should explain such loss as an integral part of the learning mechanism itself. Thus, older and older episodes become less and less influential as new ones are encoded. Older episodes contribute to a “cumulative” learning effect, however, which is important for some of the simulations reported in the current manuscript (e.g., Simulations 2–4).

Episode nodes also record the response time, which is used for anticipating when to respond (Schmidt, 2013c; Schmidt & Weissman, 2016). This mechanism is unchanged from the last version of the model, but described again here. On each trial, the most recently-encoded episodes are retrieved and used to bias the *threshold* for Response nodes. In particular, the *threshold* was determined on each processing cycle with Formula (10),

$$threshold = baseline - \left(\sum_{i=1}^n (proximity_i \times strength_i) - 0.05 \right) \quad (12)$$

where *baseline* is the maximum threshold of 0.5 and the result of the formula is restricted between 0.25 and 0.5. The *proximity* value for each episode i is calculated with Formula (11),

$$proximity_i = 1 - \left(\frac{(cycle - rt_i)^2}{10,000} \right) \quad (13)$$

where *cycle* is the current processing cycle and *rt_i* is the stored response time. With this formula, the *threshold* is strongly biased the closer the current cycle time is to the stored response time. Finally, *strength* is determined with Formula (12),

$$strength_i = \left(\frac{(6 - lag_i)^3}{500} \right) \quad (14)$$

where *lag* is how many trials back an episode i was encoded. The result of this formula, only applied to the most recent five episodes (i.e., because *strength* reaches 0 by the sixth), is that the most recently encoded episode has a much larger effect on the *threshold* than older ones.

A.6. Summary

Table A1 provides a summary of the various parameters in the PEP model. Most of these parameters are not directly relevant for the phenomena simulated in the manuscript. That is, only a few are capable of producing systematic differences between conditions, generally having to do with episodic retrieval processes. Other parameters are only related to overall processing in the model (e.g., general noise, speed of information transmission, etc.).

Table A1
Parameters, weights, and thresholds in the PEP model.

Parameter	Value	Description
<i>Activation function</i>		
Decay	0.01	Determines how fast activation accrues
Noise	$x = 0.25^a$ (Input), $x = 0.005^a$ (Identity, Response)	Random cycle-level noise
<i>Input</i>		
Bias	$x = 0.4^a$	Initial random bias toward input, produces more realistic response time distributions
Signal	1 (presented target), ~ 1 (presented distracter), 0 (unpresented)	Signal for presented and unpresented items, allows stimulus identification, makes distracter signal more variable
Momentum	0.004	Rate presented item dominates input, produces more realistic response time distribution
<i>Weights</i>		
Input-identity	3.2 (target), 1.8 (distracter)	Input-to-Identity connect weights, stronger to the (attended) target
Identity (within)	-0.5	Within-layer inhibition, produces main effect of congruency
Identity-response	1.12–1.33	Connections to responses, strengthens slightly with time
<i>Thresholds</i>		
Input, identity	0.5	Threshold to pass on activation
Response	0.25–0.5	Varies with temporal expectancies
Episode	0.02	Threshold to activate responses
<i>Episode</i>		
Write	0.01	Rate stimuli/responses written to memory
Strength	5 (distracters), 2 (targets), 1 (responses)	Episode search rate (output multiplier)/Encoding rate
Loss	0.25	Rate old memories are forgotten, critical for setting the learning rate of the model
Retrieval	0–0.15	Amount of input to Response node from episodes, maximum value important for preventing premature responding
Proximity	0–1	Increases as anticipated time approaches actual time, decreases response threshold
Strength	0–0.25	Amount of influence of episode on temporal expectancy, stronger for recent memories

^a All random variables are computed with a standard deviation equal to 40% of the mean, and are restricted between ± 2.5 standard deviations from the mean (i.e., between 0 and $2x$).

Appendix B. Cognitive control simulations

Here, we briefly consider the cognitive control paradigm simulations as backward compatibility tests to previous modelling papers (Schmidt, 2013a, 2013c, 2016a; Schmidt & Weissman, 2016). We provide only brief summaries of the experiments modelled and the relation to the novel simulations in the present report. Full details of the simulations can be found in the original report.

B.1. Simulation 7: Item-specific proportion congruent effect

In the item-specific proportion congruent paradigm (Jacoby, Lindsay, & Hessels, 2003), some colour words are presented most often in the congruent colour (e.g., “blue” 75% in blue) and other colour words are presented most often in an incongruent colour (e.g., “green” 75% in red). Because the items are randomly intermixed, a participant cannot know whether the word is mostly congruent or mostly incongruent until the word identity is known. Despite this fact, it has been argued that participants adjust their attention to the word on the basis of the identity of the word rapidly. Alternatively, the effect might be simply due to participants learning the contingencies between the presented word and the likely response (e.g., that “green” probably indicates a red response), as argued by Schmidt and Besner (2008). Thus, the PEP model produces an item-specific proportion congruent effect with the *exact same* mechanism responsible for practice curves, colour-word contingency learning effects, and binding effects. The simulated RTs are presented in Fig. B1. As in the original simulation (Schmidt, 2013a), the congruency effect is largest in the high proportion congruent (mostly congruent) condition, smallest in the low proportion congruent (mostly incongruent) condition, and intermediate in the medium proportion congruent (chance) condition.

B.2. Simulation 8: List-level proportion congruent effect

A list-level proportion congruent effect is a proportion congruent effect for contingency-unbiased *diagnostic items* that are intermixed with manipulated (contingency-biased) *inducer items*. In particular, an experiment of Hutchison (2011) was

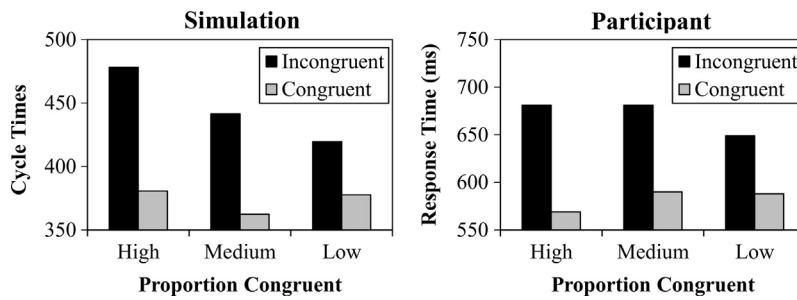


Fig. B1. Simulation 7 cycle times for congruent and incongruent items of high, medium, and low proportion congruency, with original data (Jacoby et al., 2003).

simulated. Obviously a proportion congruent effect for diagnostic items cannot be due to contingency learning. Instead, the PEP model simulates this effect with the temporal learning mechanism: the task pace is faster in the mostly congruent condition (working to the advantage of congruent trials) and slower in the mostly incongruent condition (working to the advantage of incongruent trials). Thus, the list-level proportion congruent effect is produced by the exact same mechanism that produces the mixing cost (and congruency sequence effect in the next section). The simulation results are presented in Fig. B2. As in the original simulation (Schmidt, 2013c), the congruency effect is larger in the mostly congruent condition than in the mostly incongruent condition.

B.3. Simulation 9: Congruency sequence effect

The congruency sequence effect is the finding that the congruency effect is reduced following an incongruent trial relative to following a congruent trial (Gratton et al., 1992). Typically, this has been interpreted in terms of adaptation of attention away from the word following a conflicting incongruent trial (e.g., Botvinick et al., 2001). Although stimulus-response binding (Hommel et al., 2004; Mayr et al., 2003) and contingency confounds (Mordkoff, 2012; Schmidt & De Houwer, 2011) have been identified in this effect, in at least some constrained scenarios (e.g., with prime-probe tasks rather than Stroop) a congruency sequence effect can still be observed in the absence of these biases (e.g., Schmidt & Weissman, 2014). However, such “confound-minimized” congruency sequence effects can still be explained by temporal learning. In particular, after making a (relatively fast) congruent response, participants expect to be able to respond similarly quickly on the following trial and are especially prepared to respond at this early time window. Congruent trials benefit from this expectancy, because participants can respond when they were expecting to. Incongruent trials are too difficult to maintain the (one-trial) fast rhythm, however, and the expectancy benefit is lost. The reverse is true following a (relatively slow) incongruent response. Participants expect to respond later on the next trial, which works to the advantage of incongruent trials (where the slower rhythm is maintained), but to the disadvantage of congruent trials (because participants are not prepared to respond as quickly as they otherwise could after a slow response). Thus, congruent trials benefit following a fast congruent response and incongruent trials benefit following a slow incongruent response. Indeed, trial-by-trial analyses confirm the expected relationship between previous trial response speed and current trial congruency effects (Schmidt & Weissman, 2016; see also, Kinoshita et al., 2011; Schmidt, 2013c). As such, the PEP model produces this congruency sequence effect using the same rhythmic responding mechanism responsible for mixing costs and list-level proportion congruent effects. The simulation results are presented in Fig. B3. As in the previous report (Schmidt & Weissman, 2016), the congruency effect was smaller following an incongruent relative to a congruent trial.




Fig. B2. Simulation 8 cycle times for congruent and incongruent diagnostic items in the mostly congruent and mostly incongruent condition, with original data (Hutchison, 2011).




Fig. B3. Simulation 9 cycle times for congruent and incongruent items following congruent and incongruent trials, with original data (Schmidt & Weissman, 2014, Experiment 1).

B.4. Simulations 10–12: Asymmetric list shifting effects

The asymmetric list shifting effect is the observation that the congruency effect decreases more when switching from a mostly congruent to a mostly incongruent block than it increases when switching from a mostly incongruent block to a mostly congruent block (Abrahamse, Duthoo, Notebaert, & Risko, 2013). This was originally argued to be due to attention being allowed to the word in an initial mostly congruent block, which is then forced away when switching to a mostly incongruent block. In contrast, when the mostly incongruent block comes first, attention is forced away from the word at the start and does not shift back when switching to a mostly congruent block. However, as Schmidt (2016a) pointed out with reanalyses of the original data and computational modelling results, this analysis is confounded by practice. The congruency effect tends to decrease with practice (see Simulation 1), so this is part of the reason why the decrease from mostly congruent to most incongruent is so large and the increase from mostly incongruent to mostly congruent is so small. All three experiments from the original report were simulated. The first (Experiment 1a in original report) is presented in Fig. B4, where the simple asymmetric switch cost is observed. The second (Experiment 1b) is also presented in Fig. B4, where no list shift was manip-




Fig. B4. (a) Simulation 10 cycle time congruency effects for the MC-MI and MI-MC lists and (b) Simulation 11 cycle time congruency effects for the MC-MC and MI-MI lists, with original data (Abrahamse et al., 2013, Experiments 1a and 1b). Note: MC = mostly congruent, MI = mostly incongruent.

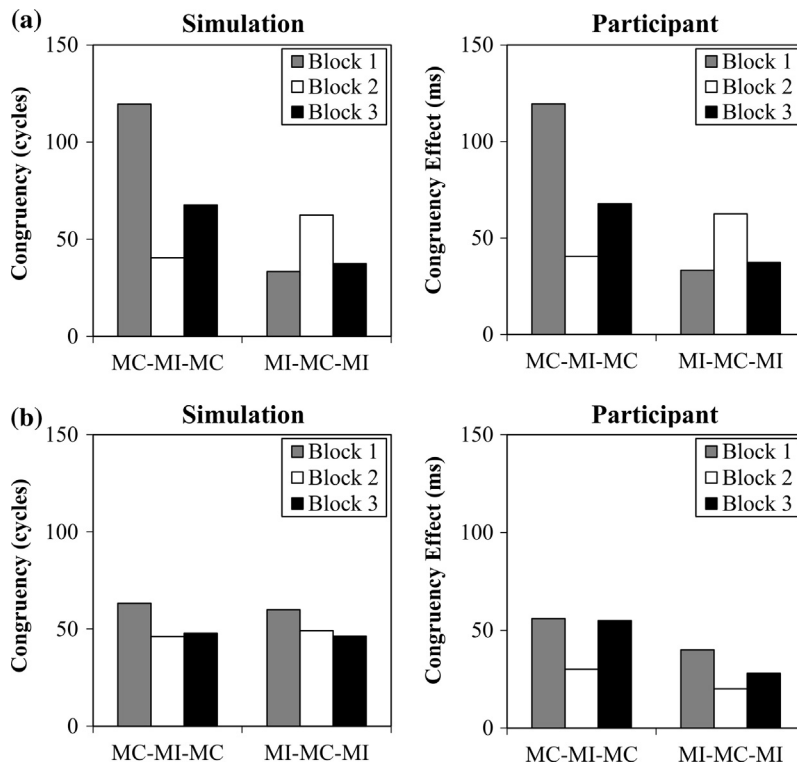


Fig. B5. Simulation 12 cycle time congruency effects for inducer (top panels) and diagnostic items (bottom panels) in the MC-MI-MC and MI-MC-MI lists, with original data (Abrahamse et al., 2013, Experiment 2). Note: MC = mostly congruent, MI = mostly incongruent.

ulated and an overall decrease in the congruency effect with practice can be observed. The final simulation (Experiment 2) is presented in Fig. B5 where proportion congruency is manipulated back and forth more than once and asymmetric effects are again observed for manipulated inducer items (top panels). Un-manipulated diagnostic items were also included in the design, but produced no significant effects in the participant data. The results are nevertheless provided (bottom panels) for information purposes. All of these effects are due to practice and therefore result from the exact same mechanism that produces practice curves, colour-word contingency effects, binding effects, and item-specific proportion congruent effects.

References

- Abrahamse, E. L., Duthoo, W., Notebaert, W., & Risko, E. F. (2013). Attention modulation by proportion congruency: The asymmetrical list shifting effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*, 1552–1562.
- Addyman, C., & French, R. M. (2012). Computational modeling in cognitive science: A manifesto for change. *Topics in Cognitive Science*, *4*, 332–341.
- Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* New York, NY: Oxford University Press.
- Anderson, M. C., Bjork, R. A., & Bjork, E. L. (1994). Remembering can cause forgetting: Retrieval dynamics in long-term memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 1063–1087.
- Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. L. (2004). An integrated theory of the mind. *Psychological Review*, *111*, 1036–1060.
- Atalay, N. B., & Misirlisoy, M. (2012). Can contingency learning alone account for item-specific control? Evidence from within- and between-language ISPC effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *38*, 1578–1590.
- Blais, C., Robidoux, S., Risko, E. F., & Besner, D. (2007). Item-specific adaptation and the conflict-monitoring hypothesis: A computational model. *Psychological Review*, *114*, 1076–1086.
- Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review*, *108*, 624–652.
- Bugg, J. M., Jacoby, L. L., & Chanani, S. (2011). Why is it too early to lose control in accounts of item-specific proportion congruency effects. *Journal of Experimental Psychology: Human Perception and Performance*, *37*, 844–859.
- Church, R. M. (1984). Properties of the internal clock. *Annals of the New York Academy of Sciences*, *423*, 566–582.
- Cleeremans, A., & McClelland, J. L. (1991). Learning the structure of event sequences. *Journal of Experimental Psychology: General*, *120*, 235–253.
- 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: Conscious and nonconscious information processing* (pp. 453–476). Cambridge, MA: MIT Press.
- Cohen-Kdoshay, O., & Meiran, N. (2009). The representation of instructions operates like a prepared reflex: Flanker compatibility effects found in first trial following S-R instructions. *Experimental Psychology*, *56*, 128–133.
- Colzato, L. S., Raffone, A., & Hommel, B. (2006). What do we learn from binding features? Evidence for multilevel feature integration. *Journal of Experimental Psychology: Human Perception and Performance*, *32*, 705–716.

- Crump, M. J. C., & Milliken, B. (2009). The flexibility of context-specific control: Evidence for context-driven generalization of item-specific control settings. *Quarterly Journal of Experimental Psychology*, *62*, 1523–1532.
- Destrebecqz, A., & Cleeremans, A. (2001). Can sequence learning be implicit? New evidence with the process dissociation procedure. *Psychonomic Bulletin & Review*, *8*, 343–350.
- Dignath, D., Pfister, R., Eder, A. B., Kiesel, A., & Kunde, W. (2014). Representing the hyphen in action-effect associations: Automatic acquisition and bidirectional retrieval of action-effect intervals. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *40*, 1701–1712.
- Dulaney, C. L., & Rogers, W. A. (1994). Mechanisms underlying reduction in Stroop interference with practice for young and old adults. *Journal of Experimental Psychology: Learning Memory and Cognition*, *20*, 470–484.
- Dutilh, G., Vandekerckhove, J., Tuerlinckx, F., & Wagenmakers, E. J. (2009). A diffusion model decomposition of the practice effect. *Psychonomic Bulletin & Review*, *16*, 1026–1036.
- Ebbinghaus, H. (1913). *Memory: A contribution to experimental psychology*. New York: Teachers College, Columbia University.
- Eder, A. B., Rothermund, K., De Houwer, J., & Hommel, B. (2015). Directive and incentive functions of affective action consequences: An ideomotor approach. *Psychological Research Psychologische Forschung*, *79*, 630–649.
- Egner, T. (2014). Creatures of habit (and control): A multi-level learning perspective on the modulation of congruency effects. *Frontiers in Psychology*, *5* Article 1247.
- Ellis, N. R., & Dulaney, C. L. (1991). Further evidence for cognitive inertia of persons with mental retardation. *American Journal on Mental Retardation*, *95*, 613–621.
- Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, *14*, 179–211.
- French, R. M., Addyman, C., & Mareschal, D. (2011). TRACX: A recognition-based connectionist framework for sequence segmentation and chunk extraction. *Psychological Review*, *118*, 614–636.
- French, R. M., Addyman, C., Mareschal, D., & Thomas, E. (2014). GAMIT – A Fading-Gaussian Activation Model of Interval-Timing: Unifying prospective and retrospective time estimation. *Timing & Time Perception Reviews*, *1* Article 2.
- Frings, C., Rothermund, K., & Wentura, D. (2007). Distractor repetitions retrieve previous responses to targets. *Quarterly Journal of Experimental Psychology*, *60*, 1367–1377.
- Gerstner, W., & Kistler, W. M. (2002). Mathematical formulations of Hebbian learning. *Biological Cybernetics*, *87*, 404–415.
- Gibbon, J., Church, R. M., & Meck, W. H. (1984). Scalar timing in memory. *Annals of the New York Academy of Sciences*, *423*, 52–77.
- Giesen, C., Frings, C., & Rothermund, K. (2012). Differences in the strength of distractor inhibition do not affect distractor-response bindings. *Memory & Cognition*, *40*, 373–387.
- Giesen, C., & Rothermund, K. (2014). Distractor repetitions retrieve previous responses and previous targets: Experimental dissociations of distractor-response and distractor-target bindings. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *40*, 645–659.
- Goschke, T. (2000). Intentional reconfiguration and involuntary persistence in task set switching. In S. Monsell & J. Driver (Eds.), *Control of cognitive processes: Attention and performance XVIII* (pp. 331–355). Cambridge, MA: MIT Press.
- Gratton, G., Coles, M. G. H., & Donchin, E. (1992). Optimizing the use of information: Strategic control of activation of responses. *Journal of Experimental Psychology: General*, *121*, 480–506.
- Grice, G. R. (1968). Stimulus intensity and response evocation. *Psychological Review*, *75*, 359–373.
- Grosjean, M., Rosenbaum, D. A., & Elsinger, C. (2001). Timing and reaction time. *Journal of Experimental Psychology: General*, *130*, 256–272.
- Hazeltine, E., & Mordkoff, J. T. (2014). Resolved but not forgotten: Stroop conflict dredges up the past. *Frontiers in Psychology*, *5* Article 1327.
- Heathcote, A., Brown, S., & Mewhort, D. J. K. (2000). The power law repealed: The case for an exponential law of practice. *Psychonomic Bulletin & Review*, *7*, 185–207.
- Hintzman, D. L. (1984). Minerva 2: A simulation model of human memory. *Behavior Research Methods Instruments & Computers*, *16*, 96–101.
- Hintzman, D. L. (1986). “Schema abstraction” in a multiple-trace memory model. *Psychological Review*, *93*, 411–428.
- Hintzman, D. L. (1988). Judgments of frequency and recognition memory in a multiple-trace memory model. *Psychological Review*, *95*, 528–551.
- Hommel, B. (1998). Event files: Evidence for automatic integration of stimulus-response episodes. *Visual Cognition*, *5*, 183–216.
- Hommel, B., & Colzato, L. S. (2009). When an object is more than a binding of its features: Evidence for two mechanisms of visual feature integration. *Visual Cognition*, *17*, 120–140.
- Hommel, B., & Colzato, L. S. (2015). Learning from history: The need for a synthetic approach to human cognition. *Frontiers in Psychology*, *6* Article 1435.
- Hommel, B., Musseler, J., Aschersleben, G., & Prinz, W. (2001). The Theory of Event Coding (TEC): A framework for perception and action planning. *Behavioral and Brain Sciences*, *24*, 849–878.
- Hommel, B., Proctor, R. W., & Vu, K. P. L. (2004). A feature-integration account of sequential effects in the Simon task. *Psychological Research Psychologische Forschung*, *68*, 1–17.
- Howard, M. W., & Kahana, M. J. (2002). A distributed representation of temporal context. *Journal of Mathematical Psychology*, *46*, 269–299.
- Hutchison, K. A. (2011). The interactive effects of listwide control, item-based control, and working memory capacity on Stroop performance. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *37*, 851–860.
- Jacoby, L. L., Lindsay, D. S., & Hessels, S. (2003). Item-specific control of automatic processes: Stroop process dissociations. *Psychonomic Bulletin & Review*, *10*, 638–644.
- Jacoby, L. L., Lindsay, D. S., & Toth, J. P. (1992). Unconscious influences revealed: Attention, awareness, and control. *American Psychologist*, *47*, 802–809.
- Jersild, A. T. (1927). Mental set and shift. *Archives of Psychology*, *14*. Whole No. 89.
- Jones, M., Mozer, M. C., & Kinoshita, S. (2009). Optimal response initiation: Why recent experience matters. In D. Koller (Ed.), *Advances in neural information processing systems 21* (pp. 785–792). Cambridge, MA: MIT Press.
- Kiesel, A., Steinhauser, M., Wendt, M., Falkenstein, M., Jost, K., Philipp, A. M., & Koch, I. (2010). Control and interference in task switching: A review. *Psychological Bulletin*, *136*, 849–874.
- Kinoshita, S., Forster, K. I., & Mozer, M. C. (2008). Unconscious cognition isn't that smart: Modulation of masked repetition priming effect in the word naming task. *Cognition*, *107*, 623–649.
- Kinoshita, S., Mozer, M. C., & Forster, K. I. (2011). Dynamic adaptation to history of trial difficulty explains the effect of congruency proportion on masked priming. *Journal of Experimental Psychology: General*, *140*, 622–636.
- Kohfeld, D. L. (1968). Stimulus intensity and adaptation level as determinants of simple reaction time. *Journal of Experimental Psychology*, *76*, 468–473.
- Kornblum, S., Stevens, G. T., Whipple, A., & Requin, J. (1999). The effects of irrelevant stimuli: 1. The time course of stimulus-stimulus and stimulus-response consistency effects with Stroop-like stimuli, Simon-like tasks, and their factorial combinations. *Journal of Experimental Psychology: Human Perception and Performance*, *25*, 688–714.
- Laird, J. E. (2012). *The soar cognitive architecture*. Cambridge, MA: MIT Press.
- Laird, J. E., Newell, A., & Rosenbloom, P. S. (1987). Soar: An architecture for general intelligence. *Artificial Intelligence*, *33*, 1–64.
- Levin, Y., & Tzelgov, J. (2016). Contingency learning is not affected by conflict experience: Evidence from a task conflict-free, item-specific Stroop paradigm. *Acta Psychologica*, *164*, 39–45.
- Lewicki, P. (1985). Nonconscious biasing effects of single instances on subsequent judgments. *Journal of Personality and Social Psychology*, *48*, 563–574.
- Lewicki, P. (1986). Processing information about covariations that cannot be articulated. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *12*, 135–146.
- Lewicki, P., Hill, T., & Czyzewska, M. (1992). Nonconscious acquisition of information. *American Psychologist*, *47*, 796–801.
- Lindsay, D. S., & Jacoby, L. L. (1994). Stroop process dissociations: The relationship between facilitation and interference. *Journal of Experimental Psychology: Human Perception and Performance*, *20*, 219–234.

- Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, 95, 492–527.
- Logan, G. D., & Bundesen, C. (2003). Clever homunculus: Is there an endogenous act of control in the explicit task-cuing procedure? *Journal of Experimental Psychology: Human Perception and Performance*, 29, 575–599.
- Logan, G. D., & Bundesen, C. (2004). Very clever homunculus: Compound stimulus strategies for the explicit task-cuing procedure. *Psychonomic Bulletin & Review*, 11, 832–840.
- Logan, G. D., & Schneider, D. W. (2006a). Interpreting instructional cues in task switching procedures: The role of mediator retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32, 347–363.
- Logan, G. D., & Schneider, D. W. (2006b). Priming or executive control? Associative priming of cue encoding increases “switch costs” in the explicit task-cuing procedure. *Memory & Cognition*, 34, 1250–1259.
- Logan, G. D., Schneider, D. W., & Bundesen, C. (2007). Still clever after all these years: Searching for the homunculus in explicitly cued task switching. *Journal of Experimental Psychology: Human Perception and Performance*, 33, 978–994.
- Los, S. A. (1994). Procedural differences in processing intact and degraded stimuli. *Memory & Cognition*, 22, 145–156.
- Los, S. A. (1996). On the origin of mixing costs: Exploring information processing in pure and mixed blocks of trials. *Acta Psychologica*, 94, 145–188.
- Los, S. A. (1999a). Identifying stimuli of different perceptual categories in mixed blocks of trials: Evidence for cost in switching between computational processes. *Journal of Experimental Psychology: Human Perception and Performance*, 25, 3–23.
- Los, S. A. (1999b). Identifying stimuli of different perceptual categories in pure and mixed blocks of trials: Evidence for stimulus-driven switch costs. *Acta Psychologica*, 103, 173–205.
- Lowe, D. G., & Mitterer, J. O. (1982). Selective and divided attention in a Stroop task. *Canadian Journal of Psychology*, 36, 684–700.
- Lupker, S. J., Brown, P., & Colombo, L. (1997). Strategic control in a naming task: Changing routes or changing deadlines? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23, 570–590.
- MacLeod, C. M. (1991). Half a century of research on the Stroop effect: An integrative review. *Psychological Bulletin*, 109, 163–203.
- MacLeod, C. M. (1998). Training on integrated versus separated Stroop tasks: The progression of interference and facilitation. *Memory & Cognition*, 26, 201–211.
- Matzel, L. D., Held, F. P., & Miller, R. R. (1988). Information and expression of simultaneous and backward associations: Implications for contiguity theory. *Learning and Motivation*, 19, 317–344.
- Mayr, U., Awh, E., & Laurey, P. (2003). Conflict adaptation effects in the absence of executive control. *Nature Neuroscience*, 6, 450–452.
- Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, 85, 207–238.
- Miall, C. (1989). The storage of time intervals using oscillating neurons. *Neural Computation*, 1, 359–371.
- Miller, J. (1987). Priming is not necessary for selective-attention failures: Semantic effects of unattended, unprimed letters. *Perception & Psychophysics*, 41, 419–434.
- Moeller, B., & Frings, C. (2014). Designers beware: Response retrieval effects influence drivers' response times to local danger warnings. *Transportation Research Part F-Traffic Psychology and Behaviour*, 24, 117–132.
- Monsell, S. (2003). Task switching. *Trends in Cognitive Sciences*, 7, 134–140.
- Mordkoff, J. T. (2012). Observation: Three reasons to avoid having half of the trials be congruent in a four-alternative forced-choice experiment on sequential modulation. *Psychonomic Bulletin & Review*, 19, 750–757.
- Mozer, M. C., Kinoshita, S., & Davis, C. (2004). Control of response initiation: Mechanisms of adaptation to recent experience. *Proceedings of the twenty sixth annual conference of the Cognitive Science Society*, 981–986.
- Myung, I. J., Kim, C., & Pitt, M. A. (2000). Toward an explanation of the power law artifact: Insights from response surface analysis. *Memory & Cognition*, 28, 832–840.
- Newell, A., & Rosenbloom, P. S. (1981). Mechanisms of skill acquisition and the law of practice. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition* (pp. 1–55). Hillsdale, NJ: Erlbaum.
- Nissen, M. J., & Bullemer, P. (1987). Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology*, 19, 1–32.
- Nosofsky, R. M. (1988a). Exemplar-based accounts of relations between classification, recognition, and typicality. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14, 700–708.
- Nosofsky, R. M. (1988b). Similarity, frequency, and category representations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14, 54–65.
- Nosofsky, R. M., Little, D. R., Donkin, C., & Fific, M. (2011). Short-term memory scanning viewed as exemplar-based categorization. *Psychological Review*, 118, 280–315.
- Nosofsky, R. M., & Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review*, 104, 266–300.
- Notebaert, W., & Verguts, T. (2007). Dissociating conflict adaptation from feature integration: A multiple regression approach. *Journal of Experimental Psychology: Human Perception and Performance*, 33, 1256–1260.
- Ollman, R. T., & Billington, M. J. (1972). The deadline model for simple reaction times. *Cognitive Psychology*, 3, 311–336.
- O'Reilly, R. C., & Rudy, J. W. (2001). Conjunctive representations in learning and memory: Principles of cortical and hippocampal function. *Psychological Review*, 108, 311–345.
- Perry, C., Ziegler, J. C., & Zorzi, M. (2010). Beyond single syllables: Large-scale modeling of reading aloud with the Connectionist Dual Process (CDP++) model. *Cognitive Psychology*, 61, 106–151.
- Pösse, B., Waszak, F., & Hommel, B. (2006). Do stimulus-response bindings survive a task switch? *European Journal of Cognitive Psychology*, 18, 640–651.
- Rabbitt, P. (1997). Introduction: Methodologies and models in the study of executive function. In P. Rabbitt (Ed.), *Methodology of frontal and executive function* (pp. 1–38). East Sussex, UK: Psychology Press Publishers.
- Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion decision model: Current issues and history. *Trends in Cognitive Sciences*, 20, 260–281.
- Ratcliff, R., Thapar, A., & McKoon, G. (2006). Aging, practice, and perceptual tasks: A diffusion model analysis. *Psychology and Aging*, 21, 353–371.
- Reber, A. S. (1967). Implicit learning of artificial grammars. *Journal of Verbal Learning and Verbal Behavior*, 6, 855–863.
- Rothermund, K., Wentura, D., & De Houwer, J. (2005). Retrieval of incidental stimulus-response associations as a source of negative priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31, 482–495.
- Schmidt, J. R. (2013a). The Parallel Episodic Processing (PEP) model: Dissociating contingency and conflict adaptation in the item-specific proportion congruent paradigm. *Acta Psychologica*, 142, 119–126.
- Schmidt, J. R. (2013b). Questioning conflict adaptation: Proportion congruent and Gratton effects reconsidered. *Psychonomic Bulletin & Review*, 20, 615–630.
- Schmidt, J. R. (2013c). Temporal learning and list-level proportion congruency: Conflict adaptation or learning when to respond? *PLoS One*, 8, e0082320.
- Schmidt, J. R. (2014). List-level transfer effects in temporal learning: Further complications for the list-level proportion congruent effect. *Journal of Cognitive Psychology*, 26, 373–385.
- Schmidt, J. R. (2016a). Proportion congruency and practice: A contingency learning account of asymmetric list shifting effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 42, 1496–1505.
- Schmidt, J. R. (2016b). Temporal learning and rhythmic responding: No reduction in the proportion easy effect with variable response-stimulus intervals. *Frontiers in Psychology*, 7 Article 634.
- Schmidt, J. R., & Besner, D. (2008). The Stroop effect: Why proportion congruent has nothing to do with congruency and everything to do with contingency. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34, 514–523.
- Schmidt, J. R., Cheesman, J., & Besner, D. (2013). You can't Stroop a lexical decision: Is semantic processing fundamentally facilitative? *Canadian Journal of Experimental Psychology*, 67, 130–139.
- Schmidt, J. R., Crump, M. J. C., Cheesman, J., & Besner, D. (2007). Contingency learning without awareness: Evidence for implicit control. *Consciousness and Cognition*, 16, 421–435.

- Schmidt, J. R., & De Houwer, J. (2011). Now you see it, now you don't: Controlling for contingencies and stimulus repetitions eliminates the Gratton effect. *Acta Psychologica*, 138, 176–186.
- Schmidt, J. R., & De Houwer, J. (2012a). Adding the goal to learn strengthens learning in an unintentional learning task. *Psychonomic Bulletin & Review*, 19, 723–728.
- Schmidt, J. R., & De Houwer, J. (2012b). Contingency learning with evaluative stimuli: Testing the generality of contingency learning in a performance paradigm. *Experimental Psychology*, 59, 175–182.
- Schmidt, J. R., & De Houwer, J. (2012c). Does temporal contiguity moderate contingency learning in a speeded performance task? *Quarterly Journal of Experimental Psychology*, 65, 408–425.
- Schmidt, J. R., & De Houwer, J. (2012d). Learning, awareness, and instruction: Subjective contingency awareness does matter in the colour-word contingency learning paradigm. *Consciousness and Cognition*, 21, 1754–1768.
- Schmidt, J. R., & De Houwer, J. (2016a). Contingency learning tracks with stimulus-response proportion: No evidence of misprediction costs. *Experimental Psychology*, 63, 79–88.
- Schmidt, J. R., & De Houwer, J. (2016b). Time course of colour-word contingency learning: Practice curves, pre-exposure benefits, unlearning, and relearning. *Learning and Motivation*, 56, 15–30.
- Schmidt, J. R., De Houwer, J., & Besner, D. (2010). Contingency learning and unlearning in the blink of an eye: A resource dependent process. *Consciousness and Cognition*, 19, 235–250.
- Schmidt, J. R., De Schryver, M., & Weissman, D. H. (2014). Removing the influence of feature repetitions on the congruency sequence effect: Why regressing out confounds from a nested design will often fall short. *Journal of Experimental Psychology: Human Perception and Performance*, 40, 2392–2402.
- Schmidt, J. R., Lemerrier, C., & De Houwer, J. (2014). Context-specific temporal learning with non-conflict stimuli: Proof-of-principle for a learning account of context-specific proportion congruent effects. *Frontiers in Psychology*, 5 Article 1241.
- Schmidt, J. R., & Liefoghe, B. (2016). Feature integration and task switching: Diminished switch costs after controlling for stimulus, response, and cue repetitions. *PLoS One*, 11, e0151188.
- Schmidt, J. R., Notebaert, W., & Van den Bussche, E. (2015). Is conflict adaptation an illusion? *Frontiers in Psychology*, 6 Article 172.
- Schmidt, J. R., & Weissman, D. H. (2014). Congruency sequence effects without feature integration or contingency learning confounds. *PLoS One*, 9, e0102337.
- Schmidt, J. R., & Weissman, D. H. (2016). Congruency sequence effects and previous response times: Conflict adaptation or temporal learning? *Psychological Research Psychologische Forschung*, 80, 590–607.
- Schneider, D. W., & Logan, G. D. (2007). Task switching versus cue switching: Using transition cuing to disentangle sequential effects in task-switching performance. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33, 370–378.
- Schneider, D. W., & Logan, G. D. (2009). Selecting a response in task switching: Testing a model of compound cue retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35, 122–136.
- Shiffrin, R. M., & Steyvers, M. (1997). Model for recognition memory: REM – Retrieving effectively from memory. *Psychonomic Bulletin & Review*, 4, 145–166.
- Simon, J. R., Craft, J. L., & Webster, J. B. (1973). Reactions toward stimulus source: Analysis of correct responses and errors over a five-day period. *Journal of Experimental Psychology*, 101, 175–178.
- Strayer, D. L., & Kramer, A. F. (1994a). Strategies and automaticity: 1. Basic findings and conceptual-framework. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20, 318–341.
- Strayer, D. L., & Kramer, A. F. (1994b). Strategies and automaticity: 2. Dynamic aspects of strategy adjustment. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20, 342–365.
- Stroop, J. R. (1935). Studies on interference in serial verbal reactions. *Journal of Experimental Psychology*, 18, 643–661.
- Treisman, A. M., & Gelade, G. (1980). Feature-integration theory of attention. *Cognitive Psychology*, 12, 97–136.
- Van Duren, L. L., & Sanders, A. F. (1988). On the robustness of the additive factors stage structure in blocked and mixed choice reaction designs. *Acta Psychologica*, 69, 83–94.
- van Maanen, L., Brown, S. D., Eichele, T., Wagenmakers, E. J., Ho, T., Serences, J., & Forstmann, B. U. (2011). Neural correlates of trial-to-trial fluctuations in response caution. *Journal of Neuroscience*, 31, 17488–17495.
- Vandierendonck, A., Liefoghe, B., & Verbruggen, F. (2010). Task switching: Interplay of reconfiguration and interference control. *Psychological Bulletin*, 136, 601–626.
- Verguts, T., & Notebaert, W. (2008). Hebbian learning of cognitive control: Dealing with specific and nonspecific adaptation. *Psychological Review*, 115, 518–525.
- Waszak, F., Hommel, B., & Allport, A. (2003). Task-switching and long-term priming: Role of episodic stimulus-task bindings in task-shift costs. *Cognitive Psychology*, 46, 361–413.
- Wickelgren, W. A. (1979). Chunking and consolidation: Theoretical synthesis of semantic networks, configuring in conditioning, S-R versus cognitive learning, normal forgetting, the amnesic syndrome, and the hippocampal arousal system. *Psychological Review*, 86, 44–60.