INCREMENTAL LEXICAL LEARNING IN SPEECH PRODUCTION: A COMPUTATIONAL MODEL AND EMPIRICAL EVALUATION

BY

GARY MICHAEL OPPENHEIM

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Psychology
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2011

Urbana, Illinois

Doctoral Committee:

Professor Gary S. Dell, Chair
Professor J. Kathryn Bock
Professor Neal J. Cohen
Professor John E. Hummel
Professor Brian H. Ross

ABSTRACT

Naming a picture of a dog primes the subsequent naming of a picture of a dog (repetition priming) and interferes with the subsequent naming of a picture of a cat (semantic interference). Behavioral studies suggest that these effects derive from persistent changes in the way that words are activated and selected for production, and some have claimed that the findings are only understandable by positing a competitive mechanism for lexical selection. This dissertation presents and evaluates a simple model of lexical retrieval in speech production that applies error-driven learning to its lexical activation network. This model naturally produces repetition priming and semantic interference effects. It predicts the major findings from several published experiments, demonstrating that these effects may arise from incremental learning. Furthermore, analysis of the model suggests that competition during lexical selection is not necessary for semantic interference if the learning process is itself competitive. Three additional experiments seek to evaluate the temporal persistence of semantic interference effects, as predicted by an incremental learning account.

ACKNOWLEDGMENTS

Many thanks to Gary Dell and Myrna Schwartz, my coauthors in the article that forms the initial portion of this dissertation. Thanks also to Xavier Alario, Aaron Benjamin, Kay Bock, Neal Cohen, Jason Crowther, Kara Federmeier, Vic Ferreira, Simon Fisher-Baum, Matt Goldrick, Catherine Hodgson, David Howard, John Hummel, Niels Janssen, Audrey Kittredge, Brad Mahon, Chad Marsolek, Randi Martin, Edu Navarrete, Lyndsey Nickels, Brian Ross, Tatiana Schnur, Eric Taylor, Sharon Thompson-Schill, three anonymous reviewers, and the Bock and Dell joint lab meeting for comments on this work and other contributions. Finally, thank you to my parents, Cal and Patti Oppenheim, and my girlfriend, Ying Wang, for their support and feedback throughout this process. This research was supported by National Institutes of Health Grants DC000191, HD44458, and MH1819990.

TABLE OF CONTENTS

CHAPTER 1: INTRODUCTION[1]

Retrieving a word from memory has consequences for later retrieval. This is particularly true when retrieval occurs in a semantic memory task such as picture naming. It is well known that the second presentation of a picture to be named speeds the naming response and diminishes the chance of error. This phenomenon, known as *repetition priming*, can be explained by the fact that each retrieval event is also a learning event, and so the second retrieval benefits from the learning that occurred the first time (*e.g.* Mitchell & Brown, 1988). Somewhat less well known is the fact that repetition priming has a "dark side". Retrieving a word has *negative* consequences for the subsequent retrieval of other words from the same semantic category (*e.g.* Belke, 2008; Belke, Meyer, & Damian, 2005; Blaxton & Neely, 1983; Brown, 1981; Damian & Als, 2005; Damian, Vigliocco, & Levelt, 2001; Howard, Nickels, Coltheart, & Cole-Virtue, 2006; Hsiao, Schwartz, Schnur, & Dell, 2009; Kroll & Stewart, 1994; Schnur, Schwartz, Brecher, & Hodgson, 2006; Vigliocco, Vinson, Damian, & Levelt, 2002; Wheeldon & Monsell, 1994). Following Oppenheim, Dell, and Schwartz

---

[1] Chapters 1-5 were produced in collaboration with Gary S. Dell and Myrna F. Schwartz. They have been published as follows:

Oppenheim, GM, Dell, GS, & Schwartz, MF. (2010). The dark side of incremental learning: A model of cumulative semantic interference during lexical access in speech production. *Cognition, 114,* 227–252.

(2007), we refer to these negative consequences as *cumulative semantic interference*. In this paper, we explain the mechanisms behind cumulative semantic interference in the domain of picture naming. This explanation takes the form of a computational model of lexical access in speech production that simulates the major phenomena in this domain. The model addresses meaning-based lexical retrieval in general, whether this is elicited by picture-naming, naming-to-definition, or spontaneous production. Our focus, however, is on persistent changes to lexical processing that result from the natural retrieval of a single word. The central theoretical point that the model implements is that repetition priming and cumulative semantic interference are two sides of the same coin. They both result from an error-based implicit learning process that tunes the language production system to recent experience.

Although our model is formally developed only for lexical access in speech production, our theoretical goals are more general. Cumulative semantic interference is a manifestation in speech production of a set of phenomena known in the memory literature as *retrieval-induced forgetting* or RIF. Retrieval-induced forgetting studies demonstrate that the episodic memory for a word or association can be impaired by the previous retrieval of a related memory (*e.g.* Anderson, Bjork, & Bjork, 1994; but see also Anderson and Neely, 1996, for a discussion of retrieval-induced forgetting in semantic memory). Currently, the explanation for such impairment is debated, with some claiming it results from suppressing previous competitors (often termed *inhibition* or *unlearning*; *e.g.* Anderson *et al.*, 1994; Melton & Irwin, 1940; Norman, Newman, & Detre, 2007; Postman, Stark, & Fraser, 1968) while others claim it stems from strengthening previous targets (*occlusion* or ‘*blocking*’[2]; *e.g.* MacLeod, Dodd, Sheard, Wilson, & Bibi, 2003; McGeoch 1932; Mensink & Raaijmakers, 1988). Our analysis of cumulative semantic interference in speech production will, we claim, speak to this debate. More generally, our model reflects a recent trend in cognition to link psycholinguistics with theories of learning and memory by developing accounts of how experience changes language processing (*e.g.* Chang, Dell, & Bock, 2006; Goldinger, 1998; Kraljic & Samuel, 2005).

Much of the theoretical importance of cumulative semantic interference hinges on an alleged property of requiring a competitive mechanism for lexical selection (*e.g.* Howard *et al.*, 2006). The most prominent theories of lexical access (*e.g.* Levelt, Roelofs, & Meyer, 1999) assume competitive lexical selection. Empirical support for this assumption has often come from picture-word interference studies (*e.g.* Schriefers, Meyer, & Levelt, 1990), in which speakers name pictures as they are presented at short offsets from distractor words. However, since Mahon, Costa, Peterson, Vargas, and Caramazza (2007) presented an analysis demonstrating that picture-word interference studies have not reliably supported the claims of competitive lexical selection, the search for empirical

[2] The term ‘blocking’ carries a quite different meaning in the retrieval-induced forgetting literature, where it refers to a hypothesis of competitor-based interference, than in the cumulative semantic interference literature, where it tends to refer to the structure of an experimental design (i.e. pictures may be presented in semantically homogeneous blocks).

support has turned to a simpler task: picture naming, specifically with regards to cumulative semantic interference.

Two serial picture-naming paradigms have been particularly common in studies of cumulative semantic interference. First is the *blocked-cyclic naming paradigm* (*e.g.* Damian *et al.*, 2001). In each block, subjects repeatedly cycle through naming a small set of pictures (*e.g.* one block might consist of four cycles through a set of six pictures). In the homogeneous condition, all the pictures in the block represent the same semantic category (*e.g.* farm animals), and in the mixed condition each picture represents a different semantic category. Cumulative semantic interference is indexed by greater difficulty naming pictures in the homogeneous condition relative to the mixed condition (the *semantic blocking effect*). Typically, the semantic blocking effect is not present in the first cycle and grows over subsequent cycles (*e.g.* Belke *et al.*, 2005). The second important serial picture-naming paradigm, used by Brown (1981, Experiment 4) and Howard *et al.* (2006), can be called the *continuous paradigm*. In this method, pictures drawn from several categories (*e.g.* animals, vehicles) are named without repeating any item, but with multiple exemplars from each category. Here, cumulative semantic interference is demonstrated by naming times that increase linearly as a function of the number of previously named pictures in that category. Importantly, the number of interspersed pictures between each category exemplar is irrelevant to the effect (Howard *et al.*, 2006). For example, in the sequence *goat, car, tomato, truck, horse*, naming time for *horse* would be slower than that for *goat*, and would be unaffected by the number of unrelated intervening items.

The nature of cumulative semantic interference: Howard et al.'s principles

Howard *et al.* (2006) argued that three specific properties of the lexical retrieval process must interact to produce cumulative semantic interference in naming latencies: shared activation, competitive selection, and priming. The idea is that each time a target word is activated, semantically related competitors are also activated (*shared activation*), and strongly activated competitors slow down the selection of target words (*competitive selection*). Retrieving a word once primes its future retrieval (*priming*), making it a stronger competitor when related words are retrieved in subsequent trials, thereby causing those subsequent target words to be retrieved more slowly. We will use these three properties to structure our review of the phenomenon and its implications for lexical retrieval.

Shared activation

When a target word such as DOG is activated during its attempted retrieval, its semantic relatives such as GOAT are also activated, thereby setting the scene for lexical competition. This principle of *shared activation* for semantically related words is what makes cumulative semantic interference specifically *semantic* in nature.

While the idea of shared activation is compatible with most current theories of semantic representation, it arises naturally from the use of distributed (or feature-based) semantic representations such as those commonly employed in connectionist models (see McClelland & Rogers, 2003, for a review). Distributed mechanisms would predict graded effects of semantic similarity, and indeed blocked-cyclic picture naming studies have demonstrated that more closely related items generate stronger interference effects than those more distant (Vigliocco *et al.*, 2002). So, for the purpose of understanding cumulative semantic interference, it may be useful to think of shared activation arising from shared semantic features rather than all-or-none category membership. That is how shared activation is implemented in our model.

As noted by Howard *et al.* (2006), however, shared semantic activation does not require distributed representations. It may occur with non-decomposed (localist) lexical concepts (*e.g.* Roelofs, 1992) provided that related concepts connect either directly (*e.g.* Collins & Loftus, 1975) or indirectly through shared category or property nodes (*e.g.* Collins & Quillian, 1969), and each activated concept sends activation to neighboring concepts. Moreover, any graded effects can be attributed to gradations in the number or strength of such connections. Thus, the finding of graded cumulative semantic interference does not allow us to distinguish between distributed and localist semantic representations.

Competitive selection

The second property of lexical retrieval that is required for cumulative semantic interference, according to Howard *et al.* (2006), is that lexical selection be *competitive*. That is, increasing the activation of non-target words should decrease the speed and accuracy with which a target word is selected. In a competitive selection process, words compete in the manner of two athletic teams during “sudden-death overtime”: the competition continues until a single winner emerges. This might be implemented via either a differential threshold (*e.g.* Levelt *et al.*, 1999) or lateral inhibition (*e.g.* Howard *et al.*, 2006), but the key is that having multiple strong competitors makes it harder to select a winner (Wheeldon & Monsell, 1994). A non-competitive selection process (*e.g.* Mahon *et al.*, 2007), in contrast, is more like a horse race that ends when the first contestant crosses a pre-determined absolute threshold. To illustrate this difference, let us imagine selecting a target word, DOG, when a competitor, GOAT, is also activated. According to a sudden-death competition method, the two compete until one clearly wins, so selecting DOG should be slower and less accurate when GOAT is more active. Thus, cumulative semantic interference in response time would occur if the semantic manipulations raise the activation of competitors. With a horse-race selection method, the speed of DOG's selection is entirely a function of DOG's own activation. The activation of GOAT does not enter into the equation, so this non-competitive selection offers no obvious way to account for cumulative semantic interference.

Locating the competition within the word production process is difficult, but several studies constrain it to a point after semantic access and before phonological access. Two findings argue for a post-semantic locus. First, performing non-verbal semantic judgments on pictures in the blocked-cyclic paradigm has proven insufficient to elicit a semantic blocking effect (Damian *et al.*, 2001). So any competition that occurs during stages before lexical access does not appear sufficient to drive the cumulative semantic interference effect. Second, bilingual continuous paradigm experiments indicate that cumulative semantic interference accumulates independently for each language, suggesting that the competitive selection process is language-specific and hence post-semantic (Castro, Strijkers, Costa, & Alario, 2008, Experiments 3 and 4).

Some evidence suggests that the competition may instead characterize the selection of abstract, pre-phonological word-forms, or lemmas. Blocked-cyclic *word* naming (reading words aloud) appears to produce semantic facilitation rather than interference, suggesting that the competition that results when naming pictures must arise before retrieving phonological word-forms (Damian *et al.*, 2001, Experiment 2a). Retrieving gender-marked determiners during word naming may bring back the semantic blocking effect, suggesting that the competition affects the post-semantic, pre-phonological retrieval of abstract lexical concepts (Damian *et al.*, 2001, Experiment 2b).

Together, the principles of shared lexical activation and competitive lexical selection are sufficient to produce the sort of semantic interference that might be seen within a single trial (*e.g.* as in the picture-word interference effect, *e.g.* Schriefers *et al.*, 1990). Shared activation causes semantically related competitors to become active, and a competitive lexical selection mechanism allows these competitors to hinder selection of a target word. Making semantic interference cumulative, however, requires some mechanism by which processes during one trial can affect subsequent trials. This is the function of priming.

Priming

Priming is Howard *et al.*'s final necessary property for cumulative semantic interference. Retrieving a word once should facilitate its future retrieval by either making the word itself more accessible or making its competitors less accessible.

While priming can be implemented in a number of ways, its effects can be characterized as either temporary or persistent. Temporary effects occur when priming is ascribed to changes in activation levels—either positive (*e.g.* Crowther, Martin, & Biegler, 2008; Howard *et al.*, 2006; Wheeldon & Monsell, 1994) or negative (inhibitory) changes (*e.g.* Brown, 1981; McCarthy & Kartsounis, 2000) that are carried over from previous trials. For example, selecting DOG may require the temporary suppression of the urge to say CAT, which might make it more difficult to access CAT for a short time. Persistent accounts (*e.g.* Damian & Als, 2005; Howard *et al.*, 2006; Schnur *et al.*, 2006) instead describe priming as a consequence of relatively permanent changes to the way words are accessed, such as incremental learning. Inspired in part by neural network models in which incremental learning is attributed to changes in connection weights rather than activation levels, persistent priming is an example of the learning that continually adjusts the cognitive system to suit its environment (*e.g.* Gupta & Cohen, 2002).

A critical property of the priming mechanism that underlies cumulative semantic interference is that the interference accumulates incrementally as a function of relevant experience (such as naming semantically related pictures), and is unaffected by irrelevant experience (such as naming unrelated pictures). In Howard *et al.*'s (2006) continuous paradigm study, naming pictures from a single category, such as DOG and then GOAT, produced the same linear accumulation of semantic interference whether the relevant pictures were separated by two, four, six, or eight unrelated items, suggesting that only the relevant experience matters. Moreover, in a variant of Howard *et al.*'s (2006) continuous paradigm, Navarrete, Mahon, and Caramazza (2008) showed that repeating an item, such as DOG, GOAT, DOG produced the same cumulative interference as accessing an additional novel exemplar from the category, thus demonstrating that each act of retrieval contributes to the effect.

Further evidence of robustness to irrelevant experience comes from the blocked-cyclic paradigm. Damian and Als (2005) showed that performing nonlinguistic tasks (Experiment 1) and naming unrelated items (Experiments 2 and 3) in between naming DOG and GOAT failed to disrupt the semantic blocking effect. Thus, the priming from each relevant experience contributes separately and robustly to the cumulative semantic interference effect, and irrelevant experience affects neither its accumulation nor diminution.

Filler trials, as used in the Howard *et al.* (2006) and Damian and Als (2005) studies required additional time to present and process, increasing the chronological time between the retrieval of related words. So these studies speak to residual activation (or inhibition) accounts of priming. Priming by residual activation should be strongly affected by the time between prime(s) and target. As argued by Bock and Griffin (2000), the activation levels that control language production must decay quickly in order for production — the rapid sequential activation of linguistic units — to succeed. For example, computer simulations of multi-word production have required that activation levels decay with time constants such that activated linguistic units lose nearly all of their activation within a second or two (*e.g.* Dell, 1986). Similarly, the effects of inhibitory processes in production are also time-bound. For example, many production theories assume that selection of a linguistic unit entails a reduction to a zero or negative activation value for the selected unit (*e.g.* Dell, Burger

& Svec, 1997; Houghton, 1990). However, the effects of this inhibition are quite temporary and are designed to prevent an immediate perseveratory error. The fact that the filler trials in Howard *et al.* (2006), and Damian and Als (2005), failed to affect cumulative semantic interference suggests that the priming that underlies this effect is reasonably persistent. Hence, cumulative semantic interference therefore is likely not largely based on the positive or negative changes in activation levels that arise solely through the spreading activation mechanisms of the production system.

A more direct demonstration of the temporal insensitivity of priming comes from Experiment 1 of Schnur *et al.* (2006), who compared naming latencies in a blocked-cyclic naming paradigm in which pictures were presented either 1-s or 5-s after the previous response (*i.e.* a 1-s or 5-s response-stimulus interval, or RSI). Any time-based decay of residual activation predicts a statistical interaction between presentation rate and semantic blocking condition, specifically less effect of the blocking with the long RSI (*e.g.*, Wilshire & McCarthy, 2002). Both presentation rates produced reliable cumulative semantic interference effects, with no interaction between RSI and semantic blocking condition, demonstrating that the priming is insensitive to the passage of time, at least at these intervals.

To summarize, the priming that causes cumulative semantic interference is temporally persistent, it accumulates with relevant experience, and it is insensitive to irrelevant experience. These properties offer an awkward fit for mechanisms based on residual activation or inhibition of linguistic units, both of whose effects should be expected to decay rather quickly. Instead, we follow Damian and Als (2005), Schnur *et al.* (2006), and Howard *et al*. (2006) by suggesting that the priming that underlies cumulative semantic interference emerges from small, persistent, experience-driven, post-selection adjustments to the mapping from semantics to words, *i.e.* incremental learning within the production system.

Speech errors and cumulative semantic interference

It appears that cumulative semantic interference does more than slow lexical retrieval; it also causes lexical selection errors. In a blocked-cyclic naming task with healthy older controls, Schnur *et al.* (2006, Experiment 1) found that naming latencies were higher and errors more frequent in the homogeneous condition, relative to the mixed condition. They also tested aphasic patients (Schnur *et al.*, 2006, Experiment 2), who made many more errors, and reported two important findings. First, patients made more semantic errors (*e.g.* naming DOG as CAT) and omissions in the homogeneous than mixed condition. Second, these semantic blocking effects increased across cycles, while other types of errors (*e.g.* phonological) showed the opposite pattern. The patients' increasing semantic blocking effects for semantic and omission errors thus resembled healthy adults' increasing blocking effects for naming latencies, suggesting that they might stem from the same underlying causes.

The link between the blocking effects on errors in patients and on latencies in unimpaired speakers also has some support from studies that attempt to associate these effects with brain regions. Neuroimaging of healthy subjects demonstrated that activation in the left inferior frontal gyrus (LIFG) correlates with increases in naming latencies due to semantic blocking and related manipulations (Moss *et al.*, 2005; Schnur *et al.*, 2009). The LIFG, for reference, corresponds to

Brodmann's areas (BA) 44, 45, and 47, the posterior part of which (BA 44/45) is Broca's area. And lesion analyses of patients from the Schnur *et al*., (2006) study revealed an association between LIFG damage and the increase in errors across blocking cycles (Schnur, Lee, Coslett, Schwartz, & Thompson-Schill, 2005; Schnur *et al*., 2009).

A second important finding from these patient studies is that patients' error effects are also robust to timing manipulations. Schnur *et al.* (2006) found that the blocking effect on errors – like the blocking effect on naming latency with unimpaired speakers – was not influenced by whether pictures were named with a 1-s or a 5-s RSI. Further examining these patients' naming errors, Hsiao *et al*. (2009) found that their within-set perseverations tended to match the words that they had used most recently. For instance, if a patient named pictures of a dog, a pig, and a goat correctly (*i.e.* saying PIG DOG GOAT) before incorrectly naming a picture of a horse, then she was more likely to name the horse as DOG than as PIG. Crucially, the key measure of recency was not time, but the number of intervening items (henceforth *item-lag*). Specifically, the chance-corrected perseveration lag functions were the same regardless of whether pictures were presented 1-s or 5-s after a patient's most recent response. Together with the temporal insensitivity of unimpaired speakers' response time effects, these results support our previous suggestion that relevant intervening experience, not timing, matters for the build-up or dissipation of cumulative semantic interference.

Modeling cumulative semantic interference

Howard *et al*. (2006) presented an elegant model of the effect of cumulative semantic interference on response time. In their model, shared activation is implemented by assuming that words receive continuous (integrated over computationally discrete timesteps) activation from semantic nodes, and each time one semantic node is activated, similar semantic nodes are also activated to a lesser degree. Lexical competition is implemented by inhibitory connections running from each word to every other word (*i.e.* lateral inhibition). A word is only selected upon reaching an absolute selection threshold, but since activated words inhibit each other, strong competitors can slow down the selection of a target word. Finally, each time a word is selected, its connection from its semantic node grows stronger, implementing a priming function.

The Howard *et al.* model is noteworthy because it instantiates the principles of shared activation, competitive selection, and priming and because it attributes the interference to processes that are insensitive to time and to unrelated interference. Our goal is to extend this approach. We do so in three respects. First, we identify the priming mechanism with error-driven connectionist learning. This learning mechanism has the natural property that each act of retrieval in a certain context strengthens the target of retrieval (repetition priming) while at the same time making it less likely that similar memories are retrieved instead in that context (similarity-sensitive interference). We will show how this mechanism is consistent with both the perseveratory gradient in the production of word errors and the insensitivity of cumulative semantic interference to unrelated items or the passage of time. By attributing these effects to error-driven learning, we link up with the many cognitive models that are based on such learning (e. g. Chang *et al.*, 2006; Gupta & Cohen, 2002; Plaut, McClelland, Seidenberg, & Patterson, 1996) and, as we later demonstrate, this attribution addresses the question of whether retrieval-induced forgetting is caused by “inhibition”.

Second, we develop the model in conjunction with theories of word production so that it can account for errors as well as response times. This requires a decision process that allows for lexical selection to play out in time, and for errors of commission and errors of omission. We propose that lexical selection, even in single-word production, is driven by the intersection of semantic activation with other (e.g. syntagmatic) constraints, and that this selection process may underlie the correlation between cumulative semantic interference and activation of Broca's area. Finally and most importantly for current debates, we offer the hypothesis that cumulative semantic interference does not, in fact, require a competitive mechanism for lexical selection. Specifically, we demonstrate that competition in the lexical selection process is unnecessary when combined with error-driven learning. The resulting non-competitive model, we claim, can explain the major findings concerned with cumulative semantic interference.

CHAPTER 2: MODEL DESCRIPTION

Overview

The key components of our model concern *lexical activation*, *lexical selection*, and *learning*. As in many models of lexical retrieval in production, retrieving a word begins with activating a set of semantic features (*e.g.* Dell *et al.*, 1997; Gordon & Dell, 2003; Rapp & Goldrick, 2000). These semantic features each connect to a number of words, and thus activate those words in proportion to the strength and number of these connections (lexical activation, Figure 1). Thus, multiple words are activated, requiring some kind of decision. For the model, we assume that the most active word is chosen. However, when more than one word is activated, it is assumed to be difficult to identify the most active one (*i.e.* if the difference in activations is slight, the winner is hard to "see"), so a 'booster' process kicks in to tease the activations apart. This booster represents a specific manifestation of a larger executive control process (such as Dell, Oppenheim, and Kittredge's, 2008, syntactic 'traffic cop'), repeatedly amplifying each word's activation until a winner can be selected (lexical selection), or until this boosting process times out. Response time is assumed to be correlated with the number of boosts needed for the winner to emerge. Errors of commission, such as semantic errors, occur when the wrong word is chosen, and errors of omission occur if the booster times out. Finally, after lexical selection has concluded, an error-driven learning process adjusts the semantic-to-lexical connections so as to facilitate future retrieval of the target word (learning). In the following sections, we describe the details of the model's architecture, and its lexical activation, lexical selection, and learning mechanisms.

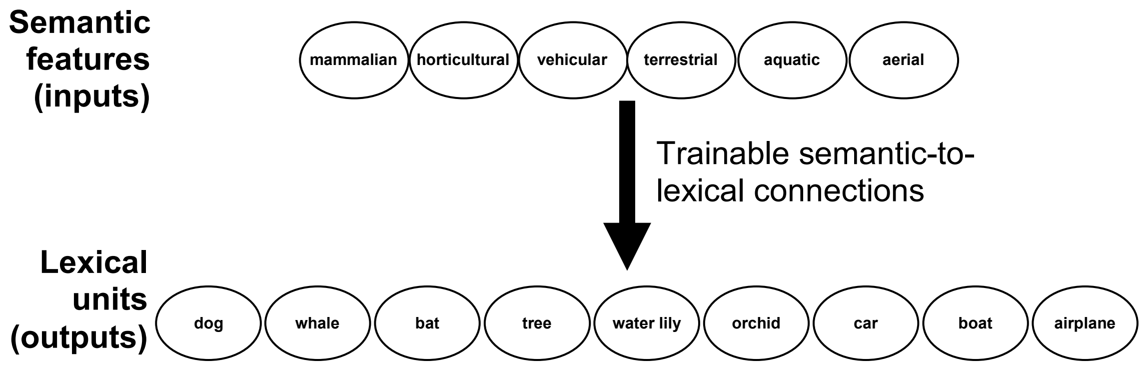

Figure 1. Words are activated by distributed semantic representations, or features. The strength of a connection from a semantic feature to a word is established by an initial period of error-based learning, and this learning continues throughout the simulations.

Model architecture

The model is a feedforward two-layer network. Semantic feature nodes (*e.g.* FURRY or AQUATIC) form the input layer of the network. Each feature node connects directly to each of the word nodes (such as DOG or BOAT) in the output layer. Connection weights are initialized at zero and are continually adjusted through an error-driven learning process, as detailed later in this description. There are no lateral connections between semantic feature nodes or between word nodes, or reverse connections from words to features.

Algorithms

Lexical activation. When semantic features are activated (as we assume happens when a picture is presented), these features in turn activate words. The net input, net_i, to any lexical node i, sums the activation, a_j, of each semantic feature, j, times the weight of its connection to the lexical node, w_{ij} (Equation 1).

Equation 1 $$net_i = \sum_j w_{ij} a_j$$

This net input, net_i, is then converted to an activation, a_i, via a logistic function (Equation 2).

Equation 2 $$a_i = \frac{1}{1 + e^{-net_i}}$$

Thus, the activations range from zero to one. We assume that lexical activation is imprecise and therefore add a small amount of normally-distributed noise, ν (with a mean of 0 and a standard deviation of θ), to the net input, net_i, yielding Equation 3.

Equation 3 $$a_i = \frac{1}{1 + e^{-(net_i + \nu)}}$$

Lexical selection. The next stage applies a competitive winner-take-all process to the lexical activations, linking increased lexical competition to increased naming latencies.

Although models of single-word production are most straightforwardly built around semantic activation, communicative speech production requires that lexical selection also take non-semantic constraints into account. A prime example of this principle is that, in sentence production, one has to not only select the right word, but also select it at the right time (e.g. Dell, Oppenheim, & Kittredge, 2008). For instance, the sentence, "The dog chased a squirrel," uses five words to communicate the idea of a dog chasing a squirrel, but one could not transpose any of the words without changing or disrupting the sentence's meaning. And empirically, the transpositions that do occur tend to respect syntactic category constraints (Ferreira & Humphreys, 2001; Garrett, 1975; Nooteboom, 1969) – for example, "The squirrel chased a dog," (a noun transposed with another noun) is a more likely error than, "The chased dog a squirrel" (a noun transposed with a verb) – suggesting that syntactic information somehow constrains lexical retrieval. Some models of sentence production have implemented this constraint by imagining lexical retrieval as involving an intersection of semantic and syntactic activations, where syntactic-sequential states drive the serial selection of words that have already been activated by a static semantic message (e.g. Chang, Dell, & Bock, 2006; Gordon & Dell, 2003). The core idea here is that lexical selection is ultimately guided by more than just semantic information: It incorporates an executive control mechanism that allows the *ad hoc* convergence of different kinds of information to influence lexical selection, for instance allowing a word to be selected according to both stable semantic features and the temporary demand to produce a noun.

Experiments demonstrating cumulative semantic interference tend to present relatively consistent demands (for example, requiring subjects to retrieve either nouns or verbs, but not to switch between nouns and verbs), so the flexibility that could be provided by a cognitive control process may not be immediately relevant to these simulations. However, we assume that speakers still use such processes to guide lexical selection in laboratory tasks like picture-naming, where words might be initially activated by semantic features, with the activation of relevant classes of words (e.g. nouns) then 'boosted' by input from a cognitive control process. Some evidence for the application of such a control mechanism to single word production comes from the maintenance of the syntactic category constraint in single-word picture naming (e.g. Dell et al., 2008), part-of-speech priming in homophone-pronunciation tasks (Melinger & Koenig, 2007), and more generally from demonstrations that lexical selection is hindered by dual task performance, indicating a dependence on central cognitive resources (Ferreira & Pashler, 2002).

Thus, we assume that lexical selection in picture naming experiments is aided by a cognitive control process which boosts the activation of task-relevant words (here, all nouns, which constitutes the entire simulated vocabulary), and for these simulations its contribution is integrated over time. After words are initially activated by semantic features, the booster function floods the network with additional activation that combines nonlinearly with the existing lexical activation until either one word grows discernibly more active than the rest or the boosting process times out. Notice that this booster process is "dumb" in the sense that it does not know which word is the target. It repeatedly boosts all of them. But because it boosts them in a multiplicative manner, the most active word gradually increases its lead on the others.

The booster is engaged only to the extent necessary to select a single word (that is, though we assume that it is always engaged during lexical retrieval, it operates more when selection is difficult), recalling Schnur *et al.*'s (2009) reports of greater LIFG activity as a function of increased lexical competition. Therefore we tentatively identify this booster with the competition-biasing mechanisms that are hypothesized to be a function of the LIFG (*e.g.* Kan & Thompson-Schill, 2004; Thompson-Schill, D'Esposito, Aguirre & Farah, 1997), but we acknowledge that our implemented booster has arbitrary properties that lack neural motivation. That is, we commit to the *functions* of the booster (aiding selection when competition is present) and its possible association to the LIFG (Broca's area), rather than to its implemented details.

The boosting process plays out over time. To determine whether a winner has emerged, at each timestep, t_n, we compare the difference between the activation of each word node, $a_{i\,t_n}$, and the mean activation of the other word nodes, $a_{others\,t_n}$, to a threshold value, τ (Equation 4).

Equation 4 $$\tau > \left(a_{i\,t_n} - a_{others\,t_n}\right)$$

If no word's activation difference exceeds the difference threshold (*i.e.* Equation 4 is false for all i), then the booster multiplies each word's current activation level, $a_{i\,t_n}$, by a constant boosting factor, $\beta > 1.0$. The result becomes its new activation level, $a_{i\,t_{n+1}}$ (Equation 5). Then this testing and boosting process repeats.

Equation 5 $$a_{i\,t_{n+1}} = a_{i\,t_n}\beta$$

A word is selected, that is, the boosting stops, if and when its activation advantage over other words (per Equation 4) is great enough. The timestep at which this selection occurs, $t_{selection}$, is treated as an index of the duration of the lexical selection process, which should correlate with naming latency. If, for the sake of simplicity, we assume no variation in the repeated boosting, this iterative process becomes computationally equivalent to Equation 6.

Equation 6 $$t_{selection} = \log_{\beta}\left(\frac{\tau}{a_{i\,t_1} - a_{others\,t_1}}\right)$$

However, if no node reaches the difference threshold within a certain number of boosts, Ω, then no word is selected and the trial is an omission. This corresponds to a simple "wait and give up" theory of omissions. So we do not consider an omission as a special state that may be achieved, but rather a lack of sufficient evidence for any particular word, making it difficult to select a word quickly enough.

Note that while the implemented boosting process may be deterministic, based on the initial activations, the target word will not necessarily be selected. The combination of a discernible-difference threshold and a selection deadline may preclude selecting any word if the difference in lexical activations is too small. Adding noise to the lexical activations, as we have done, increases this chance and further opens the possibility that a competitor will be selected instead. Furthermore, although we have not done so here, one could assume that the boosting process is subject to noise either in its normal operation, or in pathological cases (*e.g.* LIFG damage), by allowing for boosts to randomly fail for particular words at particular time steps. A noisy booster would then have properties in common with sequential stochastic decision mechanisms such as a random-walk process.

We do not implement any residual activation or inhibition in this model. When the trial ends, either by selecting a word or by failing to select a word before the deadline, all activations return to zero. Thus we assume that activation should largely or entirely dissipate within the inter-stimulus intervals that are typically used in experiments demonstrating cumulative semantic interference, and specifically claim that any activation that might persist across trials should not noticeably contribute to the observed interference.

Learning. At the end of each trial, semantic-to-lexical connection weights are adjusted according to Equation 7, which is the Widrow-Hoff or delta rule tailored for the logistic activation function (Widrow & Hoff, 1960; Rumelhart & McClelland, 1986): Δw_{ij} is the weight change for the connection to node *i* from node *j*, η is the learning rate, and d_i is the desired activation of node *i*.

Equation 7 $$\Delta w_{ij} = \eta\left(a_i\left(1-a_i\right)\left(d_i - a_i\right)\right)a_j$$

Since this equation will prove crucial for understanding the behavior of the model, we should unpack it a bit more. We have said that learning is error-driven. This means that connections are adjusted according to (d_i - a_i), the discrepancy between the desired activation of output node *i*, d_i, and its actual activation, a_i (that is, its activation before boosting). So the error in a receiving node's activation affects both the degree and the direction of the weight change. Hence, when the error for an output node (d_i - a_i) is strongly positive, connections feeding it will be greatly strengthened. When the error for an output node is strongly negative, the connections feeding it will be greatly weakened. Notice that because the logistic activation function precludes activations that are actually 0 or 1, every word unit will experience at least some error on all trials, either positive or negative, if the desired activations are 0 or 1. Next, including the $a_i(1-a_i)$ component scales weight adjustments to a_i, such that weight changes are greatest at a_i=0.5, and decrease as a_i approaches 0 or 1. Thus, weight changes are strongest for connections that contribute to moderate activations. Adding the a_j specifies that connections from that connections from input *j* should only be modified to the extent that *j* is activated. And finally, the learning rate, η, is simply an arbitrary global parameter, used to adjust how rapidly weight changes occur.

Thus the learning algorithm increases the connection weights from active semantic features to the target word, and decreases weights from those features to all other words, to the extent that those words were active before boosting. Since this learning is based on the deviation between d_i and a_i, it occurs regardless of whether the target was ultimately selected. So if the network encountered a dog (activating semantic features MAMMAL and TERRESTRIAL), then the connections from MAMMAL and TERRESTRIAL to DOG would strengthen, and the connections from MAMMAL and TERRESTRIAL to any other activated words (*e.g.* BAT) would weaken. The next time the network encounters a dog, those same semantic features will activate DOG more efficiently (*i.e.* activating DOG more and competitors less), increasing the speed and likelihood of its selection.

CHAPTER 3: AN ACCOUNT OF PUBLISHED FINDINGS

Our lexical learning model integrates several features common to theories of lexical access: lexical retrieval begins when distributed semantic features activate words (*e.g.* Dell *et al.*, 1997; Rapp & Goldrick, 2000); lexical selection uses a differential threshold (*e.g.* Levelt *et al.*, 1999); and semantic-to-lexical connections are adjusted through experience (*e.g.* Gordon & Dell, 2003; Howard *et al.*, 2006).

Can this model account for the major behavioral manifestations of cumulative semantic interference? To answer this question, we simulate several of the experiments described in the Introduction. First, Simulation 1 compares predicted selection latencies in a continuous paradigm simulation to Howard *et al.* (2006)'s behavioral data, and establishes one additional prediction from the model. In Simulation 2 we compare predicted latency patterns from blocked-cyclic presentation to data from Schnur *et al.* (2006, Experiment 1), Damian and Als (2005), and Belke (2008). Simulation 3 tests whether the simulated semantic blocking effect generalizes to new items, as reported by Belke *et al.* (2005). In Simulation 4, we turn to the aphasic patient data. Repeating the procedure from Simulation 2 with noisier lexical activations, we compare the resulting error patterns with those reported by Schnur *et al.* (2006, Experiment 2) and Hsiao *et al.* (2009). Simulation 5 explores the mechanisms behind the model's effects by distinguishing the influence of weight increases (facilitatory learning) and weight decreases (inhibitory learning) on response times, analogous to the 'blocking' versus 'inhibition' debate in the retrieval-induced forgetting literature. Finally, Simulation 6 examines the role of competitive lexical selection in creating cumulative semantic interference, demonstrating that competitive lexical selection is not, in fact, necessary for any of the effects seen in Simulations 1-5.

Our goal in these simulations is to explore how a simple, omnipresent process – incremental learning – can lead to a surprising range of behavioral effects depending on the nature and ordering of stimuli during an experiment. Toward that end, we hold all aspects of the simulations constant throughout this paper, except where we have motivated reasons to change them.

We implement a standardized vocabulary structure for these simulations (Figure 2). Equal numbers of words share each semantic feature, and each word is uniquely specified by the conjunction of two semantic features. For example, WHALE, BOAT, and WATERLILY share the feature AQUATIC, whereas BAT, PLANE, and ORCHID share AERIAL. And the features AQUATIC and VEHICULAR specify the word BOAT, while AQUATIC and MAMMALIAN specify WHALE.

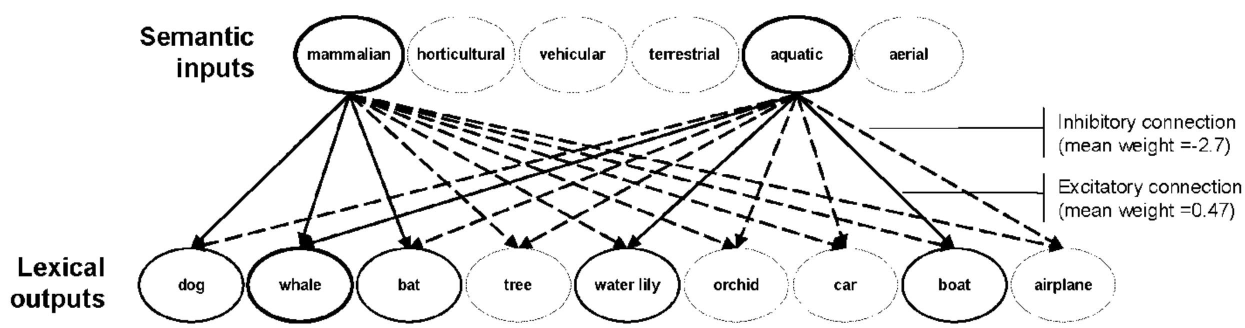

Figure 2. A scaled-down illustration of a trained network's vocabulary. Here, the feature MAMMALIAN excites three words (DOG, WHALE, and BAT) and inhibits the rest. The feature AQUATIC works the same way. When these features are activated together, they

activate all of the words that are either mammalian or aquatic. But they activate WHALE most strongly because it is both mammalian *and* aquatic.

When simulating picture naming in the model, we assume that the target picture is correctly recognized and that its semantic features are properly activated. Features that should be active get activations of 1, and those that should not be active get activations of 0. So we do not simulate errors in pre-lexical processes, though we concede that they can occur (*e.g.* Rogers et al, 2004).

Each simulation consists of two phases: training and testing. In training, we simulate the acquisition of subjects' pre-experimental lexical-semantic knowledge. Then, in testing, we simulate the learning that occurs during the experiment. We present trials that mimic the experimental conditions being simulated and continue to adjust the connections according to the same learning algorithm (including parameter values) that were used during training. So the only difference between the training and testing phases is that the testing phases focus on particular subsets of the vocabulary.

When simulating multiple testing conditions, such as homogeneous versus mixed blocks, we want to compare these directly. So we start the testing phase for each condition with the same trained weights. Thus the only difference between the conditions is the semantic relationship of the items cued in their testing phases.

To simulate testing multiple subjects with pictures from many categories, we repeat each simulation 10,000 times. Individual differences in experience are simulated by beginning each replication with a fresh network in *tabula rasa* state, and then training it with 100 randomly ordered sweeps through the vocabulary to represent the experience that this model subject brings to the experiment. Variation in the networks' performance thus comes from activation noise (present in both training and testing) and differences in the order in which words are cued.

Simulation 1 – continuous paradigm

Howard *et al.* (2006) reported that picture-naming latencies increased by a consistent amount for each same-category item that was named. The magnitude of the incremental increase was unaffected by the number of intervening items from different categories. They claimed that shared activation, competitive selection, and priming were necessary for any model to account for these findings. Because our model implements all three of these properties, it should exhibit a lag-insensitive incremental increase in selection times.

In Simulation 1a, we approximate Howard *et al.*'s protocol by cueing for production of five items from a single semantic category, like "farm animals", one at a time. A variable number of unrelated fillers (two, four, six, or eight) are cued between each critical item and the next, allowing examination of how the accumulation of semantic interference is affected by the number of intervening filler items.

Method

Parameters for this simulation are given in Table 1. Constraints on these parameters are discussed in the Results section.

Table 1. Model parameters for Simulations 1, 2, 3, 5, and 6.

Parameter	Value
Learning rate (η)	0.75
Activation noise (θ)	0.5
Boosting rate (β)	1.01
Threshold (τ)	1
Deadline (Ω)	100

Training. Networks were trained, as described above, on a vocabulary of 20 semantic features and 50 words.

Testing. A list of 25 pictures comprised the test phase. Five of the items on this list were critical items and came from the same category. Each shared one of its two semantic features with every other critical item. Following Howard *et al*. (2006), a lag of 2, 4, 6, or 8 fillers separated each critical item and the next, making a total of 20 fillers. These fillers were randomly selected with the constraint that they shared no semantic features with the critical items; they were, however, free to share features with other fillers. Each lag occurred once in each list, yielding 4!=24 possible lag sequences. Each of 10,000 networks was tested with each of these 24 sequences.

Analyses. Following Howard *et al.* (2006), mean lexical selection times were calculated for each lag and ordinal position. In our simulations, these selection times are the mean numbers of 'boosts' needed before one output node could be selected, as described in Equation 6. Following the standard practice in picture naming studies, only the selection times for correct responses were included in these calculations. Errors of omission and commission occurred in less than 1% of the trials and were excluded from the analyses.

Results and discussion

As expected, each critical item took longer to select than the previous one, with each item contributing an equal increment to the selection time. There was no systematic variation in this effect over the different lags (Figure 3a).   These results are entirely consistent with Howard *et al.*'s human data (Figure 3b).

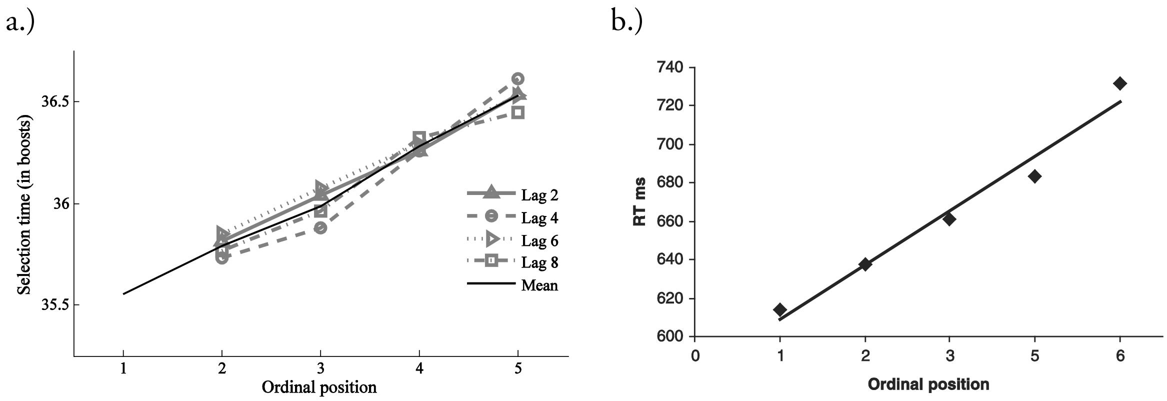

Figure 3. In Simulation 1a, the continuous paradigm, each critical item takes a bit longer to select than the previous one, regardless of the number of intervening unrelated items. a.) Model-predicted mean correct selection times in each of the four lag conditions show identical increases with ordinal position. We present only an overall mean selection time for ordinal position 1 because the lag for the first item is obviously impossible to calculate. b.) Human subjects' mean naming latencies show the same increase (reprinted from Howard *et al.*, 2006).

The proper behavior of the model requires some reasonable limits on its parameters. For instance, the boosting factor (β) must be greater than 1.00. Otherwise, the booster isn't a booster. The threshold (τ) and deadline (Ω) should be such that, given the value of the boosting factor, words are often selected before the deadline. Finally, the learning rate (η) must be sufficiently large that its effect is not obscured by activation noise (θ).

The learning rate matters because incremental learning underlies the accumulating interference in this model. Each time a word is retrieved, connections from the activated semantic features to the activated words are adjusted. Connections supporting the target word strengthen and those supporting competitors weaken. This learning event promotes repetition priming during the subsequent retrieval of the same target from these semantic features. But if one of these features is instead used to cue a different target word, this same learning means that the new target will be weaker and the previous target – now a competitor – will be stronger. So the new target is retrieved more slowly. When sequentially retrieving several words that all share a semantic feature, the learning that follows each retrieval event makes the retrieved word a strong competitor and further weakens those competitors that have not yet been named. Each newly named word from the category therefore takes a bit longer to retrieve than the previous one. Thus the incremental increase in selection times arises from the incremental adjustments to the semantic-to-lexical connections.

Priming by error-based connection adjustments also makes the model's semantic interference effect insensitive to both time- and item-based lag manipulations. Time-lag insensitivity comes from the fact that all priming in the model is persistent, at least in the sense that the weight changes do not decay with time. Furthermore, there is no residual activation to decay while unrelated items are named. Robustness to long item-lags with intervening unrelated fillers comes from the learning algorithm. The adjustments that follow each trial only affect connections from the specific semantic features that were active during that trial. DOG, for instance, relies on connections from TERRESTRIAL and MAMMAL. Thus, the only thing that would make it more difficult to retrieve DOG via these two features would be changes to connections from them. So as long as

TERRESTRIAL and MAMMAL are not activated en route to lexical selection in any filler trials, no number of filler trials will ever affect DOG's intentional retrieval.

In the introduction, we asserted that incremental learning has both a dark side (cumulative semantic interference) and a light side (repetition priming). However, the first simulation only addressed the dark side. To see how these light and dark sides interact, we turn to a related experiment by Navarrete *et al.* (2008). They added a repetition component to the continuous paradigm. Participants named related pictures separated by fillers , but the third or fourth ordinal position was filled by either a novel related picture or the same picture that appeared in the first or second ordinal position, respectively. The question here was how this repetition would affect naming latencies for subsequent novel related items. Recall that Howard *et al.* (2006) showed that naming latencies increase linearly each time a related item is named. Now if cumulative semantic interference is *type-based*, meaning that what matters is the number of unique related pictures a person has named, then repeating one of the items should not increase RTs at all. However, if cumulative semantic interference is *token-based*, meaning that what matters is the number of times a person has named related pictures (unique or not), then there should be no difference between the interference that accrues from repeating an item and that which accrues from accessing another novel item. They found the latter: Each related retrieval, whether introducing a novel item or repeating an item previously accessed, contributed separately to the cumulative semantic interference.

But what of the model? Is its cumulative semantic interference effect type-based or token-based? To address this question, we now simulate Navarrete *et al.*'s (2008) repetition experiment as Simulation 1b.

Method

Parameters, vocabulary, and training for this simulation were identical to those of Simulation 1a. The testing phase differed slightly, following Navarrete *et al.* (2008).

Testing. The testing phase was identical to that of Simulation 1a, except in two respects. First, we included filler items between the critical items, but did not manipulate item-lag. Second, either the third or fourth ordinal position represented a second cueing of one of the earlier critical items. In one set of lists, the item in the first ordinal position was repeated in the third position. In the other set of lists, the item in the second position was repeated in the fourth position. So a five-position critical item sequence might go DOG BAT DOG WHALE MOLE, or DOG BAT WHALE BAT MOLE.

Analyses. Response times from multiple model subjects were generated as in Simulation 1a. We compared the response times to name items in sequences with a repeated item in ordinal position three to those in sequences with a repeated item in position four.

Results and discussion

Repeating one of the critical items produced additional interference, just like accessing another novel item from the same semantic category (Figure 4). Specifically, we see that the linear increase in response times as a function of ordinal position, apparent in Figure 3, grows normally when one of the ordinal positions is filled by a repeated item. This simulated finding mirrors Navarrete *et al.*'s empirical data. Thus, the model's interference effect, like that in the human data, appears to be token-based. In the model, this happens because error-based learning creates additional interference each time a related item is accessed.

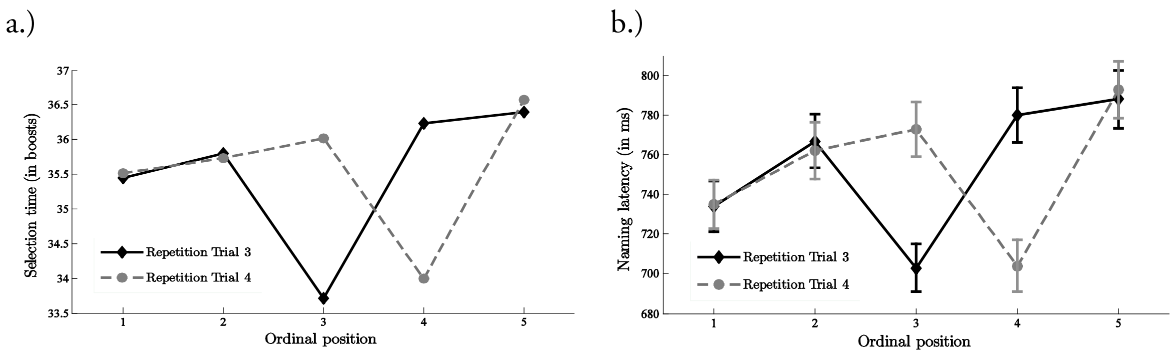

Figure 4. Repeating an item creates cumulative semantic interference similar to that from accessing a novel item. a.) Model-predicted mean selection times from Simulation 1b. b.) Empirical results from Navarrete *et al.* (2008).

It is also worth noting that the model approximates the relative sizes of benefit due to repeating an item and the cost that each similar item imparts, the former being several times larger than the latter. However, we do not want to emphasize exact quantitative properties of the model, as its representations and vocabulary are quite simplified.

Simulation 2 – blocked-cyclic paradigm

In the continuous paradigm, each subject named a large number of pictures just once. The blocked-cyclic naming paradigm, in contrast, requires subjects to repeatedly name a small number of pictures. Here we gauge semantic interference effects by comparing blocks of pictures from the same semantic category, the homogeneous condition, to blocks of unrelated pictures, the mixed condition.

There are two major findings from this paradigm. First, when pictures are presented in the homogeneous condition, they take longer to name than the same pictures presented in the mixed condition (*e.g.* Damian *et al.*, 2001). Second, the magnitude of this semantic blocking effect increases with each cycle (Belke, 2008, Experiment 1; Belke *et al.*, 2005, Experiment 1; Damian & Als, 2005, Experiment 4; Schnur *et al.*, 2006, Experiment 1). So the blocked-cyclic paradigm's semantic interference effect grows with each cycle, similar to the way that the continuous paradigm's semantic interference effect grew with each related item. Our second simulation simulates this blocked cyclic procedure.

Method

Model parameters were identical to Simulation 1, and are given in Table 1.

Vocabulary. The vocabulary for this and the remaining simulations consisted of 12 semantic features mapped onto 36 words. Each feature cued exactly six words, and each word was cued by the intersection of exactly two features. Training followed the format described in the introduction to the simulations.

Testing. Two parallel testing phases simulated the homogeneous and mixed conditions of a blocked-cyclic naming experiment. In each condition, a set of six words was repeatedly cued for four cycles, for a total of 24 trials, with words ordered randomly within each cycle. The homogeneous condition used six words from a single category, so they all shared one of their two semantic features. Words in the mixed condition each represented a different category, so none shared a feature with any other word in the set.

Both testing phases began with the same trained connection weights. But learning then continued separately in each condition. This learning proceeded at the same rate (η) as during the shared training phase. So each replication tests the same trained network in both conditions, with no carryover from one condition to the next.

Analyses. We report selection times in terms of the mean number of 'boosts' for each item position in each condition. As before, errors of omission and commission occurred in less than 1% of the trials and were excluded from the analyses. [3]

Results and discussion

The trial-by-trial predictions (Figure 5a) show both repetition priming and cumulative semantic interference. Repetition priming occurs in both the homogeneous and mixed conditions;

---

[3] Errors were rare, but we note that they were relatively more frequent in the homogeneous condition than the mixed, and in both conditions, the errors were predominantly omissions. Given the infrequency of errors in this simulation, however, we do not discuss them further here.

each time a word is cued, it is selected more quickly. This is because the incremental learning that follows each retrieval facilitates future retrievals of the target word by strengthening its connections and weakening connections from its features to competitor words. Cumulative semantic interference only occurs in the homogeneous condition, creating an incremental increase in selection times within each cycle. This is the same pattern that we saw in Simulation 1, because it arises from the same process. However, in this simulation, the relation between repetition priming and cumulative semantic interference becomes more apparent. Each time one word gets stronger, its competitors get weaker. When that word is retrieved again, it is primed. When one of the competitors is retrieved, though, it is subject to interference. Because repetition priming and cumulative semantic interference are both at work in the blocked-cyclic paradigm we get a saw-toothed function for selection times in the homogeneous condition (Figure 5a), and a small decrease in the per-cycle mean selection times (Figure 5b). A smaller per-cycle decrease is precisely what we see in human data for the same task (*e.g.* Figure 5c, from Schnur *et al.*, 2006, Experiment 1). So it appears that the model successfully extends to blocked-cyclic naming.

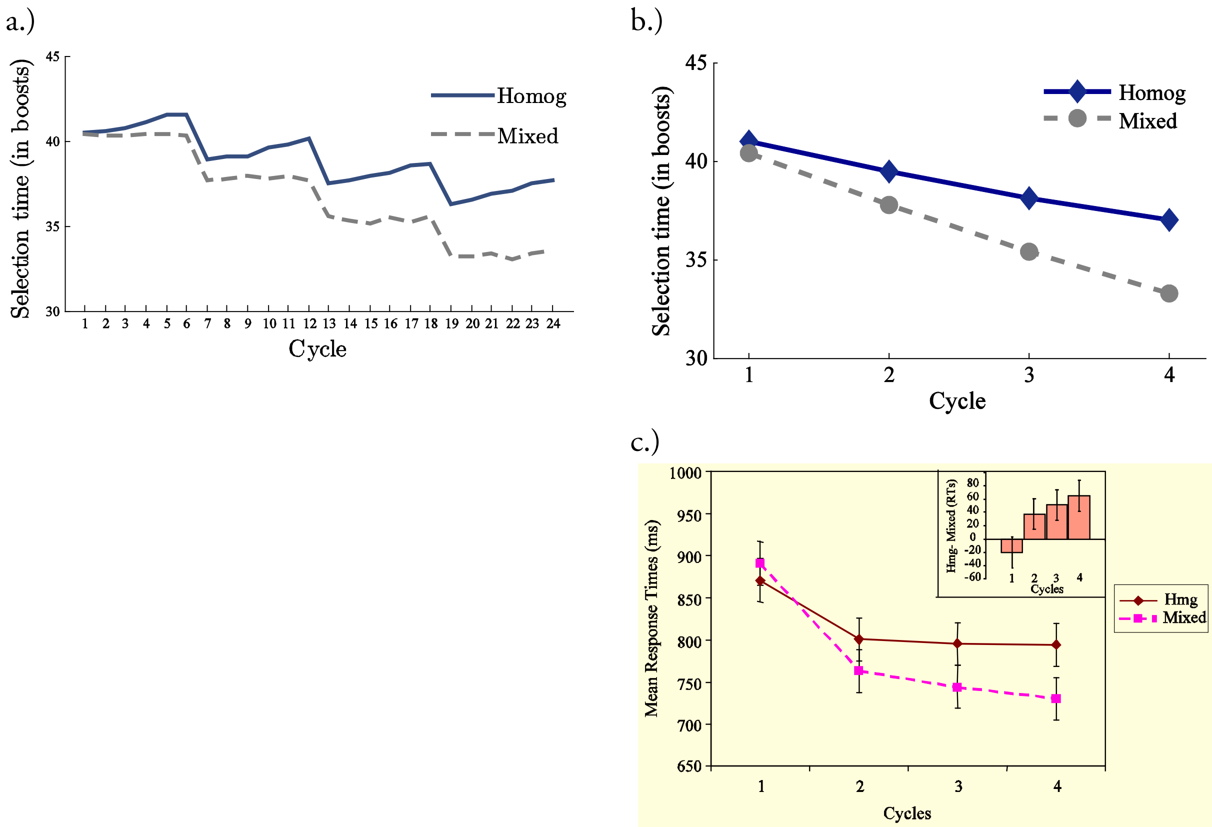

Figure 5. Simulated lexical selection times from Simulation 2 mirror human subjects' naming latency patterns for the blocked-cyclic paradigm. a.) Mean simulated lexical selection times across four cycles through sets of six words, for 10,000 networks. b.) Predicted per-cycle means, derived from (a), for comparison with the human data presented in (c). c.) Subjects' per-cycle mean naming latencies across four cycles, from Schnur *et al.* (2006).

Our model predicts some attenuation of the repetition priming across cycles, because of the learning algorithm. Each retrieval is a learning event, thus decreasing error for the next retrieval. And since the magnitude of each weight change depends on the magnitude of the erroneous activation for the relevant output node, this error-reduction reduces the weight change (and hence repetition priming) that will result from the next repetition. We acknowledge, though, that with our current parameter values for the learning and decision processes, the non-linearity is smaller than it is in the human data.

There is one feature of the data that this simulation does not exhibit. While our main effect of blocking begins in the first cycle, Belke *et al.* (2005) observed that in human data it tends not to appear until the second cycle. In fact, human data sometimes shows a brief semantic facilitation effect, for example in the first cycle depicted in Figure 5c. However, there is evidence that this early facilitation represents a conscious strategic process rather than an integral part of lexical retrieval, and is therefore beyond the scope of a model of lexical access. Three findings support this conclusion. First, Wheeldon and Monsell (1994) reported a similar brief semantic facilitation in a continuous paradigm experiment, and argued that it followed a different time course than the longer-lasting interference effect, suggesting a separate process. Extrapolating to the blocked-cyclic paradigm, a brief semantic facilitation could delay the appearance of an interference effect until the second cycle. Second, Damian and Als (2005, Experiment 4) demonstrated that the semantic interference appears in the first cycle, and grows thereafter, if homogeneous and mixed sets are interwoven a single block. This suggests that subjects' expectations may play a role in the early facilitation. Finally, and most conclusively, Belke (2008, Experiment 1) showed that adding a working memory load to the blocked-cyclic naming paradigm led to semantic interference in the first cycle, with no sign of early facilitation. Since the dual-task disrupts semantic facilitation while leaving the growing semantic interference effect intact, we can conclude that the facilitation represents a resource-demanding process that is distinct from whatever underlies the cumulative semantic interference effect and is not integral to the process of lexical retrieval.

Simulation 3 – generalization of interference to new pictures

An interesting empirical property of the semantic blocking effect is that it extends to naming new pictures from the same category (Belke *et al.*, 2005, Experiment 3). For instance, repeatedly naming a small set of birds, such as CROW, FINCH, and GULL, creates a substantial semantic blocking effect that will carry over, without interruption, to naming an entirely new set of birds. Belke *et al.* (2005) argued that this carryover reflected a refractory behavior where related words become temporarily inaccessible due to residual activation and/or inhibition.

In order to test whether a learning model could account for Belke *et al.*'s finding without representing residual activation or inhibition, we now test the model using their procedure. They compiled sets of pictures of from one category, and had subjects name half of these pictures for four cycles before switching over to name the other half for four cycles. For example, given a set of birds {CROW FINCH GULL JAY ROBIN SPARROW} subjects might name CROW, FINCH, GULL, FINCH, CROW, GULL, FINCH, GULL, CROW, GULL, FINCH, CROW, and then immediately switch over to JAY, ROBIN, SPARROW, ROBIN, SPARROW, JAY, ROBIN, JAY, SPARROW, JAY, SPARROW, ROBIN. The crucial question is whether the naming latency difference between related and unrelated conditions that develops while cycling through the first subset will continue in the first cycle of the new subset.

Methods

We followed the methods of Simulation 2, with just one change in the testing procedure. Instead of cuing one set of six words for four cycles during testing, we now cued one set of three words for four cycles, and then a set of three different words for four additional cycles. In the homogeneous condition, all six words shared a single semantic feature. None of the six words in the mixed condition shared any semantic features.

Model parameters were identical to Simulation 1 and 2, and are given in Table 1.

Results and discussion

The simulated blocking effect carried over from the first to the second subset without interruption (Figure 6a). Cumulative semantic interference built up while naming the first subset in the first four cycles, resembling the findings from Simulation 2. Upon switching to the second subset, in cycles five through eight, the interference effect continued with the entirely new items. The blocking effect in the fifth cycle exceeded that in the first cycle, demonstrating that the accumulated semantic interference transferred to new items from the same category.

In the model, interference generalizes to new items from the same category because the incremental weight changes that follow each selection affect all competitors. Indeed, this is required for the model to simulate the results of Howard *et al.*'s (2006) continuous paradigm in which same-category items do not repeat. Importantly, the fact that our simulated data (Figure 6b) mirrors Belke *et al.*'s (2005) human data (Figure 6c) demonstrates that the human results do not depend on the residual activation or inhibition that are normally associated with refractory behaviors (*e.g.* Forde & Humphreys, 1997; McCarthy & Kartsounis, 2000). Rather, this behavior can arise from persistent incremental learning.

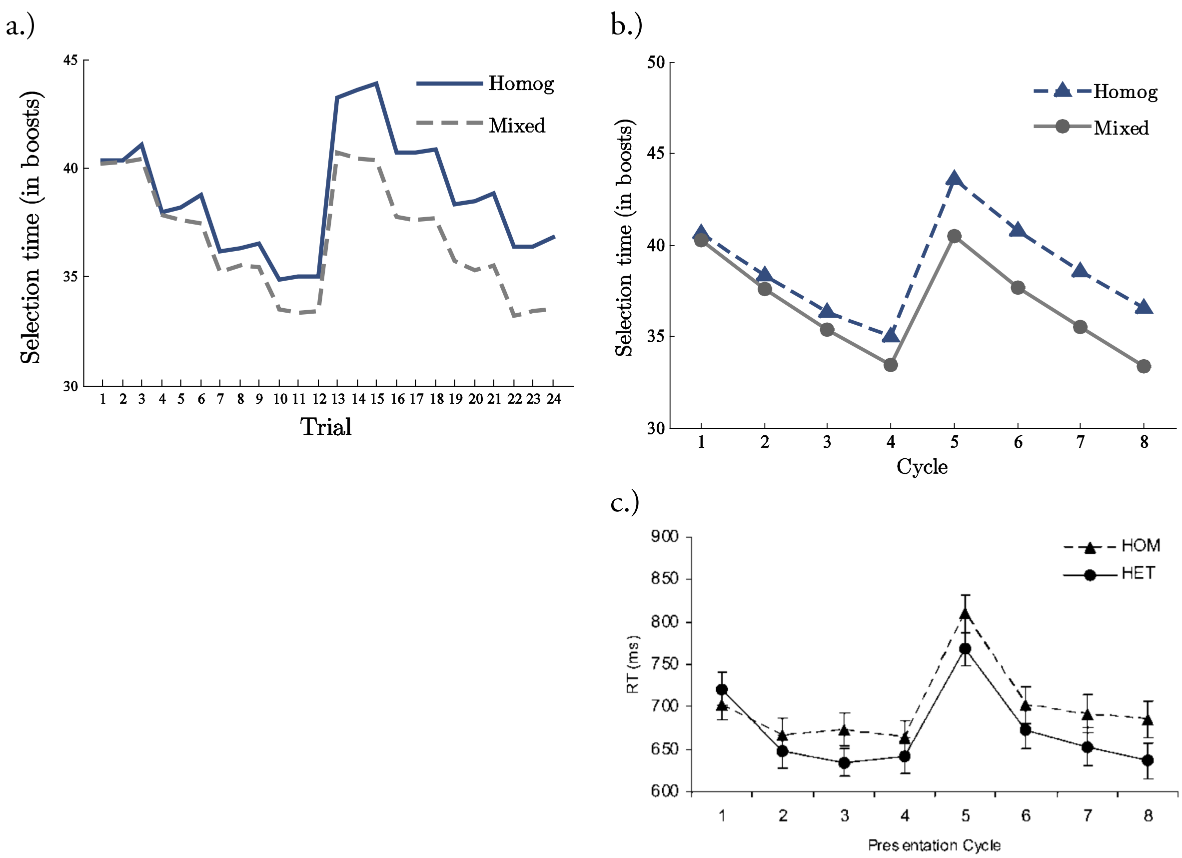

Figure 6. Results from Simulation 3. a.) Lexical selection times for items in homogeneous and mixed sets presented in three-trial cycles. The second subset replaced the first after the twelfth trial. b.) Per-cycle mean selection latencies, derived from (a), for comparison with (c). c.) Human subjects' per-cycle mean naming latencies, from Belke *et al.* (2005).

Simulation 4 – aphasic errors

Semantic blocking manipulations elicit longer picture naming latencies from healthy subjects, and they also affect lexical selection errors made by individuals with aphasia (Schnur *et al.*, 2006). Semantic errors and omissions become increasingly likely in the homogeneous condition, compared to the mixed baseline (Schnur *et al.*, 2006, Experiment 2). Other types of errors (*e.g.* phonological or unrelated errors) do not show such effects. Furthermore, when patients name pictures incorrectly, their within-set substitutions tend to match the words that they have used more recently, and there was no difference between this perseverative recency effect at 1-s and 5-s inter-stimulus intervals (Hsiao *et al.*, 2009).

To simulate the patient error effects, we need a theory of the differences between aphasic and nonaphasic lexical access in this task. Here we follow the approach of several researchers (*e.g.* Dell *et al.*, 1997; Rapp & Goldrick, 2000) and attribute impaired performance to the alteration of one or more model parameters, but without changing any of the model's processes. Specifically, we increase the activation noise parameter, with the result that errors become much more likely. Rapp and Goldrick specifically simulated aphasia by activation-noise lesions, and Dell *et al.* demonstrated that noise lesions mimic the decay lesions that they used to simulate some patients.

Schnur *et al.* (2006) compared unimpaired (Experiment 1) and aphasic (Experiment 2) performance in the blocked-cyclic paradigm by running both groups in essentially the same experiment. We do likewise, repeating Simulation 2, which we had based on their Experiment 1 procedure, with noisier lexical activations. Our analyses follow Schnur *et al.* (2006, Experiment 2) and Hsiao *et al.* (2009), so that we can compare the model's predictions directly to the data. So, if successful, the predictions should show Schnur *et al.*'s semantic blocking effects for semantic errors and omissions, and Hsiao *et al.*'s recency gradient for perseveration errors.

Methods

The vocabulary, training, and testing followed Simulation 2, only increasing the amount of noise in the lexical activations from 0.5 to 1.0. The new parameters are given in Table 2.

Table 2. Model parameters for Simulation 4. Except for Activation noise, these parameters are identical to those given in Table 1.

Parameter	Value
Learning rate (η)	0.75
Activation noise (θ)	1
Boosting rate (β)	1.01
Threshold (τ)	1
Deadline (Ω)	100

The analyses followed those that Schnur *et al.* (2006) and Hsiao *et al.* (2009) used for their human data. Schnur *et al.*'s (2006, Experiment 2) major findings concerned omissions and semantic errors, so our analyses focused on these as well. As described in the Model Description section, omissions occurred when the activations of lexical nodes were so similar that no winner could be selected before the deadline. Incorrect selections were classified as either *semantic errors* or *other errors*, according to whether the target and error shared a semantic feature. Following Schnur *et al.*, we calculate per-cycle means for each error type.

Perseveration error analyses followed Hsiao *et al.* (2009). These errors occurred when an erroneously selected word matched another target in the block whose name had been selected previously. These errors were rare in the mixed condition, and so perseveration error analyses were restricted to the homogeneous condition. For each such error, we recorded the number of trials back that the erroneously selected word was most recently selected. So if a dog was named as DOG, a bat was named as BAT, and a whale was then named as DOG, then the whale→ DOG error would be counted as a perseveration at lag-2 because DOG was last produced two items back. Actual target-error pairs from the testing phase were then randomly reshuffled within each cycle, with the results coded as above, in order to estimate the probability that such a lag distribution might occur by chance. Following Hsiao *et al.*, (2009, pp.136-137), and Cohen and Dahaene (1998, p. 1643) before them, we thus derived a chance-corrected estimates of perseveration frequencies at each lag.

Hsiao *et al.* (2009) reported perseveration functions for each of two ISIs. Our model assumes no time-based decay, and should therefore predict the same effects for any ISI. However, we give some measure of the variability of our predicted perseveratory effect by presenting data from four 10,000-block replications.

Results and discussion

Doubling the activation noise substantially increased the predicted error rates, producing errors in 12.4% of homogeneous trials and 10.9% of mixed trials.[4] Errors of commission were largely semantic, with less than 0.01% unrelated errors occurring in either condition.[5] The homogeneous condition elicited more semantic errors (Figure 7a) and omissions (Figure 7b) than the mixed condition. These effects increased across cycles, recalling Schnur *et al.*'s (2006, Experiment 2) aphasic patients' error patterns (Figure 7c).

---

[4] These totals included 2.7% semantic errors and 8.9% omissions overall, but these relative proportions depend on the values of the selection deadline and activation noise parameters. That is, a greater value for the deadline would decrease the number of omissions and increase the number of semantic errors, so the fact that there are more omissions than semantic errors is merely an accident of parameters. Therefore we focus on these error types individually, but avoid comparing them.

[5] Schnur *et al.* (2006) also reported that "Other errors" (*e.g.* phonological slips and unrelated word errors) become increasingly likely in the Mixed condition, relative to their frequency in the Homogeneous condition, as shown in Figure 7c. Although our model does not currently generate such errors, adding some noise to semantic-input activations increases the likelihood of unrelated word errors, and these do in fact show such reverse blocking effects.

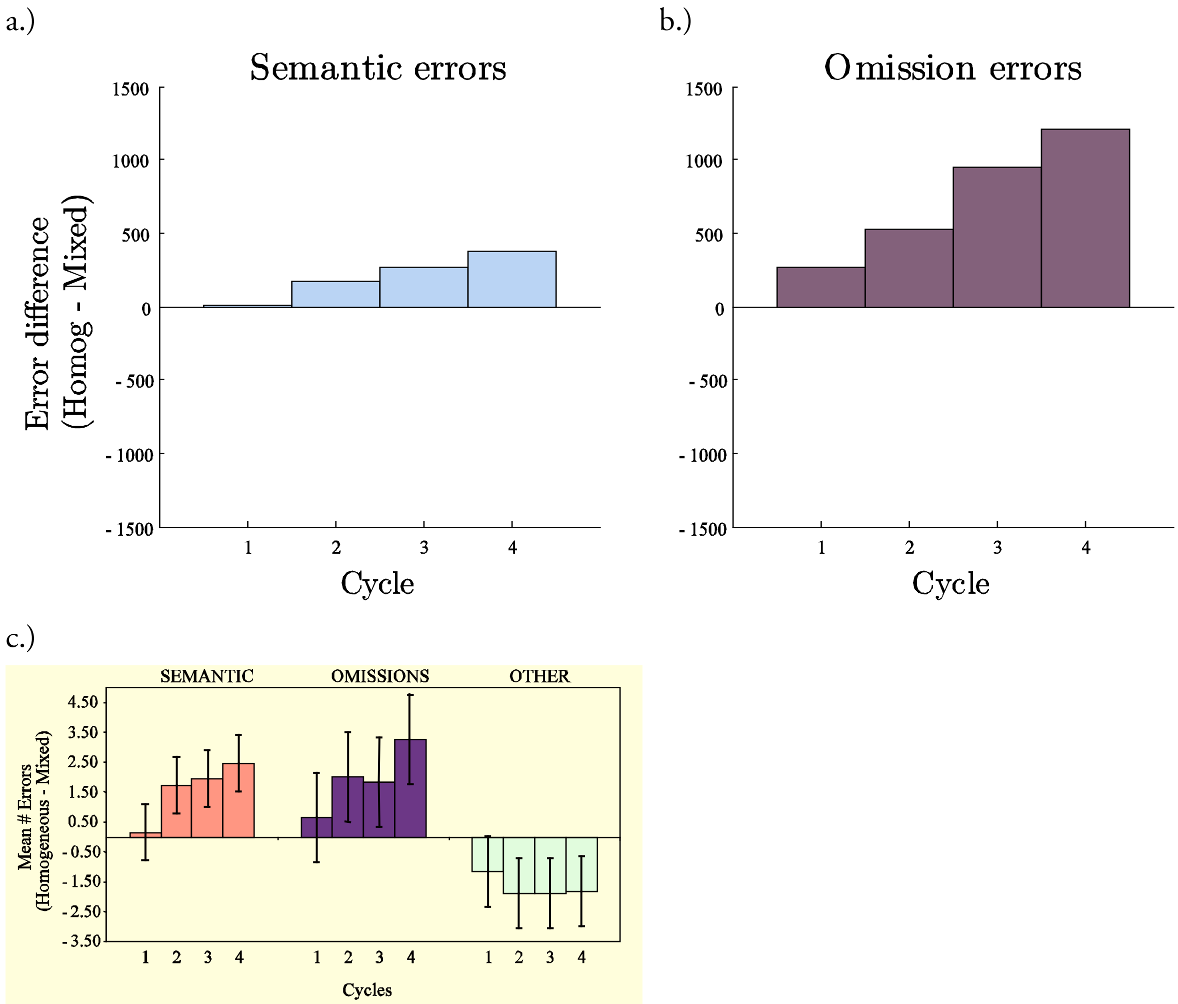

Figure 7. Predicted semantic error and omission patterns from Simulation 4, compared to patient error data for the same task. a.) Simulated semantic errors become increasingly more frequent in the Homogeneous condition, compared to the Mixed condition. b.) The differences in the omission error rates between the Homogeneous and Mixed conditions also increase across cycles. c.) Patient error data from Schnur *et al.* (2006).

Within-set substitution errors accounted for 51% of the semantic errors in the homogeneous condition, in agreement with Schnur *et al.* who reported 53%. In our simulated data, most (82.2%) of these within-set substitutions were identified as perseverations. As illustrated in Figure 8a, the tendency to perseverate a word exceeds chance at short lags, but falls to near-chance levels thereafter. This pattern recalls Hsiao *et al.*'s (2009) report that aphasic patients' were more likely to perseverate words that they had used more recently (Figure 8b).

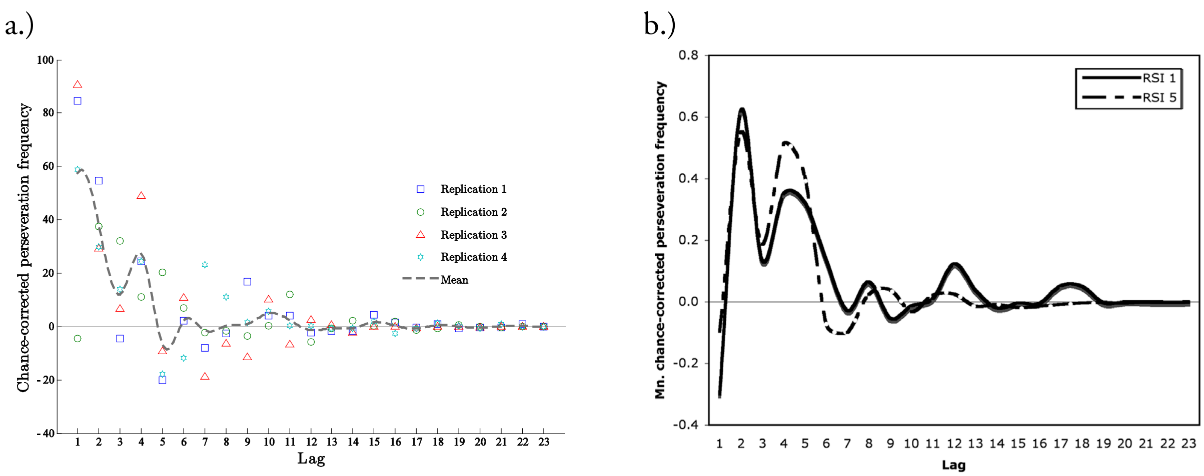

Figure 8. Chance-corrected perseveration frequencies in the homogeneous condition, plotted as a function of item-lag. a.) In Simulation 4, perseveration errors at short lags are more frequent than would be expected by chance. b.) Patients' perseverations, from Hsiao *et al.* (2009), are also more frequent at short lags.

One apparent difference between our predicted lag function (Figure 8a) and Hsiao *et al.*'s (2009) mean lag function (Figure 8b) is that our model is most likely to perseverate words at lag-1, whereas their patients were not. However, Hsiao *et al.* noted that this was not actually a reliable feature of their patient data, and actually excluded lag-1 from their statistical analyses. Roughly half of their patients showed this lag-1 dip, while half did not. Given this lack of consistency across patients, we follow one of Hsiao *et al.*'s proposed explanations, positing that their lag-1 dip may merely reflect a repetition-avoidance strategy that some patients chose to employ. Such a strategy has some precedence in non-patient work. For instance, Anderson and Neely (1996) cite two early naming-to-definition experiments demonstrating that preceding a trial with a semantically related prime produced facilitation if the prime was never the correct answer, but interference if the prime was sometimes the correct answer (Brown, 1979; Roediger, Neely, & Blaxton, 1983). This change suggests that a brief facilitation or anti-perseveration effect may merely index participants' sensitivity to the sequential-statistical properties of the testing paradigm (*i.e.* participants learn to expect that no picture will be immediately repeated, allowing them to constrain their predictions for the next item). Consistent with applying such an interpretation to this patient data, Hsiao *et al.* noted that targets only immediately repeated in 2.2% of their trials, so a repetition-avoidance strategy would have helped performance on the vast majority of trials.

While the precise omission-to-commission error ratio depends on the selection deadline, all of the error predictions that we have presented here hold true for a wide range of parameters. As long as there is sufficient noise to produce errors, but not so much that selection becomes completely random, the semantic error and perseveration effects emerge. And as long as the selection deadline parameter cuts off some, but not all, lexical selections, the omission effects also show up. Thus, the effects are not just a matter of fitting the model to the data.

So where do these error effects come from? Since activations return to baseline at the end of each trial, any errors necessarily result from the interaction of activation noise and persistent changes in connection weights. Recall that omissions occur when target and competitor activations are too similar for a winner to be resolved before the selection process times out. Semantic errors happen when a non-target word from the target category becomes substantially more active than its competitors, including the target. The conditions underlying these errors are rare, under normal circumstances. But adding noise to the model increases the likelihood that errors will arise from target and competitor net inputs that are merely "somewhat similar". And these somewhat similar net inputs drove the selection latency effects in Simulations 1-3. Thus, the error effects derive from the same basic process that led to latency effects in the previous simulations: incremental learning.

Incremental learning (Equation 7) also drives the recency effect for perseverations. Here it may be helpful to work through an example. Let us imagine a patient sequentially naming three pictures, DOG, BAT, and WHALE, that all share a MAMMAL feature. When he first names DOG, the MAMMAL feature will support each of these words more-or-less equally. After naming DOG, however, the link from MAMMAL to DOG is strengthened and the links from MAMMAL to BAT and to WHALE are weakened. Thus, when the patient encounters the picture BAT, the MAMMAL feature will activate DOG more strongly than BAT or WHALE, making him more likely to name the picture as DOG than as WHALE. This is the primary basis of the recency effect:

words tend to form stronger competitors when their connections from shared features have been strengthened more recently, and weakened less recently.

We should note, however, that other learning models could explain a recency effect, provided that they can explain errors in addition to response times. In Howard *et al.*'s (2006) priming model, for instance, a semantic-to-lexical connection is strengthened by a fixed increment each time its associated word is selected. Thus, the DOG-BAT-WHALE explanation for recency outlined above is also consistent with models that learn just by strengthening each item by a fixed amount.

Does our error-based learning algorithm offer any advantage in explaining the recency gradient for perseverations? To investigate this question, we repeat the perseveration simulation with three different learning rules applied at testing. First, we present results from a model using our standard delta-rule learning algorithm. To gauge the contribution of error-based learning over other approaches, we compare these results to those from a version of the model that uses a non-error-based connection-strengthening algorithm. Finally, to assess how much of the perseveration effect might reflect random variation, structural peculiarities of the blocked-cyclic paradigm, or persistent biases that cannot be attributed to either learning process, we present results from a version that does not learn at all during the testing phase.

Networks were trained as normal, using a delta-rule learning algorithm and the parameters listed in Table 2. Perseveration simulations followed, as described above, but with two changes.

First, after training each network as normal, we applied one of three learning algorithms at the time of testing. The *delta-rule simulation* implemented the same error- and activation-proportional learning algorithm (Equation 7) that we have used elsewhere in this paper. The *priming simulation* instead implemented a fixed-increment unsupervised learning algorithm. Each time a word was selected, connections to it from the active features were strengthened by a fixed increment.[6] Finally, the *non-learning simulation* did not strengthen or weaken any connections during the test phase.

Second, to preclude the possibility that omission errors might obscure differences in the perseveration effects, we removed the timeout parameter (Ω) during the test phase. So these simulations do not allow for omission errors, meaning that they will not be directly comparable to our previous simulations. However, removing the omission behavior allows us to keep all parameters constant across simulations, without further concern for matching error and omission rates.

The delta rule and priming simulations produced equivalent rates of semantic errors (5.79% and 5.80%, respectively). This means that we can compare the strength of these recency effects more or less directly. The non-learning simulation produced slightly more semantic errors (6.60%),

---

[6] We set this increment equal to the average weight increase for the first target in the test phase of the delta rule simulation (*i.e.* $\Delta w_{ij} = \eta\left(a_i\left(1-a_i\right)\left(d_i - a_i\right)\right)a_j = 0.75 \times \left(0.72 \times \left(1-0.72\right)\left(1-0.72\right)\right) \times 1 = 0.042$).

owing to the fact that its mappings never improved after training. As in the previous simulation, non-semantic errors were exceedingly rare, so we describe perseveration effects only for the homogeneous condition.

In the delta-rule simulation, perseverations were again much more likely to match recent responses (Figure 9). The priming simulation also produced some recency effect for perseverations, though this was much smaller. And the non-learning simulation demonstrates that both effects exceed what we might expect to emerge from any persistent biases in model weights. It is thus clear that incremental learning is crucial for the recency effect.

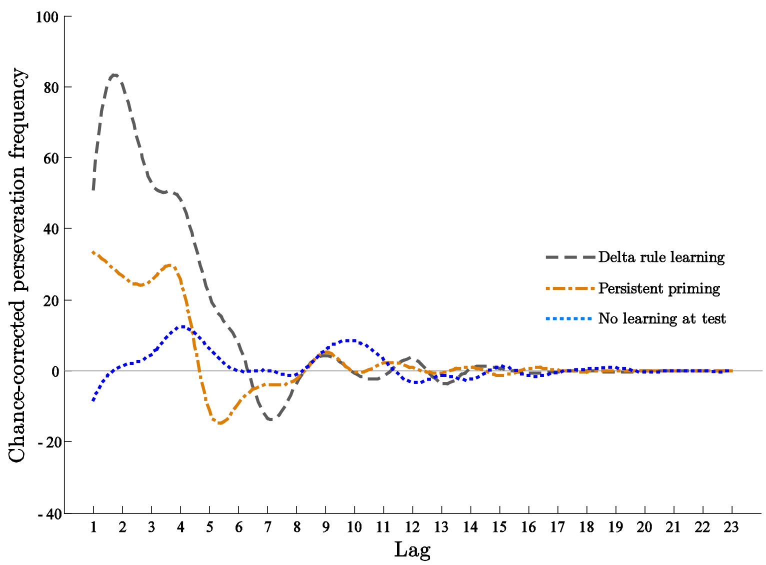

Figure 9. Chance-corrected perseveration frequencies as a function of item-lag, using three different learning rules during the testing phase: 1.) Error-based delta rule learning, 2.) Non-error based fixed-increment strengthening (priming), and 3.) No learning. As in Figure 8a, each spline represents a mean of four 10,000-block replications, though here we have omitted the individual data points in favor of visual clarity.

The delta rule creates the stronger recency gradient by causing connection weakening as well as connection strengthening. As in the priming account, the more recent a potential competitor is, the more likely that it will have been strengthened more times by the delta rule in the blocked-cyclic paradigm. However, the delta rule amplifies the gradient because the more recent a competitor is, the fewer times it will have been weakened by the action of the rule, since its last strengthening. So, in the sequence BAT, DOG, WHALE, when WHALE is the target, both BAT and DOG have both been strengthened, and so both are potential perseverates. But DOG, the more recent item, is a more likely error for WHALE because BAT was weakened during the DOG trial. BAT was especially active at that time (because it had just been strengthened), and hence the delta rule will have weakened it in proportion to its activation. Thus, DOG will be a more powerful competitor than BAT during the WHALE trial.

CHAPTER 4: MODEL ANALYSES

Simulation 5 – facilitation versus inhibition

Our learning algorithm specifies changes in connection weights whenever there is some degree of error and it does not distinguish between the learning that strengthens connections and the learning that weakens them. For example, we saw that, in the analysis of the model's perseverations, both strengthening and weakening of connections contribute to the effect of recency on these kinds of errors. But both processes arise from the same weight-change equation.

Several explanations of cumulative semantic interference, however, attribute it either to a process that strengthens competitors or a process that inhibits targets, but not both, and thus there is a debate about whether this effect is, at core, one of facilitation or inhibition. This is analogous to the occlusion versus inhibition debate in the retrieval-induced forgetting literature. The central idea in the facilitation (or occlusion) explanation is that retrieving an item facilitates its future retrieval, making it a stronger competitor when related words are cued. Thus, lexical competition grows more intense when retrieving words in a homogeneous context because each new target faces a growing number of prepotent competitors. This facilitatory account dominates much of the cumulative semantic interference literature. Damian *et al.* (2001) and Belke *et al.* (2005) both suggested temporary facilitation mechanisms, while Wheeldon and Monsell (1994), Damian and Als (2005), Schnur *et al.* (2006), and Howard *et al.* (2006) argued for more persistent facilitation. Some memory researchers have also offered facilitation-based explanations of retrieval-induced forgetting (*e.g.* Macleod *et al.*, 2003; McGeoch, 1932; Mensink & Raaijmakers, 1988).

However, inhibition has played a more prominent role in discussions of retrieval-induced forgetting (*e.g.* Anderson, 2003; Anderson *et al.*, 1994; Norman *et al.*, 2007; Postman *et al.*, 1968). The basic idea here is that retrieving a target necessarily harms its competitors, whether by inhibiting concepts directly (*e.g.* Postman *et al.*, 1968) or merely making them less accessible via retrieval cues (*e.g.* Melton & Irwin, 1940). For cumulative semantic interference, an inhibition account would entail the active suppression of lexical competitors during each naming attempt, with the lasting consequence of making them less accessible. Each new target in a homogenous set then becomes increasingly difficult to retrieve by virtue of being repeatedly suppressed each time one of its competitors is retrieved instead. Though pure inhibitory explanations are relatively rare in discussions of cumulative semantic interference, we note one example in McCarthy and Kartsounis' (2000) report of patients' omission errors.

So in all these cases we see semantic interference attributed to priming that either strengthens recent targets or weakens recent competitors, but not both. And, as Schnur *et al.* (2006) pointed out, these two mechanisms generate different predictions. We can use our model to separate the strengthening and weakening processes to see what each contributes to predicted cumulative semantic interference effects. In this simulation, we selectively apply either connection strengthening or connection weakening while disrupting the other. If the model realizes the notion that cumulative semantic interference arises from strengthening targets, then a version of the model that only allows for connection strengthening should exhibit the interference effect. If cumulative semantic interference arises from weakening competitors, then it should be seen in a model version that only allows for connection weakening.

Methods

Model parameters were identical to Simulations 1-3, and are given in Table 1.

Methods for these simulations followed Simulation 2 exactly, but with one additional manipulation during the testing phase. Each replication began with the standard training phase, where learning was applied according to the delta rule. So the training was identical to Simulation 2. During the testing phase, however, the learning algorithm was modified to only apply either weight increases or decreases. In Simulation 5a, these weight changes only increased connection weights (Equation 8) and in Simulation 5b, they only decreased the weights (Equation 9).

Equation 8
$$\Delta w_{ij} = \begin{Bmatrix} \Delta w_{ij} & \text{if } \Delta w_{ij} \geq 0 \\ 0 & \text{if } \Delta w_{ij} < 0 \end{Bmatrix}$$

Equation 9
$$\Delta w_{ij} = \begin{Bmatrix} \Delta w_{ij} & \text{if } \Delta w_{ij} \leq 0 \\ 0 & \text{if } \Delta w_{ij} > 0 \end{Bmatrix}$$

Results and discussion

The simulations that only increased or only decreased weights had distinct effects on simulated lexical selection times (Figure 10). Recall that, in Simulation 2, selection times in the homogeneous condition increased within a cycle and decreased across cycles, while the blocking effect (homogeneous-mixed difference) increased both within and across cycles (Figure 10a). This pattern represents the combined effects of increased and decreased weights. In a, weight increases alone produced a step-like decrease for selection latencies in both homogeneous and mixed conditions (Figure 10b). This is what we would expect from a repetition priming effect. However, as shown by the close overlap of curves, the semantic blocking effect was minimal, suggesting that weight increases play a much stronger role in repetition priming than cumulative semantic interference. Learning through weight decreases alone (b) generated a robust semantic blocking effect with only a very weak repetition priming effect (Figure 10c). Selection times steadily increased in the homogeneous condition and decreased only slightly in the mixed condition.

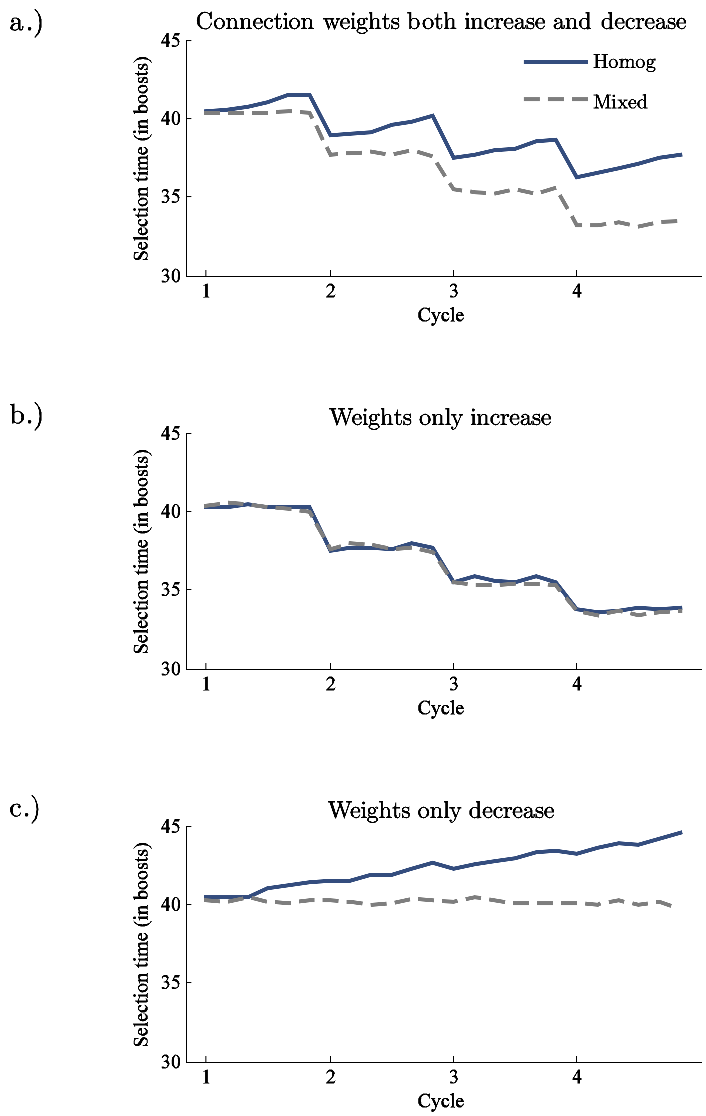

Figure 10. Strengthening connections to target words produces a repetition priming effect, while weakening connections to competitors produces a semantic interference effect. a.) Applying both weight increases and decreases (Simulation 2) generates selection latencies resembling a saw-tooth pattern. b.) Increasing connection weights to target words (a) decreases selection times in each cycle, creating a step-like pattern. c.) Decreasing weights to competitors (b) continually increases selection times in the homogeneous condition.

Errors were rare, occurring in less than 1% of the trials in either simulation. But, as in Simulation 2, these errors were predominantly omissions and were relatively more likely in the homogeneous conditions than in the mixed. As one would expect from Figure 10, this blocking effect was stronger when learning via weight decreases than when learning via increases. Moreover, learning via weight decreases produced a blocking effect for omission errors that increased across cycles, thus tracking the interference effects in response times seen in Figure 10c.

From these findings, we can characterize the model's repetition priming as a facilitatory (strengthening of weights) effect, and its cumulative semantic interference as an inhibitory (weakening of weights) effect. Thus, the model's account clearly differs from a popular characterization of cumulative semantic interference as a facilitation-based effect, adopting instead an inhibition explanation that better resembles recent accounts of retrieval-induced forgetting (e.g. Norman et al., 2007).

Simulation 6 – competition in selection and learning

In the previous simulation we explored the contributions of weight strengthening and weakening to repetition priming and cumulative semantic interference. Crucially, we found that the weight-strengthening version of the model failed to produce cumulative semantic interference. This finding is unexpected because this restricted version still implements Howard *et al.*'s (2006) necessary and sufficient principles of shared activation, priming, and competitive selection. Shared activation was clearly present in the distributed semantic features. The experience-based weight increases facilitated access to recent targets, fulfilling Howard *et al.*'s priming function. And the lexical selection algorithm was competitive by virtue of implementing a differential threshold. So why didn't it work?

The answer may lie in the notion of competitive selection that our model implements. In a nutshell, our version of competitive selection may not have been competitive enough. Recall that our selection rule (Equation 6) compares each word's activation to the mean activation of all its competitors. Facilitation-based explanations of cumulative semantic interference often focus on the role of individual prepotent competitors in prolonging lexical selection times. But, with a large vocabulary that includes lots of unrelated words, using an average activation for the differential threshold could minimize the impact of these prepotent competitors. That is, having a competition between the target and the mean activation of other words is quite close to having a competition between the target and some absolute standard or threshold, which is not technically a competition.

If this analysis is correct and lexical selection in our model truly functioned as a non-competitive process, then by Howard *et al.*'s logic the full version of the model with both strengthening and weakening of weights should not have been able to generate cumulative semantic interference in any of the previous five simulations. But it did. In the following simulations we explore this puzzle and offer a solution, a solution that forces us to change our conceptions of what constitutes competition. Specifically, we first verify our intuition that the competition involving the mean of the competitors is functionally like that of using a non-competitive absolute-threshold decision rule by repeating Simulation 5 with an absolute threshold. From this, we then develop a new hypothesis about the nature of competition and test one of its ramifications by repeating Simulation 5 with a clearly competitive selection rule in which the target is compared against only its most potent competitor.

Methods

We repeated Simulation 5 with an absolute (*i.e.* non-differential) threshold, replacing Equation 6 (reprinted below as Equation 8) with Equation 9, below. Thus, lexical selection becomes a non-competitive horse race. Competitors' activations do not affect selection times, and the first word to reach the threshold wins.

Equation 8 $$t_{selection} = \log_{\beta}\left(\frac{\tau}{a_{i\,t_1} - a_{others\,t_1}}\right)$$

Equation 9 $$t_{selection} = \log_{\beta}\left(\frac{\tau}{a_{i\,t_1}}\right)$$

Otherwise, this simulation replicates Simulation 5 exactly.

Results and discussion

The results from the absolute threshold model (Figure 11) were virtually indistinguishable from those from the differential-threshold model that we used in the first five simulations (Figure 10). The full-learning condition produced the familiar saw-toothed function (Figure 11a), indicating repetition priming and cumulative semantic interference. Weight increases carried a repetition priming effect (Figure 11b), while decreases carried the interference effect (Figure 11c).

In fact, the absolute threshold version of the model is not just indistinguishable from the differential-threshold model with regard to the saw-tooth function. It can simulate every one of the phenomena that the differential-threshold model did. Figure 12 shows the results of new simulations that repeat Simulations 1-4 while replacing the competitive selection rule with a non-competitive absolute threshold.

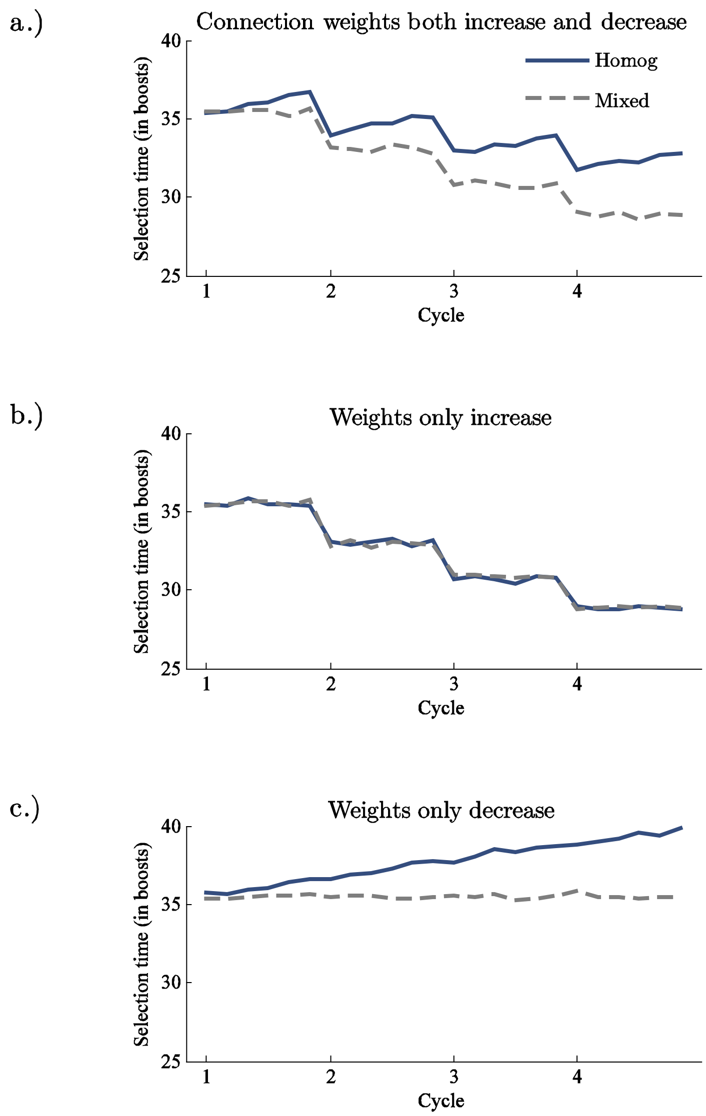

Figure 11. Predictions from the non-competitive selection model (Equation 9) mirror those from our competitive model in Simulation 5. a.) Implementing both weight increases and decreases produces the familiar saw-toothed pattern. b.) Weight increases alone produce the step function associated with repetition priming. c.) Weight decreases still lead to a cumulative semantic interference effect.

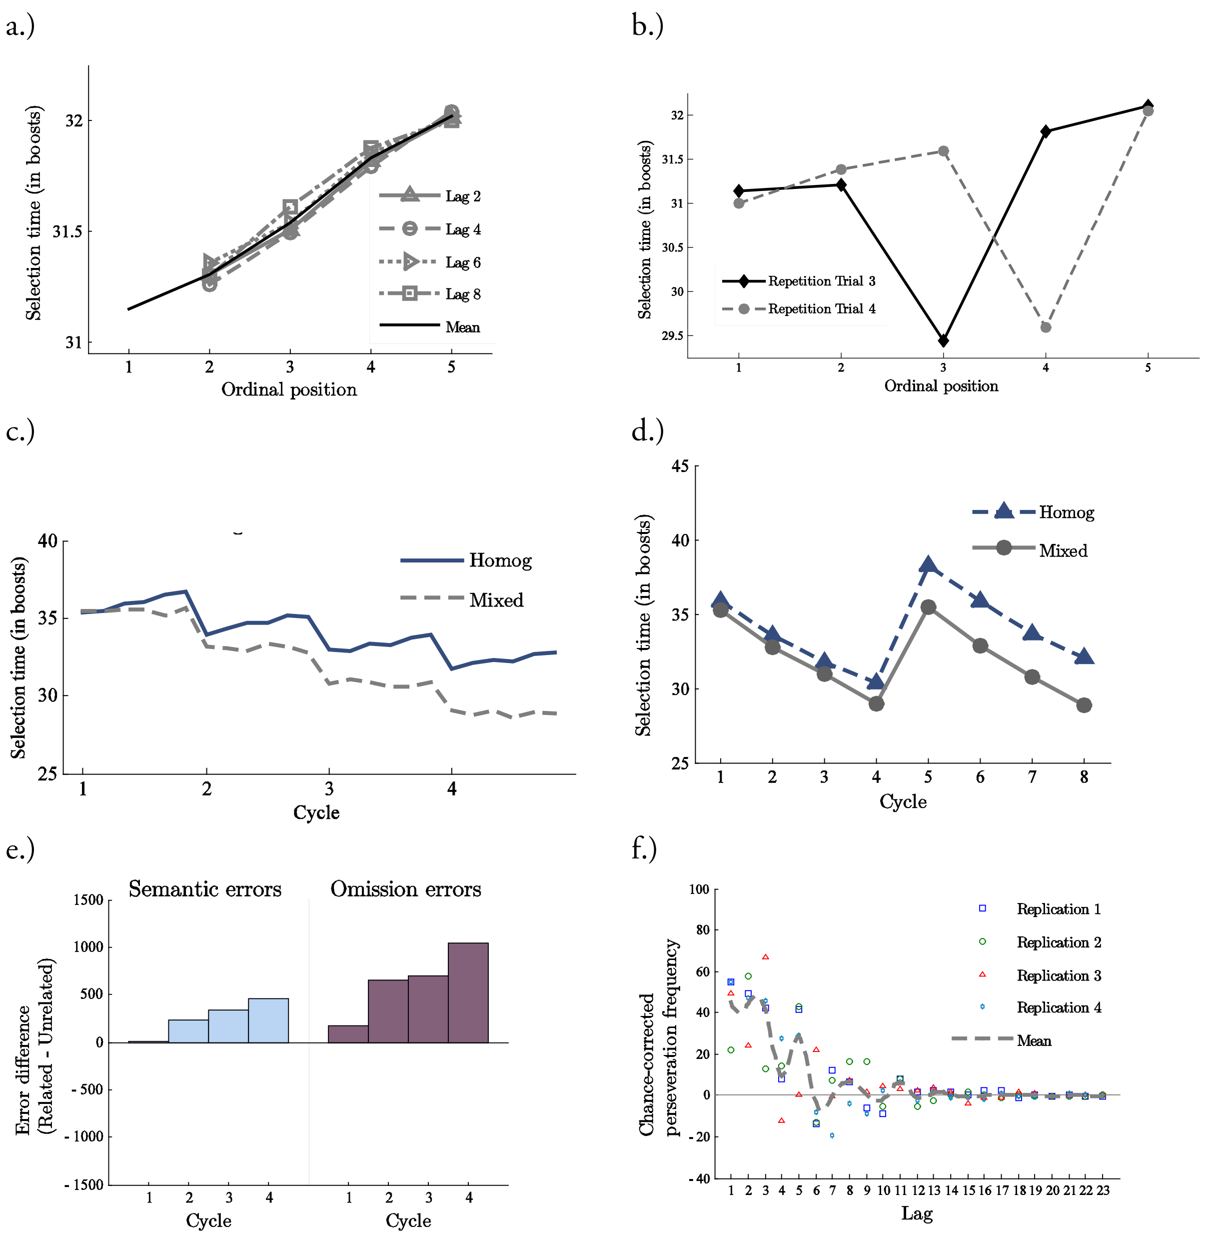

Figure 12. Replicating Simulations 1-4 with a non-competitive selection mechanism gives results identical to those already presented. a.) Lag-insensitive increases in selection times as a function of ordinal position (Simulation 1a replication). b.) Repeated items also

contribute to the interference effect (Simulation 1b replication). c.) Selection times increase within a cycle and decrease across cycles (Simulation 2 replication). d.) Accumulated interference generalizes to other same-category items (Simulation 3 replication). e.) Error effects increase across cycles (Simulation 4 replication). f.) More recently used words are more likely to be perseverated (Simulation 4 replication).

These results are important for two reasons. First, they fit our previous results to a tee. This consistency means that not only was competitive selection unimportant for our previous findings, but it did not even contribute to them. Second, demonstrating that we can get cumulative semantic interference effects without a competitive selection mechanism proves that cumulative semantic interference does not require competitive lexical selection.

So how does the weight-reduction version of the model (Figure 8c) produce cumulative semantic interference without competitive selection? We claim that the interference effect requires competition in a broader sense, not necessarily limited to the selection process. If the essence of competitive selection is to cause the facilitation of one word to slow down retrieval of other words, then our model *already does that through its learning algorithm*. The learning process naturally involves weakening connections to competitors while strengthening connections to targets. Weakening connections slows down the subsequent retrieval of competitor, thereby creating cumulative semantic interference. So, in this context, error-based learning is competitive.

Howard *et al.*'s (2006) model required competitive lexical selection because it offered a wholly facilitation-based account of cumulative semantic interference. Their priming mechanism strengthened connections to targets without weakening connections to competitors. So their model implemented something like our facilitatory learning condition. And within those confines, they needed competitive selection to convert repetition priming into a competitor-inhibition effect and achieve semantic interference (*e.g.* Mensink & Raaijmakers, 1988).

We can test this analysis. Having established our original mean-based selection rule as essentially non-competitive, we can try making it more competitive. Instead of comparing each word's activation to the mean activation all its competitors, we might hone in on just the strongest. The idea here is to amplify the effect of the strongest prepotent competitor, so that it has a greater impact on selection times. Thus, we replace Equation 6 with Equation 10, below.

Equation 10 $$t_{selection} = \log_{\beta}\left(\frac{\tau}{a_{i\,t_1} - a_{strongest\ competitor\ t_1}}\right)$$

Now we repeat the previous batch of simulations. If this tweak of the competitive selection rule is sufficient to turn repetition priming into cumulative semantic interference, then we should see it in the facilitatory learning condition.

And we do (Figure 13). In the facilitatory learning condition, focusing competitive selection on a single close competitor translates the repetition priming into cumulative semantic interference (Figure 13b). Interestingly, the selection time pattern in this facilitatory learning condition resembles that of the simulation including both weight increases and decreases (Figure 13a). Thus, competitive selection allows facilitatory changes (as in Howard *et al.*'s model) to account for both repetition priming and cumulative semantic interference effects. We should note, however, that inhibitory learning (Figure 13c) still plays a major role in creating the combined interference effect, demonstrating its continued relevance to explanations of cumulative semantic interference.

a.)

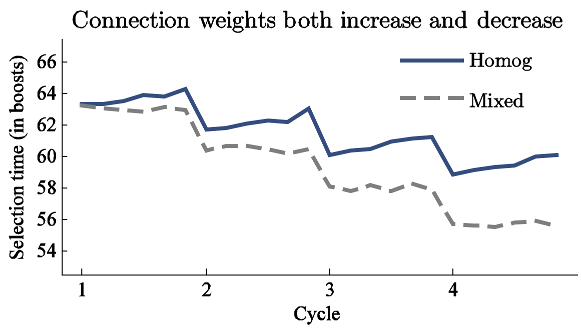

b.)

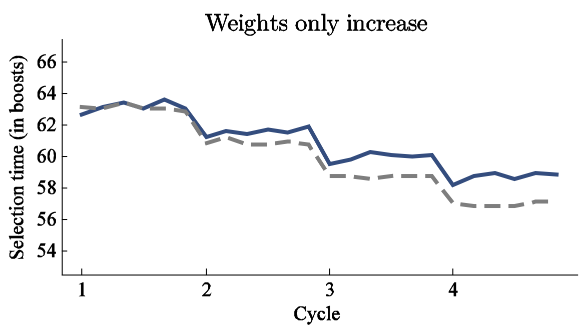

c.)

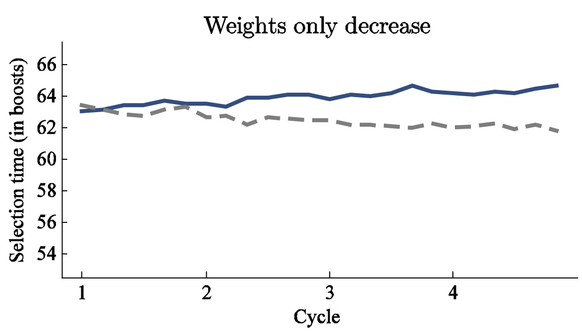

Figure 13. A more competitive selection rule (Equation 10) compares each word's activation with just its strongest competitor. a.) Applying both weight increases and decreases generates

selection latencies resembling the saw-tooth pattern from Simulation 2. b.) Increasing connection weights to target words produces a less dramatic version of this saw-tooth. c.) Decreasing weights to competitors continually increases selection times in the homogeneous condition while decreasing times in the mixed condition.

The results of Simulations 5 and 6 suggest that cumulative semantic interference requires competition, but not necessarily during lexical selection. Incremental learning can also be competitive, and competitive *learning* is sufficient to produce cumulative semantic interference without competitive *selection*. Therefore, the phenomenon of cumulative semantic interference cannot be claimed to uniquely support the existence of a competitive mechanism for lexical selection in speech production. We expand on this important conclusion in the general discussion.

CHAPTER 5: GENERAL DISCUSSION OF SIMULATIONS

Summary and implications of findings

Lexical retrieval leads to lexical learning. The light side of learning is well known. Retrieving the same word again becomes faster and more accurate. But learning also has a dark, competitive, side that hinders the subsequent retrieval of semantically related words. In our theoretical framework, this dark side of learning leads to the behaviors identified with cumulative semantic interference.

Remarkably, this framework does not require competitive lexical selection. Competitive learning obviates the need for competition in the selection process. Thus, we can modify Howard *et al.*'s (2006) explanation to identify four properties that must be true of lexical retrieval in order to account for cumulative semantic interference: shared activation, activation-dependent selection time, persistent priming, and competition. First, activating one word must also activate similar words. Though, as Howard *et al.* (2006) pointed out, this *shared activation* can be accomplished in many ways, it is an inherent property of models such as ours that rely on distributed semantic representations. Second, more activated words should be selected more quickly than less activated words. Such *activation-dependent selection time* is common to most theories of lexical selection, both competitive (*e.g.* Roelofs, 1992) and non-competitive (*e.g.* Mahon *et al.*, 2007). Third, retrieving a word must persistently prime its future retrieval. Such *persistent priming* is most readily understood as incremental learning (*e.g.* Damian & Als, 2005). Finally, facilitating retrieval of one word must have negative consequences for related words. This *competition* could be implemented in many ways, including competitive lexical selection, but we note that it follows naturally from implementing priming as error-based learning.

Given these principles, our learning model offers a parsimonious account of the empirical aspects of cumulative semantic interference. It incorporates distributed semantic representations, error-based learning, and a booster mechanism that produces activation-dependent lexical selection times and errors.

Response time effects

The model's selection times reflect all the response-time hallmarks of cumulative semantic interference. Lag-invariant incremental increases in response times, in the continuous paradigm (Simulation 1a), showed that the model can account for Brown's (1981) and Howard *et al.*'s (2006) main findings. Moreover, the model simulated the dual effects of repeating an item (Navarrete et al., 2008): that item is named much more quickly (repetition priming), but repeating it creates additional semantic interference, just about as much interference as two different related items would (Simulation 1b). Applied to blocked-cyclic naming (Simulation 2), the model's selection times showed both repetition priming and cumulative semantic interference effects (*e.g.* Damian *et al.*, 2001; Schnur *et al.*, 2006, Experiment 1). And the accumulated interference transferred to novel items from the same semantic category (Simulation 3), recalling Belke *et al.*'s (2005) report. Thus, we have demonstrated that the model accounts for the major reaction time manifestations of cumulative semantic interference.

Error effects

The learning model can also account for aphasic patients' error patterns (Simulation 4). One way of understanding aphasic brain damage is to assume that processes work normally, but are more prone to error. The model was made more error-prone by adding a small amount of noise to each word's net input. This noise led to blocking effects for semantic errors and omissions that increased with each cycle, reminiscent of findings from Schnur *et al.*'s (2006, Experiment 2) patient work. And the model's perseveration errors recalled Hsiao *et al.*'s (2009) lag-based recency gradient, suggesting a further match to the empirical patient data.

Although we are heartened by the degree to which the model simulates the major response time and error data effects in the literature, we acknowledge the model's limitations. We do not, indeed we cannot, 'fit' the data quantitatively. This is because the model's lexical base is small and its treatment of semantics is rudimentary. Most humans, for example, know more than 36 words, and would describe a dog as something more than merely a terrestrial mammal. Moreover, we have constrained the scope of the model to lexical activation and selection. We have not represented any processes leading up to the activation of semantic features, nor any processes that follow lexical selection. And, although we postulate that the experimental paradigms that are modeled may induce strategies, we do not simulate these. All of these factors compromise the model's ability to fit the actual numbers. What is not compromised by the model's simplifications, we would argue, is its ability to make theoretical issues more transparent. This was a strength of Howard et al.'s (2006) model and we hope to achieve the same with our model, particularly with regard to the issue of the role of competition, to which we turn next.

Competition and the occlusion versus inhibition debate

We provided an analysis of the model's account of cumulative semantic interference in Simulations 5 and 6. In contrast to several previous accounts, facilitatory processes contributed almost nothing to the model's interference effect (Simulation 5). Learning-based weakening of semantic-to-lexical connections, however, reliably produced cumulative semantic interference effects (Simulation 5), even without competitive lexical selection (

Simulation 6). Thus, competitive selection was not needed for cumulative semantic interference when weight decreases occurred during learning. In fact, Simulation 6 demonstrated that, although the selection mechanism employed in the first five simulations was functionally non-competitive, the model still exhibited cumulative semantic interference.

Competitive lexical selection did, however, allow the model to generate cumulative semantic interference via weight strengthening alone (

Simulation 6), as in Howard *et al.* (2006), by effectively converting an occlusion process into an inhibitory one (*cf.* Mensink & Raaijmakers, 1988). Thus cumulative semantic interference can derive from either *competitive selection*, where strong competitors interfere with target retrieval, or *competitive learning*, where strengthening a target involves weakening competitors. Therefore, we can derive a more general principle that lexical competition affects lexical selection, without constraining the point at which this competition comes into play. And error-based learning provides sufficient competition to explain cumulative semantic interference.

Our conclusion that competitive selection is not required to explain cumulative semantic interference has ramifications for the current debate about the necessity of such a selection process. It is fair to say that a majority of production researchers hold that lexical selection is competitive; that is, activated competitors retard target retrieval (*e.g.* Levelt *et al.*, 1999). Competitive selection was thought to have been demonstrated in the picture-word interference paradigm, in which a seen or heard distractor item slows the naming of a picture presented at about the same time (Schriefers *et al.*, 1990). Mahon *et al.* (2007), however, presented findings suggesting that the influence of external distractors on naming response times occurs post-lexically, and so the relevance of semantic interference from this paradigm for competitive lexical selection can be questioned (see, *e.g.* Abdel Rahman & Melinger, 2009; Janssen, Schirm, Mahon, & Caramazza, 2008; Mahon & Caramazza, 2009, for recent discussion). Thus, cumulative semantic interference, in which interference is generated from previous naming trials rather than external stimuli, might be seen as better evidence for competitive lexical selection (Dell, Oppenheim, & Kittredge, 2008; Howard *et al.*, 2006; Navarrete, *et al.*, 2008; Schnur *et al.*, 2006). Here is where our modeling exercise matters. Although the model implicates a role for competition in explaining semantic interference, it does not require a competitive selection process to simulate the data. On the one hand, the model reinforces the conclusion of Howard *et al.*, (2006), that some kind of competition is required to explain cumulative semantic interference. On the other hand, the model offers the novel hypothesis that this competition arises through learning, rather than through a lexical selection mechanism that is slowed by any activated competitor. If this hypothesis is true (as well as concerns about the relevance of the picture-word paradigm), then we are back at square one on the question of competitive lexical selection.

Incremental learning as an account of cumulative semantic interference

The model uses the delta rule, an error-based learning algorithm, to explain cumulative semantic interference. The important aspect of this algorithm is that it creates both strengthening of the connections to the target, and the weakening of the connections to competitors. Connection strengthening is clearly required for repetition priming, and connection weakening appears to explain at least a component of semantic interference and the strength of the perseveratory lag effect. Moreover, connection weakening is necessary for semantic interference if it is assumed that lexical selection is not competitive. Thus, the delta rule motivates an account of the data that combines the two principal hypothesized explanations for semantic interference and retrieval-induced forgetting in general—the occlusion and the inhibition hypotheses. The fact that the weakening and strengthening is directly proportional to error, is not, as far as we can tell, directly relevant to the model's account of the data. But this aspect of the delta rule is motivated by research on Pavlovian conditioning (*e.g.* Rescorla & Wagner, 1972), episodic memory encoding (*e.g.* McClelland, McNaughton, & O'Reilly, 1995), frequency sensitivity in priming (*e.g.* Chang et al., 2006), greater repetition priming in more error-prone conditions (*e.g.* Anderson, 2008), and division of labor effects in multi-component computational systems (*e.g.* Harm & Seidenberg, 2004).

Extending and testing the model

Though not formally simulated, the model is also compatible with several other empirically-established effects of cumulative semantic interference on lexical retrieval times:

1. **Interference is robust to timing manipulations** (*e.g.* intervening non-verbal fillers: Damian & Als, 2005, Experiment 1; RSI manipulations: Hsiao *et al.*, 2009; Schnur *et al.*, 2006, Experiments 1 and 2; *cf.* simultaneous presentation: Belke *et al.*, 2005, Experiment 2). In our model, cumulative semantic interference derives from incremental learning and therefore follows the same time course as the learning process itself. That is, it persists without regard to time.
2. **Interference is robust to filler material in other paradigms** (*e.g.* naming pictures from other categories in the blocked-cyclic paradigm: Damian & Als, 2005, Experiments 2-4). In line with Simulation 1, this robustness comes from the fact that only relevant experience leads to relevant learning, and hence priming or interference. As long as filler material is sufficiently orthogonal to the critical items, it should never affect the build-up or resolution of semantic interference.
3. **Interference effects are graded as a function of semantic similarity** (Vigliocco *et al.*, 2002). In other words, more similar sets of items produce more interference than less similar sets. Such graded effects naturally emerge from the use of distributed semantic

representations, where similarity is a function of feature overlap instead of discrete category membership.

4. **Interference is task-dependent** (*e.g.* Damian *et al.*, 2001). It requires mapping from shared semantic representations to separate lexical representations. Therefore, tasks that engage either type of representation, but involve no such mapping, should not elicit interference. For instance, non-verbally categorizing pictures according to visual or semantic features (*e.g.* Damian *et al.*, 2001) should not involve the semantic-to-lexical mapping, and should therefore not show our cumulative semantic interference. Similarly, orthographically cued word naming (*e.g.* Damian *et al.*, 2001, Experiment 2; Kroll & Stewart, 1994, Experiment 1) should only elicit semantic interference to the extent that utterance planning requires semantic access.
5. **Interference is cue-independent.** For example, Wheeldon and Monsell, (1994) used naming-to-definition to prime picture-naming, suggesting that the prime affected mappings from amodal semantic representations to lexical items. Our model is consistent with this cue-independence because semantic interference effects are carried in the semantic-to-lexical connections. Any process that uses these connections should therefore show cumulative semantic interference, regardless of the instigating stimulus.

These five response time effects derive from the fact that the model attributes semantic interference to incremental learning during lexical access, as opposed to some kind of time-dependent facilitory or inhibitory priming. Indeed, these findings could also be readily explained by the model of Howard et al. (2006), which was the first implemented account of semantic interference in production based on persistent changes. We consider our model to be a descendant of that model. However, our model differs from its ancestor in four respects. First, it links up with speech-error based models of production and hence accounts for errors, including perseverations and omissions. This allows the model to simulate aphasic error data, and to offer a mechanism for how the brain chooses among active lexical candidates. Second, it ascribes an important role to connection weakening in explaining semantic interference effects on RT's and errors, and specifically the perseveration lag function. Connection weakening is a natural consequence of the delta rule, an algorithm for incremental learning. Third, the model represents lexical retrieval for speech production as a controlled process. The booster function links selection in the model's representation of a picture naming task to syntax-driven selection in models of sentence production, and suggests a means of representing other executive constraints on retrieval. And fourth, the model represents the possibility that the competition needed to explain interference may not occur during lexical selection; it can arise from learning. As a result, non-competitive lexical selection becomes a viable account for data in this paradigm, and for production in general.

Additional predictions of our model

Our model offers additional predictions, mostly stemming from the idea that cumulative semantic interference reflects incremental learning. First, because the model's learning is error-based and not dependent on actual selection, cumulative semantic interference should accrue even when a name is not correctly retrieved, as in the case of errors and omissions. Although this property has not been tested for cumulative semantic interference, it appears to hold true for retrieval-induced forgetting. Using a part-set cueing paradigm, Storm, Bjork, Bjork and Nestojko (2006) demonstrated that providing misleading practice cues (*i.e.* they either elicited novel associations or offered no valid response) created retrieval-induced forgetting equivalent to that from valid cues. Therefore we expect, and our model predicts, that semantic interference should even accrue from naming trials that elicit omission errors.

Also, because cumulative semantic interference is attributed to persistent changes in connection weights, the effects should not spontaneously dissipate. Several picture-naming studies (*e.g.* Howard *et al.*, 2006; Nickels, Howard, Dodd, & Coltheart, 2008) have suggested that cumulative semantic interference persists at relatively long item-lags (minutes). But longer periods of persistence have not been examined in the picture-naming paradigm. One episodic memory study (Anderson & Spellman, 1995) concluded that retrieval-induced forgetting persisted over lags of twenty minutes and another (Storm *et al.*, 2006) even reported that retrieval-induced forgetting effects may remain detectable after a 1-week lag (but see MacLeod & Macrae, 2001, and Postman *et al.*, 1968, for some evidence to the contrary). So, with the caveat that relevant experience is easier to define in a model than in the real world, we should find that cumulative semantic interference dissipates largely as a function of relevant experience, rather than as a function of time.

Finally, since our model derives from domain-general principles, it should formally extend to non-linguistic processes. Any system that involves shared activation, activation-dependent selection, and competitive learning should exhibit similar effects, and may be examined through similar experimental paradigms. Throughout this paper, we have referred to retrieval-induced forgetting (RIF), a well-known episodic memory effect, where the process of retrieving one association leads to impaired recall of competing associations. Several prominent accounts of RIF (*e.g.* Anderson, 2003; Norman *et al.*, 2007) suggest that it arises from the process of new information overwriting the old. In other words, they ascribe the memory effect to the process of learning. Such explanations have also surfaced for effects in visual object recognition (*e.g.* Marsolek, Schnyer, Deason, Ritchey, & Verfaellie, 2006). Thus, we can identify cumulative semantic interference with a more general theme in the way the mind operates, independent of the particular types of representations in use.

It may also be possible to extend the model's links with brain mechanisms and regions. The successful resolution of semantic interference may depend on processes subserved by the LIFG and/or left temporal lobe (LT) and thus may be impaired by damage to these regions. Evidence for this localization comes from both neuroimaging studies of healthy subjects (Maess, Friederici, Damian, Meyer, & Levelt, 2002; Moss *et al.*, 2005, Schnur *et al.*, 2009; see also Hocking, McMahon, & de Zubicaray, 2008) and lesion-mapping studies in patients (Schnur *et al.*, 2005; 2009). The involvement of frontal processes is further supported by single-case and group studies that associate exaggerated blocking effects with nonfluency and other symptoms of anterior aphasia (Biegler, Crowther & Martin, 2008; McCarthy & Kartsounis, 2000; Schnur *et al.*, 2006; Wilshire & McCarthy, 2002). Consistent with the LT localization, Simulation 4 showed that adding noise to lexical activations succeeded in reproducing the aphasic blocking pattern. In a follow-up test (not reported in this article), we simulated LIFG damage by instead adding noise to the booster mechanism and found that this gave the same result as the simulated lexical damage. While these simulations of brain damage are encouraging, we note that they fail to capture some subtle but potentially important distinctions. For one thing, it seems that while lesions to either the LIFG or LT exaggerate the blocking effect, only LIFG lesions produce the pattern of increased errors across cycles (Schnur *et al.*, 2005; 2006); in our simulations, noisy lexical activations and a noisy booster both had this effect. Also, whereas our simulations of aphasic damage focused on errors (in keeping with the data from Schnur *et al.*, 2006), recent evidence suggests that some frontal aphasics manifest exaggerated blocking interference in naming latencies rather than errors (Biegler *et al.*, 2008). To resolve these issues, it may be useful to measure the time course of LIFG and LT activation with imaging methods with good temporal resolution (e.g. the EROS optical imaging technique, Tse *et al.*, 2007). If interference in, say, the blocking paradigm is present in lexical-semantic areas (LT), but resolved in the LIFG, then manipulations of degree of interference should manifest first in the former region, and then the latter.

Conclusion

The model instantiates a dynamic view of lexical knowledge. Shared semantic representations put competing words in a dynamic equilibrium where no semantic feature connects too strongly to any one word. Each act of lexical retrieval produces persistent, competitive, learning that perturbs this balance. It facilitates repeating the same word and impairs access to competing words. But retrieving a competitor shifts the balance back again. So not only are we capable of learning new words every day, but we are constantly adjusting which words, of the ones we know, are more or less available for use in speaking.

We believe that our model offers a parsimonious account of the lexical production data dealing with persistent effects, and it has implications for other issues in production, learning, and memory. Similarity-based interference that results from learning is not unique to lexical retrieval (*e.g.* Anderson, 2003; Marsolek *et al.*, 2006; Norman *et al.*, 2007). But, when words are retrieved in a semantically-manipulated context, the resulting impairment is what we know as cumulative semantic interference.

CHAPTER 6: EXPERIMENTS WITH HUMANS

Introduction

Cumulative semantic interference refers to the phenomenon that retrieving multiple words from a semantic category makes it increasingly difficult to retrieve other semantically related words. For instance, when naming a series of pictures, like HAMSTER, GERBIL, MOUSE, naming the second picture will take longer than naming the first, and naming the third will take longer still (*the incremental interference effect*; e.g. Brown, 1981; Howard, Nickels, Coltheart, & Cole-Virtue, 2006; Navarrete, Mahon, & Caramazza, 2010; Nickels, Howard, Dodd, & Coltheart, 2008), and the increase in naming latencies may be accompanied by increases in error rates (Navarrete et al., 2010). When repeatedly naming a small set of pictures, a paradigm called blocked-cyclic naming, semantic interference manifests as an attenuation of repetition priming in blocks where the pictures are semantically related, compared to blocks where they are not (*the semantic blocking effect*). Again, this interference increases naming latencies and semantic error and omission rates for both healthy (Belke, 2008; Belke & Meyer, 2007; Belke, Brysbaert, Meyer, & Ghyselinck, 2005; Belke, Meyer, & Damian, 2005; Damian & Als, 2005; Damian, Vigliocco, & Levelt, 2001; Schnur, Schwartz, Brecher, & Hodgson, 2006; Vigliocco, Vinson, Damian, & Levelt, 2002) and impaired speakers (Biegler, Crowther, & R. C. Martin, 2008; Hsiao, Schwartz, Schnur, & Dell, 2009; Jefferies, Baker, Doran, & Ralph, 2007; Schnur et al., 2006, 2009; Wilshire & McCarthy, 2002), and the accumulated interference generalizes to new words from the same semantic category (Belke et al., 2005). Neural signatures of cumulative semantic interference include increased activation in the left inferior frontal gyrus and left anterior temporal lobe, and patient work has associated stronger interference with damage to these areas (Schnur et al., 2009). Computational accounts of cumulative semantic interference generally agree that it reflects some form of competition between lexical entries (e.g. Howard et al., 2006; Oppenheim, Dell, & Schwartz, 2010). So cumulative semantic interference is empirically robust and stands to become a theoretically and clinically useful tool for investigating lexical access during semantically driven speech production.

Yet one crucial question about the phenomenon remains: how long does cumulative semantic interference last? A short-lived effect could suggest a by-product of normal speech production (e.g. spreading activation), offering a glimpse of the means by which words are retrieved, while a more persistent effect might better reflect changes to the way that speakers structure their vocabularies.

Earlier inquiries aimed only at establishing if interference might be better characterized as residual activation or as something more persistent (e.g. Damian & Als, 2005; Howard et al., 2006; Hsiao et al., 2009; Schnur et al., 2006). Residual activation could create semantic interference if within-trial changes in activation were able to affect retrieval in subsequent trials. For instance, nodes representing the word HAMSTER might have resting activation of 0 that increases to 1 when it is retrieved. If HAMSTER remains active for some time after being retrieved – perhaps its activation decays exponentially towards zero as a function of time – then its activity could disrupt attempts to retrieve GERBIL (e.g. via competitive lexical selection mechanisms). While this account has some appeal, the serial ordering of words in fluent speech requires that such residual changes in activation should rapidly decay following target selection (Bock & Griffin, 2000; Dell, 1986), meaning that activation-based effects should be bound to millisecond timescales.

Several studies have now demonstrated that cumulative semantic interference is not actually characterized by such rapid decay, thus ruling out the residual activation explanation. For instance, the magnitude of cumulative semantic interference for naming latencies remains essentially the same whether pictures are named 1,000 or 12,000 milliseconds apart (Damian & Als, 2005; Schnur et al., 2006, Experiment 1), with similar findings for error patterns in aphasic data (Schnur et al., 2006, Experiment 2; Hsiao et al., 2009). The interference is also resilient to manipulation of item-lags: a related competitor hinders target retrieval just as much if it comes two items earlier or 32 (Howard et al., 2006; Nickels et al., 2008). So these findings show that interference persists much longer than would be expected for phenomenon based on residual activation, and does not require continued focus on a particular task or semantic category, leading some researchers (e.g. Damian & Als, 2005; Howard et al., 2006; Oppenheim et al., 2010) to characterize CSI as reflecting more permanent changes to the way that speakers retrieve words, such as incremental learning.

In this context, *incremental learning* refers to small, experience-driven changes in the mapping from semantic concepts to words, and would thus link cumulative semantic interference – a laboratory phenomenon – to the real-world processes by which people incrementally acquire and lose their vocabulary. As an explanation of cumulative semantic interference, the incremental learning account assumes that entries in a speaker's lexicon are maintained in a dynamic equilibrium where each experience makes one word persistently more accessible and its competitors less so, but the next experience may partially reverse that change. This account has been most fully explored in Oppenheim, Dell, and Schwartz's (2010) *Dark Side* model of lexical access, which demonstrated that adding incremental leaning to a basic connectionist framework, without incorporating short-lived changes, was sufficient to account for both error and response time manifestations of CSI for both intact and impaired speakers. In that model, words are activated by shared semantic features, and then repeatedly boosted until a clear winner emerges. After each retrieval, an error-driven learning algorithm adjusts the mapping from meanings to words. Strengthening connections to the previous target produces repetition priming, and weakening connections to previous competitors produces semantic interference. For instance, naming a mammal as HAMSTER would strengthen a connection from the concept MAMMAL to the word HAMSTER and weaken the connection from MAMMAL to GERBIL, thus making HAMSTER more accessible and GERBIL less so, but retrieving GERBIL in the next trial would do the opposite.

A key property of incremental learning accounts is the time course that they assume for cumulative semantic interference. While a learning mechanism does not, in principle, require that weight changes (and hence interference) remain as strong after one-second as after one-hour (e.g. McClelland & Rumelhart, 1985, assumes time-based weight decay, where a weight change comes on strong before falling back to a smaller persistent level), it does imply that learning, and hence its behavioral consequences, should remain in effect until overwritten. Thus an incremental learning mechanism implies that cumulative semantic interference should persist over lags much longer than the seconds and minutes that researchers have previously investigated. Such true persistence may be problematic on both empirical and theoretical grounds. Empirically, researchers have previously failed to demonstrate any effect of semantic interference at lags longer than a few minutes (e.g. Wheeldon & Monsell, 1994), which could suggest that cumulative semantic interference may not actually fit the expected time course for incremental learning any more than it fits that for residual activation. Fully persistent changes would also imply that semantic blocking effects should (under typical counterbalancing strategies) reverse over the course of a blocked-cyclic naming experiment, yielding semantic interference near the start of the experiment and facilitation near the end, but no such reversal has yet been reported. Theoretically, fully persistent changes raise the question of how speakers might overcome persistent interference in order to produce fluent and varied speech outside the laboratory. That is, if semantic interference were truly permanent, accumulating with 30ms for each competitor (Howard et al., 2006), then the interference that a speaker accumulates over their lifetime should quickly diminish any possibility of speaking fluently on a variety of related topics. Although repetition priming could offset this effect (e.g. Oppenheim et al., 2010), it obviously requires the opportunity for repetition. Thus, some have argued that CSI must decay, perhaps lasting for less time than it takes to boil an egg (i.e. entirely disappearing within 80s-4 minutes, e.g. Damian & Als, 2005; Wheeldon & Monsell, 1994).

There is, however, reason to suspect that cumulative semantic interference should last longer. For one thing, cumulative semantic interference seems to have much in common with a human memory phenomenon known as retrieval-induced forgetting (Oppenheim et al., 2010). Retrieval-induced forgetting (RIF; e.g. Anderson, Bjork, & Bjork, 1994) has been shown to affect episodic recall at least 20 minutes later (Anderson & Spellman, 1995), and by one report, may remain evident even a full week after encoding (Storm et al., 2006). However, other studies have failed to detect such interference at longer lags (MacLeod & Macrae, 2001; Postman et al., 1968), casting some doubt on these claims. Moreover, the persistence of RIF has only been examined for episodic memories of explicit cue-word associations (e.g. FRUIT:A___ --> APRICOT), so it is not clear that this property should extend to more naturalistic meaning-based word retrieval.

This chapter presents several picture-naming experiments that consider the persistence of cumulative semantic interference effects at much longer lags than ever before, thus addressing an important implication of incremental learning accounts of cumulative semantic interference. Specifically, they ask whether the semantic interference that accumulates when naming a small set of items (e.g. HAMSTER, GERBIL, MOUSE) generalizes to new items from the same category (e.g. SQUIRREL, GUINEA PIG, RAT) approximately an hour later. This one-hour (3,600,000ms) delay is approximately 130 times longer than the longest delay between category exemplars in any published report of cumulative semantic interference (27,250ms in Howard et al.'s, 2006, 8-item lag condition), 34 times longer than in any unpublished report that I know of (105,250ms in Nickels et al.'s, 2008, 32-item lag condition), and 15-45 times longer than current estimates for the maximum lifespan of the effect (80,000-240,000ms, according to Damian & Als, 2005; Wheeldon & Monsell, 1994).

In these experiments, persistent consequences of semantic interference should manifest as longer naming latencies in the second session for pictures whose semantic competitors were named in the first. The first two experiments use variants of a blocked-cyclic naming procedure to ask whether the semantic blocking effect – that is, the contrast between naming latencies for Homogeneous- versus Heterogeneous-context pictures – in the second session exceeds that of the first. Both experiments do suggest some carryover, but the trends prove difficult to statistically confirm. Experiment 3 takes a more direct approach to examine the persistence of semantic interference and its possible decay. Together, these experiments clearly demonstrate that interference persists over at least a one-hour delay, providing reasonably strong support for the notion that the cumulative semantic interference seen in laboratory studies reflects the persistent restructuring of semantic-to-lexical mappings.

A secondary goal of these experiments is to consider a hypothesized continuity between the incremental interference effect and the semantic blocking effect. Briefly, the incremental interference effect typically presents in non-blocked picture naming studies as an increase in picture naming latencies for each semantically related picture that has been named before (Brown, 1981; Howard et al., 2006; Navarrete et al., 2010). Crucially, that increase is always evident the first time that a picture is named (i.e. in the first, and often only, cycle). This pattern appears inconsistent with the canonical development of the semantic blocking effect in blocked-cyclic naming (as described by Belke, Meyer, & Damian, 2005): the blocking effect is typically absent or even reversed in the first cycle of naming, then emerges in the second cycle and remains relatively stable thereafter. I attempt to reconcile these apparently inconsistent manifestations of cumulative semantic interference by considering analogues of each manifestation within both experimental paradigms.

Experiments 1 and 2 demonstrate that an analogue of the incremental interference effect is actually evident in the first cycle of these blocked-cyclic naming experiments. Experiment 3 uses concept-similarity ratings to demonstrate an analogue of the semantic blocking effect in a non-blocked paradigm, and further show for the first time that the slope of the incremental interference effect is sensitive to gradations of semantic similarity.

Experiment 1

The incremental learning account predicts that cumulative semantic interference should continue to impair word retrieval long after the original interference-generating episode has passed. There is a challenge in figuring how to test this hypothesis, though. Practically, we want a design that creates a lot of interference, so it seems appropriate to use a cyclic naming design, assuming that each trial will contribute at least a bit more interference, thus increasing our opportunity to detect a persistent effect. Following this idea, Experiment 1 examines whether repeatedly naming a set of three pictures from a single semantic category, in a blocked-cyclic picture naming paradigm, might produce an exaggerated semantic blocking effect when naming three new pictures from the same category an hour later.

This experiment adapts a design that Belke et al. (2005, Experiment 3) originally used to demonstrate that the semantic blocking effect *immediately* generalizes to new exemplars from the same semantic category. In that blocked-cyclic naming experiment, participants repeatedly named four pictures representing either a single semantic category (homogeneous context; e.g. DOG, COW, MOUSE, SHEEP) or four different semantic categories (heterogeneous context; e.g. DOG, TOASTER, ARM, BALL). After cycling through this set four times, the pictures were replaced; in the homogeneous context, the new pictures represented the same semantic category as the previous set (e.g. CAT, HORSE, PIG, GOAT), and in the heterogeneous context they represented four new semantic categories (e.g. TREE, SQUARE, HAT, NUN). Crucially, the blocking effect that developed when naming the first set of homogeneous pictures (DOG, COW, MOUSE, SHEEP) appeared to carry over to the second (CAT, HORSE, PIG, GOAT; Figure 14a-b).[7] Oppenheim et al. (2010) showed that such a pattern would naturally arise if each act of word retrieval were also an act of word learning (Figure 14c).

---

[7] Belke et al. statistically demonstrated this carryover by comparing naming latencies in the experiment described above (the 'switch' condition) to those in an experiment where participants named the same four pictures through all eight cycles (the 'no-switch' condition). If semantic interference *did not* generalize to new exemplars, then it would have to re-accumulate after the switch, yielding a weaker interference in the switch condition. Their analyses showed significant main effects of semantic context and switch, but no significant interaction between the two (as would be expected if the interference did not generalize). My experiments, however, seek to avoid

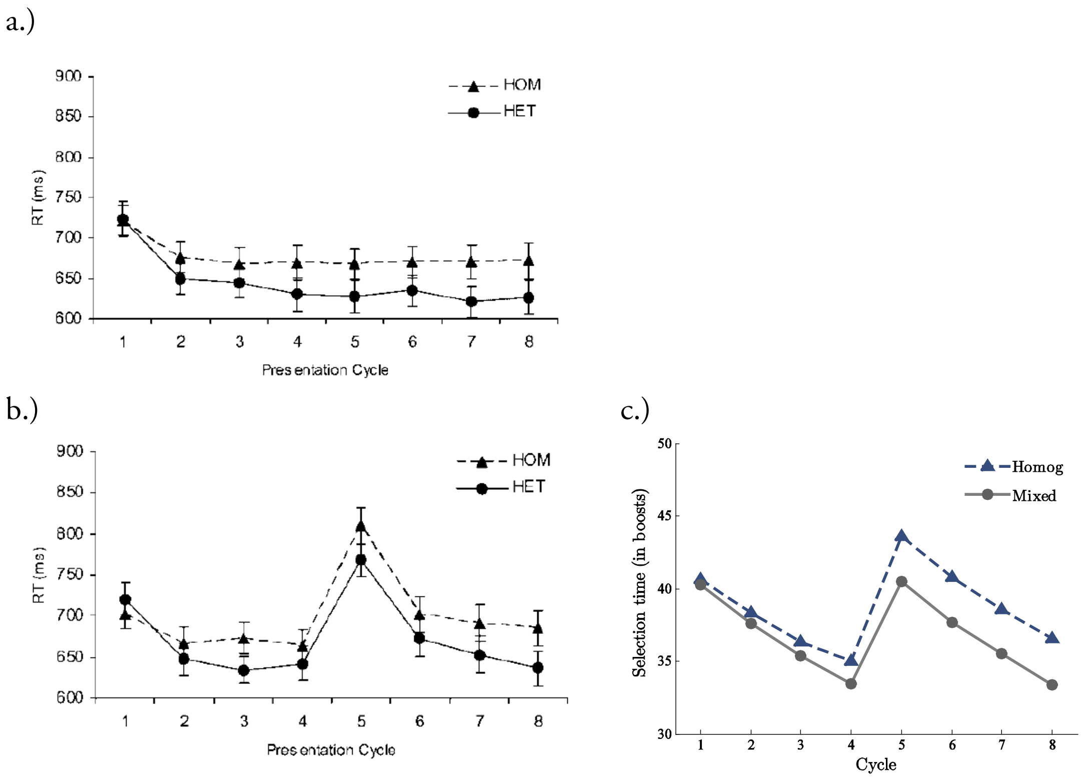

Figure 14. Results from Belke et al.'s (2005) Experiment 3, which showed that the semantic blocking effect generalized to new items from the same category. a.) Naming latencies in a condition where the same set of four items was named for eight cycles. b.) Naming latencies in a condition where the stimuli were replaced by new items (from the same semantic category in the Homogeneous condition; different categories in the Heterogeneous condition) after the fourth cycle. c.) The *Dark Side* simulation of (b), from (Oppenheim, Dell, & Schwartz, 2010).

Experiment 1 presents a structural twist on Belke et al.'s experiment, by presenting the first and second set from each category in two separate sessions, one hour apart (Figure 15). In each session, participants repeatedly named small sets of pictures in the context of either pictures from the same semantic category (homogeneous context) or pictures from different semantic categories (heterogeneous context). In the second session, homogeneous-context items always represented novel items from the same semantic categories as the homogeneous-context items in the first session, while heterogeneous-context items represented categories that had not appeared in the first session. As in other blocked-cyclic naming experiments, cumulative semantic interference is primarily measured here as the difference in naming latencies for homogeneous versus heterogeneous pictures.

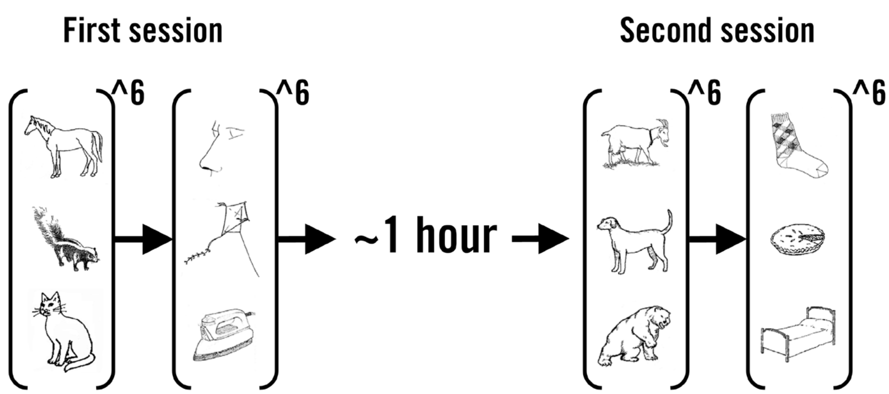

Figure 15. A simplified depiction of the structure of Experiment 1. In each block (Homogeneous or Heterogeneous), participants name a set of three pictures six times in a random order (sampling randomly without replacement). One hour later, they continue the task with new pictures; in the Homogeneous condition, these pictures represent the same semantic categories as in the first session. In the Heterogeneous condition, they represent novel categories.

Dividing the task into two sessions necessitates some changes to the experiment design. The major complication is that, if semantic interference persists indefinitely, then we have to consider semantic relations not just within blocks and sessions, but across them, thus making an adaptation of

Belke et al.'s control condition (shown in Figure 14a) infeasible.[8] Consequently, assessing the generalization of semantic interference – and now its persistence – cannot rely on the *lack* of a significant difference in the magnitude of interference between a condition where the items changed (Figure 14b) and a condition where they did not (Figure 14a), which served as Belke et al.'s (2005) major evidence. Therefore, the test of whether semantic interference persists and generalizes in Experiment 1 draws on both a prediction of the Dark Side model and a trend in Belke et al.'s data to establish a new criteria: since semantic interference in blocked-cyclic naming experiments is known to accumulate over multiple cycles, a *persistent* interference effect should produce a larger semantic blocking effect in the second session than in the first.

As noted earlier, this experiment also provides an opportunity to consider incremental interference effects in the context of a blocked-cyclic naming paradigm. Oppenheim et al.'s (2010) simulations suggest that incremental interference effects should underlie the semantic blocking effect, and thus predict the presence of incremental interference effects in blocked-cyclic naming

[8] This is because continuity in the Heterogeneous context of the control condition requires that any items presented in that context be semantically orthogonal to any other items that are presented in the trials intervening between their first and second session presentations. Even assuming that semantic orthogonality could be approximated by binary category membership, each block for the control condition would thus require five categories that could not be reused in any other block, thus severely limiting the amount of data that could be collected from each participant.

experiments, beginning in the very first cycle. In contrast, most explanations for the typical absence of a semantic blocking effect in the first cycle of blocked-cyclic naming (as reviewed by Damian & Als, 2005) postulate mechanisms that should attenuate incremental interference effects in that cycle. So two important questions here are whether an incremental interference effect can be obtained in the first cycle of blocked-cyclic naming, and whether its presence lines up with the emergence of the semantic blocking effect.

Methods

The experiment was divided into two sessions. In the first session, after a familiarization phase, participants repeatedly named eight blocks of three pictures for six cycles each. Half of these blocks consisted of three pictures from a single semantic category (homogeneous sets) and half consisted of three pictures from three different categories (Heterogeneous sets). One hour later, eight new blocks of pictures were named, where each Homogeneous set consisted of three novel exemplars from a semantic category that had been presented as a homogeneous set in the first session, and each heterogeneous set consisted of three exemplars from categories that had not appeared in the first session. Assignment of items to the various conditions (Homogeneous vs. Heterogeneous context, first vs. second session) was counterbalanced across subjects.

Participants

Twenty-four University of Illinois undergraduate students participated in exchange for course credit. All were native monolingual US English speakers with normal or corrected-to-normal vision and no reported history of language impairments. None had participated in any of the computer simulations or other experiments reported in this dissertation, nor had they cited Oppenheim, Dell, and Schwartz (2007; 2010) in any published work.

Materials

This experiment employed the set of 72 black-and-white line drawings previously used and described by Schnur et al. (2006; also used in Biegler, Crowther, & R. C. Martin, 2008; Hsiao, Schwartz, Schnur, & Dell, 2009; Schnur et al., 2009; Thompson-Schill, Schnur, Hirshorn, Schwartz, & Kimberg, 2007), which were originally drawn from Snodgrass & Vanderwart (1980) and other similar sources. These depicted six exemplars (e.g. DOG, CAT, HORSE, GOAT, SKUNK, BEAR) from each of 12 semantic categories (e.g. mammals, shapes, appliances).

Design

The experiment consisted of two 144-trial sessions, separated by a one hour delay. Each session included eight six-cycle blocks, with three pictures each appearing once in each cycle (for a total of 36 trials per block). Figure 16 illustrates the assignment of pictures to the blocks and conditions. Within each block, the three pictures always included either three pictures from a single semantic category (e.g. horse, skunk, dog, i.e. Homogeneous set) or three pictures from three different semantic categories (e.g. glass, top, nun, i.e. Heterogeneous set). Homogeneous and Heterogeneous blocks alternated, with their order counterbalanced across participants. In the second session, the Homogeneous blocks always presented novel exemplars from one of the semantic categories that had appeared as a Homogeneous-item category in the first session (e.g. skunk, bear, goat, to follow the horse et al. example), while the Heterogeneous blocks always presented novel exemplars from semantic categories that had not appeared in the first session (e.g. nurse, toe, flower). Thus, a participant always named items from the same categories in the Homogeneous conditions in both sessions, and never named items from the same categories in the Heterogeneous conditions in both sessions.

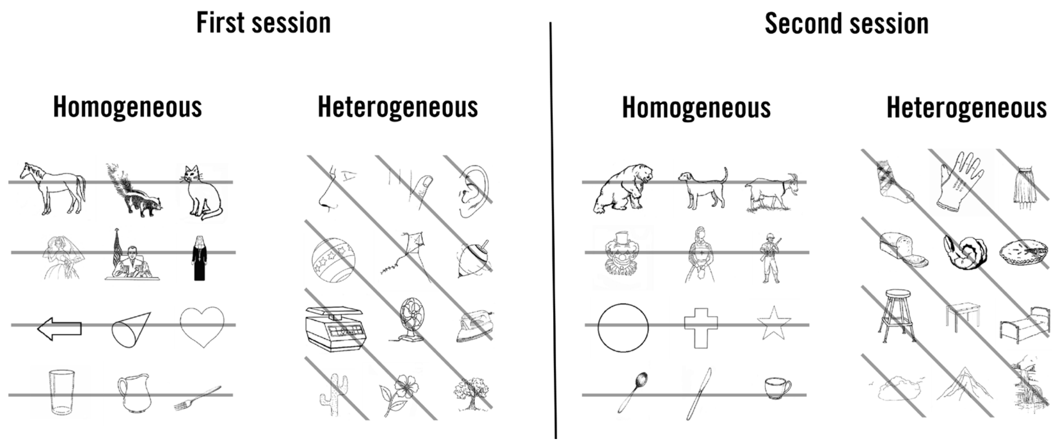

Figure 16. Illustration of the assignment of items to blocks and conditions for a single participant in Experiment 1. Grey lines connect items assigned to the same block. Homogeneous blocks group items within categories (rows); Heterogeneous blocks group items across categories (diagonals). Note that identical categories (but different items) are represented in the Homogeneous blocks of each session, but nonidentical categories are represented in the Heterogeneous blocks of the two sessions. Assignment of category sets to the various Session and Semantic context slots was counterbalanced across participants, such that each slot was occupied by each set an equal number of times across the experiment.

To populate this design, the 12 semantic categories were split into three groups of four categories each. In each run of the experiment, one group served as the source for homogeneous items in both Sessions 1 and 2, a second group served as the source for heterogeneous items in Session 1, and the remaining group served as the source for heterogeneous items in Session 2. To facilitate counterbalancing, each semantic category was further split in half (attempting to match log frequency across halves and avoid strong paired associations within a half), yielding two subgroups for each group of categories. Homogeneous sets used these subgroups directly, while heterogeneous sets were created by drawing sets of items from across categories. Assignments of category groups and sets to the Session and Semantic context slots were counterbalanced across participants, such that each slot was occupied by each set an equal number of times across the experiment.

Procedure

The experiment began with a familiarization phase, immediately followed by the first of two test sessions. During familiarization, each picture was presented once, one at a time, in the center of a 17" CRT monitor. Pictures were ordered randomly without replacement, with each sequence of 12 pictures sampling one picture from each category. Each picture was preceded by a fixation point for 500ms, then a blank screen for 500ms, and then appeared on the screen for 2000ms or until the voicekey detected a response. Participants were instructed to name each picture as it appeared. The desired name then appeared below the picture, as confirmation or correction. This familiarization procedure is similar to that used in many other picture-naming studies (e.g. Alario, Ferrand, Laganaro, et al., 2004).

Then the first test session began. Each test session consisted of eight 36-trial blocks, with each block consisting of either three homogeneous- or three heterogeneous-condition pictures.

Before each block, participants were prompted to press the spacebar when ready to proceed. They then named three pictures for six cycles. Presentation order within a cycle was random. Similar to the familiarization phase, each trial began with a fixation point for 500ms, then a blank screen for 500ms, then the picture appeared on the screen for 2000ms or until the voicekey detected a response, at which point the screen went blank for 1000ms. This sequence/timing is similar to what has been reported for most previous blocked cyclic naming experiments (Table 3).

Table 3. Summary of previous experimental designs for published experiments investigating cumulative semantic interference via picture naming. Timing for current experiments approximates those published previously. A '-' denotes cells where no information was included in the published article. AR&M=Abdel Rahman & Melinger. *The Exemplars column lists the number of unique exemplars of a category that were named in a single block or cycle. †These times are voicekey-modified such that, for example, a stimulus displays for a maximum of 2000ms, but will disappear immediately (ending the trial) if the voicekey detects a response.

			Trial timing (in ms)					RT-based exclusion criteria				
Paradigm	Language	Exemplars*	Fixation	Blank	Stimulus	ISI	Total	Lower	Upper	SD	SD basis	Comments
Blocked cyclic picture naming												
Damian et al (2001, Expt 1)	German	5	500	400	500	1500	2900	200	1500	>3	*Ss* cond. $\hat{\mu}$	
Vigliocco et al (2002)	English (UK)	8 or 4	300	450	2500†	200	3450†	250	1500	>2	*Ss* cond. $\hat{\mu}$	
Maess et al (2002)	German	5	200	600	1300†	1300	3400†	250	1300			MEG
Belke et al (2005a)	English (UK)	4	500	100	1100	800	2500	-	-	-		
Belke et al (2005b, Expts 1&3)	English (UK)	4	800	100	1100	650	2650	-	-	-		
Damian & Als (2005, Expts 1-4)	English (UK)	4	500	500	1500	1500	4000	250	1200			Interleaved
Schnur et al (2006, Expt 1 short)	English (US)	6	-	-	5000†	1000	6000†	300	2000			Older *Ss*
AR & M (2007, Expts 1&2)	German	5	500	-	2000	1000	3500	-	-	>2.5	*Ss* cond. $\hat{\mu}$	
Belke & Meyer (2007, Expt 1a)	English (UK)	4	800	100	1100	650	2650	-	-			1/2 older *Ss*
Belke (2008)	German	5	800	50	1100	970	2920	-	-	-		
Ganushchak & Schiller (2008)	Dutch	5	500-800	500	500†	var.	2000	300	1500			ERP
Schnur et al (2009, Expt 1)	English (US)	6	500	200	650	1150	2500	n/a	n/a			fMRI; no RT
AR & M (2010, Expts 1-4)	German	5	500	-	2000	1000	3500	-	-			
Other designs												
Howard et al (2006)	English (UK)	5	500	250	2000	500	3250	250	2000			Unblocked, 1-cycle
Hocking et al (2009)	English (AUS)	5	2200	-	800	-	3000	300	1500			fMRI; 1-cycle blocks
Navarrete et al (2010, Expt 1)	Italian	5	500	250	1500	700	2950	350	-	>2	*Ss* $\hat{\mu}$	Unblocked, 4-cycles
Average	-	5	650	292	1541	923	3198	272	1563			
Current studies												
Experiment 1	English (US)	3	500	500	2000†	1000	4000†	250	1500			
Experiment 2	English (US)	3	500	500	2000†	1000	4000†	250	1500			Interleaved
Experiment 3	English (US)	3	500	500	2000†	1000	4000†	-	2000	>2.5	Residuals	Semi-blocked, 6-cycles

After the first session ended, participants undertook unrelated activities for approximately one hour. They first participated in an inner speech tonguetwister experiment (now published as Oppenheim & Dell, 2010). None of the stimuli in the two experiments overlapped for any participant.[9] With any remaining time, participants were offered sudokus and given the opportunity refresh themselves (e.g. bathroom, water break), with the instruction to avoid discussing the experiments. After at least 60 minutes had passed since the end of the first block of the first session (mean time: 74:41 minutes, range: 55:50-88:42), participants began the second session.

The second test session proceeded exactly as the first, but included entirely different pictures, as described in the Design section. Until the onset of the second session, participants were given no indication that they would return to the picture naming task, as opposed to participating in a third unrelated experiment.

Equipment

Experiment presentation was controlled via E-Prime v1.1. Responses were digitally recorded via a headmounted microphone, and transcribed offline. Naming latencies were collected via an E-Prime serial response box, with erroneous naming latencies (e.g. due to lipsmacks and equipment malfunctions) later identified via visual waveform analysis of the recorded audio.

Data analysis

Participants' utterances were digitally recorded and verified offline, with naming latencies assessed via E-Prime's voicekey function. Four types of outcomes were excluded from the naming

[9] Though two pictures – a head and a mouth – did appear as cues in the tonguetwister experiment, they only served as nonverbal instructions, and never appeared in the naming experiment. The data show essentially the same patterns if the data from the affected category is excluded.

latency analyses: 1.) Utterances that did not match the expected response (e.g. dog → goat, dog → mutt, dog → go…, dog → [omission]); 2.) Utterances that were preceded by or included unexpected sounds (e.g. lipsmacks, hesitations, movement noise, laughter); 3.) Utterances that were generated too early (<250ms) or too late (>1500ms); and 4.) Equipment malfunctions. Four participants were replaced because these exclusions left less than 80% of their responses for the naming latency analysis.

Speech errors. Speech errors were primarily classified according whether the resulting utterance matched, in whole or part, the name of another item in the block (i.e. a "within-block error") or another item from the same semantic category (a semantic error), in that order. This classification particularly affects interrupted errors (e.g. DOG --> CA...), where the same utterance may be classified differently depending on the other items in the block (that is, the strategy avoids assuming a semantic basis for errors a priori), and avoids the problem of trying to identify the source of an interrupted error (where open-ended classifications could depend on the coder's imagination).

These semantic and within-block errors were statistically analyzed via mixed effects logistic regression with crossed random effects of participants and items, implemented via the *glm* function in the *lme4* package (Bates, Maechler, & Bolker, 2010) in *R* (R Development Core Team, 2010). Random slopes are not included due to the sparseness of the data (~2 observations per participant). All fixed effect contrasts are centered with a mean of 0, to minimize collinearity between interactions and their constituent main effects. Binomial contrasts are coded with a range of 1; continuous or ordinal predictors are coded as continuous, with their native units (e.g. Cycle is coded as a continuous predictor with a mean of 0 and a range of -2.5 to +2.5). Non-directional p-values for interesting effects are calculated via Wald statistics. Other errors were descriptively classified but not statistically examined.

Naming latencies. Linear mixed effects regressions (via the *lmer* function, also in Bates & Maechler's, 2010, *lme4* package for *R*) also allowed concurrent subject and item analyses of the naming latency data. Since all previous blocked cyclic naming experiments have used traditional ANOVA methods, there is some question of how to adapt the analytical methods to mixed effects regression. For this and Experiment 2, I followed the ANOVA studies in analyzing un-transformed RTs[10]; this seems to be an important (and often conservative) point with respect to detecting the characteristic emergence of the Homogeneity effect over cycles, because typical data transformations (e.g. log, inverse) would discount RT differences in the long-latency first cycle. Given an observed nonlinearity in the effect of the Cycle predictor, as well as previous analyses' treatment of it as a nominal factor, I analyze it as a five-level ordinal predictor, coded via orthogonal Helmert contrasts. Helmert contrasts compare the mean at each level to the mean for all subsequent levels, thus asking whether there is a substantial change between the current data and that which follows. A full

[10] Virtually all blocked cyclic naming studies have used untransformed naming latencies. The only exception that I know of is that Biegler et al. (2008) used a log-transform to facilitate comparing patients' semantic blocking effects (mean RTs: ~800-3000ms) to those of neurally intact controls (mean RT ~700ms).

Helmert coding opens concerns about multiple comparisons, so I constrain my interpretation of individual Helmert contrasts by evaluating them sequentially: first I consider the Cycle 1 contrast, and if that reaches significance, I consider the Cycle 2 contrast, and so on. All other ordinal and continuous predictors are coded as centered linear effects (mean=0), and all binomial predictors are coded with a mean of 0 a range of 1. To reduce some variance in the model, I also include as a predictor the naming latency for the previous valid response (e.g. Baayen & Milin, 2010); it amounts to a recognition that sometimes people are fast because they are fast, and sometimes they are slow because they are slow (e.g. Wagenmakers et al., 2004). All other fixed effects in the model are theoretically motivated and orthogonal, so I retain them in the model regardless of their statistical significance.

Models include crossed random effects for participants and items. Since a full complement of random slopes – including within-participants and -items random slopes for each fixed effect – prevented the regression from converging, I used a forward-selection algorithm (*ffRanefLMER.fnc*, from Tremblay, 2011) to incrementally add random slopes that would improve model fit (at $p<.05$) without disrupting its convergence. The random effects structure did not include random correlations, due to a limitation in the Markov chain Monte Carlo algorithm that currently governs the estimation of p-values for linear mixed effects regression models (i.e. *pvals.fnc*, from Baayen, 2011). All reported p-values are non-directional.

Results

Errors

Participants named pictures correctly on the vast majority of trials, generating errors on just over one percent of total responses (Table 4). The 45 semantic errors (37 Homogeneous, 8 Heterogeneous) are of particular interest, since these are expected to reflect cumulative semantic interference (Oppenheim et al., 2010; Schnur et al., 2006), and these certainly appear more numerous in the Homogeneous blocks. This bias alone is not strong evidence for a semantic blocking effect on speech errors, though, since any slip to the name of another picture in the block (a "within-block error") would necessarily be classified as a semantic error in the Homogeneous condition, but not in the Heterogeneous one. In this experiment, errors often seemed to match block-coordinates regardless of the semantic context (possibly reflecting strategic preparations or statistical learning of picture-to-picture transition probabilities). Therefore it may be more appropriate to estimate semantic effects by collapsing across two categories of word errors: those that resulted in semantic category coordinates and those that resulted in the names of other block coordinates. This measure thus controls for a generalized tendencies towards anticipations or perseverations. A logistic regression of the resulting error distribution offers unambiguous evidence that participants were nearly twice as likely to produce an incorrect name in the Homogeneous blocks as in the Heterogeneous (OR=2.01, p=.025[11]). This semantic effect fits with several previous demonstrations of such effects in both healthy and impaired speakers (e.g. Belke, Brysbaert, Meyer, & Ghyselinck, 2005; Belke, Meyer, & Damian, 2005; Ganushchak & Schiller, 2008; Jefferies, Baker, Doran, & Ralph, 2007; Schnur, Schwartz, Brecher, & Hodgson, 2006; Wilshire & McCarthy, 2002).

Table 4. Naming errors from Experiment 1.

		Session1		Session 2	
Outcome	Example	Homog	Heterog	Homog	Heterog
Correct	*dog↦dog*	1640	1640	1622	1634
Semantic error		17	5	20	3
Completed	*dog↦goat*	6	5	5	3
Interrupted	*dog↦go...*	11	0	15	0
Other within-block error		0	9	0	4
Completed	*sock↦bed*	0	1	0	0
Interrupted	*sock↦be...*	0	8	0	4
Total within-block errors		*17*	*14*	*20*	*7*
Omission	*dog↦...*	1	0	0	0
Other error		16	11	26	24
Perspective	*dog↦animal*	7	2	11	9
Phonological	*dog↦tog*	3	2	4	1
Disfluency	*dog↦d... uh... dog*	3	2	3	3
Other or uncodable	*dog↦ti... @#%!*	3	5	8	10
Lipsmacks, *etc.*	*dog↦[pop]... dog*	52	60	58	60
Equipment failure	*[static]*	2	3	2	4

[11] Since a detailed regression with only 58 errors presents some power issues (assuming that one would typically want at least 10-20 errors per fixed effect), I note that the main effect of semantic context would remain significant at a similar odds-ratio and p-value if omitting the other factors.

Table 5. Logistic regressions for combined Semantic and Within-block errors (n=58):

	Coef β	SE(β)	p	OR (exp(β))
(Intercept)	−5.37	0.24	**<.001**	0.00
SemanticContext	0.70	0.31	**0.025**	2.01
Cycle	−0.08	0.09	0.35	0.92
Session	−0.31	0.30	0.31	0.73
SemanticContext:Cycle	0.23	0.17	0.19	1.26
SemanticContext:Session	0.91	0.61	0.13	2.49

We can also note two encouraging (n.s.) trends in the semantic/within-block error distributions. First, the semantic blocking effect in this experiment emerged across cycles (Cycle X Semantic context: OR=1.26, p=.19), recalling similar demonstrations for error patterns of aphasia (e.g. Jefferies et al., 2007; Schnur, Schwartz, Brecher, & Hodgson, 2006) and more generally supporting the notion that error patterns of brain-damaged individuals show continuity with those of unimpaired speakers (Dell, Schwartz, N. Martin, Saffran, & Gagnon, 1997). Second, the semantic blocking effect in the second session appears stronger than that in the first (Session X Semantic context: OR=2.49, p=.13), which is what we might expect if the semantic interference that accumulated in the first session persisted across the one-hour break. So, while the small error counts must caution against serious detailed consideration of the error data, the observed distributions are very much in line with previous findings for normal and impaired speakers, and appear consistent with the notion that cumulative semantic interference hinders lexical access up to an hour later.

Naming latencies

Participants named pictures correctly on 94.6% of trials. Following practices from previous experiments (Table 3), I excluded 9 responses (0.13%) with latencies of less than 250ms or more than 1500ms, thus leaving 6,527 (94.4%) observations for the naming latency analyses. Table 6 summarizes the latency data by Session, Cycle, and Semantic context, and Figure 17a conveys the same information graphically. In a nutshell, the main effect of Semantic context was approximately twice as large in the second session as in the first – which is consistent with the hypothesis the cumulative semantic interference reflect long-lasting changes to lexical access – but this difference was still relatively small and did not reach significance.

Table 6. Mean naming latencies for Experiment 1. To factor out subject and item effects, these are estimated, per session/semantic context/cycle, via a very simple linear mixed effects regression model that includes exactly those factors, plus crossed random effects for subjects and items.

	Session 1									Session 2									Interaction		
	Heterogeneous			Homogeneous			Diff $(Hom_1 - Het_1)$			Heterogeneous			Homogeneous			Diff $(Hom_2 - Het_2)$			$Diff_2 - Diff_1$		
Cycle	RT (ms)	[95% HPD]		RT	[95% HPD]		Est.	[95% HPD]		RT	[95% HPD]		RT	[95% HPD]		Est.	[95% HPD]		Est.	[95% HPD]	
1	720.6	689.5	753.2	706.1	673.8	735.0	-13.7	-33.7	4.8	737.6	704.9	767.0	730.3	696.8	761.6	-6.6	-24.8	12.9	7.3	-19.4	33.3
2	568.4	538.5	598.7	563.7	533.8	594.5	-4.8	-21.2	15.1	565.2	532.3	594.6	566.5	532.6	595.9	1.3	-17.7	18.2	6.7	-18.4	31.4
3	555.7	523.3	587.3	572.8	540.8	603.4	17.0	-2.3	35.1	556.8	524.0	586.8	573.0	542.7	604.0	15.9	-5.0	33.0	-1.2	-28.2	23.0
4	567.6	535.6	598.7	579.8	543.4	606.8	12.1	-5.5	30.5	566.5	533.3	595.5	579.3	549.3	612.7	12.6	-6.1	30.0	1.1	-25.0	27.0
5	567.1	530.7	595.1	582.0	547.0	611.5	14.3	-3.3	31.3	559.8	528.2	591.6	580.6	547.9	610.6	20.2	2.2	38.6	5.5	-20.0	31.5
6	565.6	536.7	597.8	567.2	534.3	596.9	1.8	-16.4	18.3	564.7	535.6	596.7	581.7	548.9	612.0	17.3	-1.1	37.3	16.0	-10.1	43.0

a.) b.)

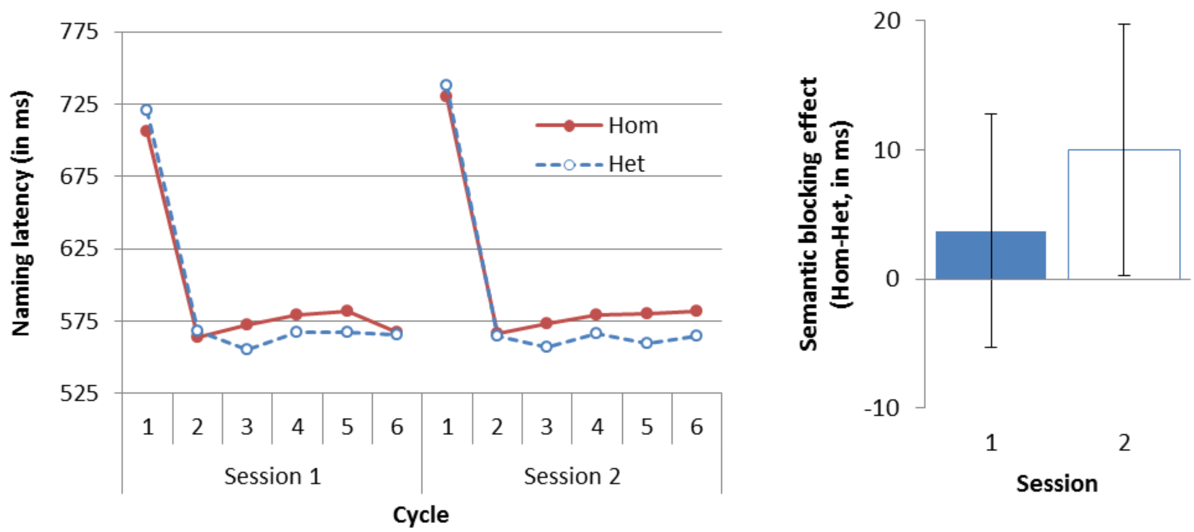

Figure 17. Mean naming latencies (a) and estimated semantic blocking effects (b) for Experiment 1. Estimates in (b) are simple main effects, drawn from a linear mixed effects model; error bars depict 95% highest posterior density intervals.

Table 7 summarizes the results of a linear mixed effects regression of the naming latency data. Mean naming latencies were quite similar for the two sessions (First session: 592.81ms, Second session: 595.31; main effect of Session: p=.75), and sequential trials showed significant autocorrelation (main effect of Previous Trial RT: Estimate= 0.10, p<.001). Unsurprisingly, naming latencies showed substantial repetition priming, dropping ~145ms after the first cycle (main effect of the Cycle 1 Helmert contrast: $p<.001$).[12] Within a cycle, each picture was named, on average, ~26ms *faster* than the previous (main effect of Position in cycle: $p<.001$), suggesting that the small cycles may have allowed participants to anticipate the presentation of certain pictures before the pictures actually appeared[13] (cf Elman, 1990). The trend is essentially chaff here, suggesting that blocked-cyclic naming tasks may reflect task-specific strategies in addition to normal speech production processes (Belke, 2008; Hsiao et al., 2009; Oppenheim et al., 2010), but it more generally supports assertions that speakers can prepare utterances at a relatively fine level of detail before actually triggering them for overt production (e.g. Griffin & Bock, 1998).

[12] An increase after the Cycle 2, though significant, seems to be driven specifically by longer RTs in the Homogeneous condition.

[13] The appearance of this effect in Cycle 1 seems to run counter to this interpretation, but trial-level data suggests that the Cycle 1 effect (unlike those in Cycles 2:6) is driven entirely by its first trial – the first trial of the block – so it seems reasonable to assume that it simply reflects a cost of initiating the naming task for each block.

Table 7. Summary of the linear mixed effects regression for Experiment 1.

	Estimate	[95% HPD Int.]		p
(Intercept)	594.06	569.05	619.84	**<.001**
Previous trial RT	0.10	0.06	0.14	**<.001**
Session (1:2)	2.50	−13.15	17.26	0.75
Cycle (1:6, Helmert)				
$\Delta Cycle_1 \rightarrow Cycle_{>1}$	−145.44	−161.95	−128.49	**<.001**
$\Delta Cycle_2 \rightarrow Cycle_{>2}$	10.91	4.15	18.19	**<.001**
$\Delta Cycle_3 \rightarrow Cycle_{>3}$	5.82	−1.15	12.81	0.11
$\Delta Cycle_4 \rightarrow Cycle_{>4}$	−2.29	−9.72	5.16	0.55
$\Delta Cycle_5 \rightarrow Cycle_{>5}$	−2.59	−11.45	5.93	0.55
Position in cycle (1:3)	−25.69	−32.54	−19.25	**<.001**
× Cycle				
$\Delta Cycle_1 \rightarrow Cycle_{>1}$	15.89	−0.41	31.78	**0.05**
$\Delta Cycle_2 \rightarrow Cycle_{>2}$	12.30	3.61	21.31	**0.01**
$\Delta Cycle_3 \rightarrow Cycle_{>3}$	2.43	−6.28	11.28	0.58
$\Delta Cycle_4 \rightarrow Cycle_{>4}$	5.62	−3.63	14.70	0.23
$\Delta Cycle_5 \rightarrow Cycle_{>5}$	1.92	−8.54	12.55	0.73
Semantic context (Het,Hom)	7.00	0.09	13.80	**0.04**
× Session	5.63	−4.45	15.99	0.29
× Cycle				
$\Delta Cycle_1 \rightarrow Cycle_{>1}$	17.62	3.22	33.16	**0.02**
$\Delta Cycle_2 \rightarrow Cycle_{>2}$	14.58	0.13	27.81	**0.04**
$\Delta Cycle_3 \rightarrow Cycle_{>3}$	−1.80	−15.48	12.90	0.8
$\Delta Cycle_4 \rightarrow Cycle_{>4}$	0.42	−15.24	14.90	0.96
$\Delta Cycle_5 \rightarrow Cycle_{>5}$	−8.47	−25.76	8.74	0.34
× Position in cycle	0.20	−5.71	6.87	0.97
× Cycle				
$\Delta Cycle_1 \rightarrow Cycle_{>1}$	−27.06	−45.02	−9.17	**<.001**
$\Delta Cycle_2 \rightarrow Cycle_{>2}$	−1.81	−19.78	14.30	0.83
$\Delta Cycle_3 \rightarrow Cycle_{>3}$	14.00	−2.78	32.19	0.12
$\Delta Cycle_4 \rightarrow Cycle_{>4}$	2.06	−16.49	20.88	0.84
$\Delta Cycle_5 \rightarrow Cycle_{>5}$	2.00	−19.47	22.82	0.86

Semantic interference effects

Overall, participants took ~7ms longer to name pictures in blocks with other pictures from the same semantic category (Homogeneous; e.g. DOG, BEAR, GOAT) than in blocks without (Heterogeneous; e.g. SOCK, PIE, BED; main effect of Semantic context: p=.04). Though significant, this manifestation of the blocking effect is smaller than what has be shown in other studies, presumably due to the inclusion of so few (n=3) Homogeneous condition trials per cycle; in the Dark Side model, this would produce smaller blocking effects by increasing the ratio of connection-strengthening experiences to connection-weakening experiences. It did, however, increase across cycles, with Semantic context yielding significant interactions for both the first (p=.02) and second (p=.04) cycle Helmert contrasts. Some increase in the blocking effect across cycles is typical (Belke, Meyer, et al., 2005; Damian & Als, 2005; Schnur et al., 2006), although it is remarkable that in the current experiment, the expected main effect– longer naming latencies in Homogeneous blocks – did not emerge until after the second cycle. I will return to this point in the Discussion, as it offers novel constraints on the development of the semantic blocking effect.

Since the semantic blocking effect emerged later than expected, it is worth asking whether the semantic context manipulation may have elicited some more subtle consequence beforehand (Table 8). Howard et al (2006) demonstrated that, when naming pictures from a single semantic category, interspersed with filler items, naming latencies increased by ~30ms (±8.2, revised to ~24ms by Alario & Moscoso Del Prado, 2010; Nickels et al.'s, 2008, similar Experiments 1 & 3 reported effects of 24ms and 20ms, respectively) for each member of the category that had previous been named. Oppenheim et al.'s (2010) Dark Side model showed how such a pattern could also produce semantic blocking effects in blocked cyclic naming. Following that suggestion, we can ask whether a homologue of Howard et al.'s ordinal position effect (i.e. an incremental interference effect) might be found in the blocked cyclic naming paradigm. In blocked-cyclic naming, such an effect would manifest as an interaction of semantic context with the ordinal position of a trial within the current cycle. Although the current experiment offers no evidence (Estimate[14]=0.2ms, p=.97) of the sustained sawtooth that the Dark Side model predicts (e.g. Figure 5; cf Cycle 1 Helmert contrast for Semantic context X Position in cycle, from Table 7: Estimate=-27.1ms, p<.001), the semantic context-by-ordinal position interaction (i.e. the incremental interference effect) appears in the first cycle exactly as expected: a ~23ms increase with ordinal position (Semantic context X Position in cycle, for a regression restricted to Cycle 1 data: Estimate=22.6ms, p=.03), which neatly fits the revised estimates for the incremental interference effect. This provides the first direct evidence that

[14] Curiously, after the first cycle the data actually show a non-significant trend in the opposite direction: within each cycle, naming latencies appear to speed up more in the Homogeneous condition. Given the overall acceleration of naming latencies within each cycle, such a pattern could plausibly reflect strategic response preparation (with the assumption that the strategy is more effective for related items), though it could also reflect some kind of more automatic incremental semantic facilitation (Damian & Als, 2005; Wheeldon & Monsell, 1994). In any case, the trend disappears in Experiment 2, and is replaced by a significant effect in the opposite direction (i.e. the predicted sawtooth) in Experiment 3, consistent with the idea that a separate process may obscure cumulative semantic interference at short lags or in small cycles. Note that this point is separate from the point elsewhere that such temporary semantic facilitation is not sufficient to explain the absence of a blocking effect in the first cycle.

the incremental interference effect is present in blocked-cyclic naming data, consistent with the Dark Side's claim that the semantic blocking effect develops incrementally (Oppenheim et al., 2010).

Table 8. Summary of a linear mixed effects regression for Experiment 1, restricted to the data from Cycle 1. This regression suggests two additional patterns that do not seem immediately relevant to the present investigations. First, even though mean naming latencies are similar across sessions, they differ quite a lot in the first cycle (a Session * Cycle interaction but that does not seem theoretically interesting). Second, there is a little support for the autocorrelation predictor in Cycle 1, though this seems to likely reflect rather mundane factors.

	Estimate	[95% HPD Int.]		p
(Intercept)	724.4	690.31	757.77	**<.001**
Previous trial RT	0.0	−0.08	0.09	0.95
Position in cycle	−35.0	−49.46	−20.72	**<.001**
Session	20.5	2.00	38.49	**0.03**
Semantic context	−8.8	−24.88	7.52	0.28
× Position in cycle	22.6	1.68	42.53	**0.03**
× Session	7.5	−23.86	38.24	0.64

However, the major effect of interest – a hypothesized larger semantic blocking effect in the second session than in the first – received only suggestive support (interaction of Semantic context with Session: Estimate=5.63ms, p=.29). Though not significant, we can still consider the estimated interaction because it provides the only data that we have so far that is relevant to the longer-term persistence of cumulative semantic interference. The ~6-ms estimate for the interaction indicates that that the blocking effect was approximately twice as large in the second session as in the first (i.e.

Session 1: 4.19; Session 2: 9.82ms; see Figure 17b), which is reasonably consistent with the trends from Belke et al.'s (2005) experiment where the second batch of items was named immediately after the first. Although the main effect of blocking (i.e. Homogeneous>Heterogeneous) did not survive the switch to new items, as it did in Belke et al.'s experiment, the reverse blocking effect Homogeneous<Heterogeneous) in the first cycle of the second session was at least numerically smaller than that in the first cycle of the first session (Semantic context X Session, for a regression restricted to Cycle 1 data: Estimate=7.5ms, p=.64). Thus, although the experiment failed to detect the expected interaction that would have demonstrated persistent consequences of cumulative semantic interference, the patterns in the data remain encouraging with respect to its possible existence.

Experiment 1 discussion

In line with the hypothesis that cumulative semantic interference reflects a persistent restructuring of meaning-driven word retrieval, Experiment 1 offers some indication that cumulative semantic interference lasts long enough and remains strong enough to impair lexical access at least an hour later. This persistent interference effect could not be statistically confirmed, but the trend offers some hope.

This experiment also demonstrated, for the first time, the presence of a Howard et al (2006)-style incremental interference effect in a blocked-cyclic naming paradigm, thus supporting claims that the semantic blocking effect develops incrementally. This finding is noteworthy because it has previously been claimed that semantic blocking effects tend to emerge only in the second cycle because they are initially nullified by some short-lived incremental process (Damian & Als, 2005), for example either an incremental semantic facilitation (e.g. persistent activation or learning at the semantic level; e.g. Wheeldon & Monsell, 1994) or a post-selection inhibition that stops a recently used word from competitively inhibiting the selection of a related word (e.g. Vitkovitch, Rutter, & Read, 2001). Finding an apparently intact incremental semantic interference effect in the first cycle – in the absence of a main effect of semantic blocking seems to argue against such explanations. What else might disrupt the appearance of a blocking effect in the first cycle? One possibility, suggested by the current work, is that semantic interference effects also affect items named in Heterogeneous blocks, since their semantic competitors have often appeared in previous blocks. Taking such overlap into account, the Dark Side model would predict that the expected semantic blocking effect – indexing the accumulation of semantic interference in Homogeneous blocks – should initially be obscured by persisting interference that slows retrieval in the baseline Heterogeneous blocks.

Remarkably, while the incremental interference effect was quite strong in the first cycle of naming the main effect of semantic blocking did not emerge in this experiment until the *third* cycle (Figure 17a). A far more typical pattern, characterizing the results from every previously published blocked-cyclic naming experiment, is for the effect to emerge in the *second* cycle (Belke, Brysbaert, et al., 2005; Belke, Meyer, et al., 2005; Damian & Als, 2005; Schnur et al., 2006); but see Damian & Als, 2005, Expt 4b, and Belke, 2008, Experiment 1 for specific manipulations that led to its emergence in Cycle 1). It has even been claimed that, once established in the first cycle, the blocking effect remains stable for all subsequent cycles (Belke, Meyer, & Damian, 2005; but see Biegler et al., 2008, Experiment 1a, Belke, Meyer, & Damian, 2005, Experiment 3; Schnur, Schwartz, Brecher, & Hodgson, 2006; Schnur et al., 2009 for examples to the contrary). The standard explanation is that a semantic context only begins to influence lexical access after it has been established in the first cycle. In fact, many analyses of semantic blocking effects exclude data from the first cycle, on precisely those grounds (Belke & Meyer, 2007; Belke, Brysbaert, et al., 2005; Belke, Meyer, et al., 2005; Damian, 2003; Damian et al., 2001; Vigliocco et al., 2002). The question is, how exactly does the first cycle of a block establish the semantic context? The presence of an incremental interference effect in the first cycle, as noted above, suggests that the semantic blocking effect is established incrementally. But we might also make a distinction between type- and occurrence-based accounts of this incrementality. Type-based accounts (Abdel Rahman & Melinger, 2009; Nickels et al., 2008) would argue that what matters for semantic interference is the number of unique types competing for selection, in this case the number of activated competitors: each time you name a *novel* competitor, it adds another competitor to the arena and thus slows the lexical selection process a bit more for the next related target. Occurrence-based accounts (Navarrete, Mahon, & Caramazza, 2008; Oppenheim et al., 2010), by contrast, argue that what matters is the number of times that any competitor has been named (novel or not).[15] The incremental interference

[15] A consequent distinction between these accounts is that type-based accounts can afford to posit

effect in Cycle 1 is of course compatible with both accounts, but the delayed emergence of the semantic blocking effect in Experiment 1 seems to argue against the type-based explanation. All of the competitors that would be named in a block had already been named by the end of the first cycle, so the second cycle could not contribute any novel competitors, and therefore, a type-based account would have predicted that the semantic blocking effect should have already reached its maximum, whereas an occurrence-based account might predict a continued increase. This experiment showed a significant increase in the size (and sign) of the blocking effect after all competitors had been established, thus supporting an occurrence-based account over a type-based one.[16]

---

linear effects for the number of novel competitors, whereas occurrence-based accounts seem to require *nonlinear* effects for the number of competitor trials (or assume some mediating factor, e.g. processing difficulty).

[16] One mediating possibility, though, is that the semantic blocking effect actually requires time to develop. For instance, an initial semantic facilitation (e.g. Wheeldon & Monsell, 1994) my require several seconds to decay, or the weight changes that underlie cumulative semantic interference may need to time to 'sprout' after being planted by experience (though the first cycle incremental interference effect in this experiment seems to argue against that possibility). Another possibility, suggested above, is that the growth of the semantic blocking effect across cycles reflects, in part, the remediation of persistent semantic interference effects in Heterogeneous blocks, on the assumption

Experiment 2

Although Experiment 1 revealed several interesting patterns, it could not demonstrate significant carryover of cumulative semantic interference from the first session to the second. Part of the problem may have been that its semantic blocking effect was itself relatively weak, making it difficult to detect whether it might be influenced by other factors. Previous studies reporting large interference effects have all used cycles of at least four to six related pictures (Table 3). It may be that the three-picture cycles in this experiment approached a boundary condition where cumulative semantic interference is less robust, or more mundane factors may have disrupted the effect. For instance, the task may have lent itself too easily to response strategies that have little to do with normal word retrieval processes. Other researchers have previously suggested that the blocked-cyclic paradigm may generally lend itself to such strategies (e.g. Belke, 2008) and this potential surely increases when a person has fewer pictures to name. For instance, by taking the cyclic structure of a block into account, subjects may be able to prepare their responses before a picture actually appears (cf. Bock & Griffin, 1998). In fact, the naming latencies give some indication that participants may have anticipated the presentation of certain pictures before the pictures actually appeared. Within a

---

that semantic interference can be quickly overcome via repetition in a manner similar to that described for lexical frequency effects (D. L. Scarborough, Cortese, & H. S. Scarborough, 1977) and partly captured by the error-proportional learning implemented in the Dark Side model.

cycle, each picture was named, on average, 22ms faster than the last, a very strong effect and precisely the sort of pattern that one would expect if a participant were using each stimulus to constrain their predictions for the next (cf Elman, 1990). Such advance preparation would obviously be problematic for an experiment that considers response time relative to a stimulus onset, because response times may not consistently reflect the full time needed to plan and execute a response. With this in mind, Experiment 2 modifies the procedure from Experiment 1, in the hopes that, by reducing the opportunity for task-specific strategies, naming latencies might better reveal consequences of the implicit learning that is proposed to underlie cumulative semantic interference.

To reduce the potential for strategic response preparation, Experiment 2 directly interleaved items from the Homogeneous and Heterogeneous conditions to create 6-item cycles (Figure 18).[17] Damian and Als (2005, Experiment 4) previously used such a manipulation to explore both the short term (Lag-1) persistence of the semantic blocking effect and possibility that it might be attenuated in the first cycle by a short-lived semantic facilitation or self-inhibition process. Unlike Damian and Als' interleaving however, here I sought to reduce predictability of the trials by adding some randomness into the interleaving. In Damian and Als' experiment, the Homogeneous and Heterogeneous slots alternated (e.g. Hom, Het, Hom, Het, Hom, Het). In the present experiment,

[17] Although this manipulation obliterates the distinction between Homogeneous and Heterogeneous blocks or contexts, the blocked-cyclic naming terminology remains useful in describing the experiment.

the slot sequence was randomized, such that for each sequence of two slots, the first would be assigned to either the Homogeneous or Heterogeneous condition, and the second would be assigned to the other (e.g. Hom, Het, Het, Hom, Hom, Het).[18] Aside from integrating the Homogeneous and Heterogeneous conditions, Experiment 2 basically replicated 1.

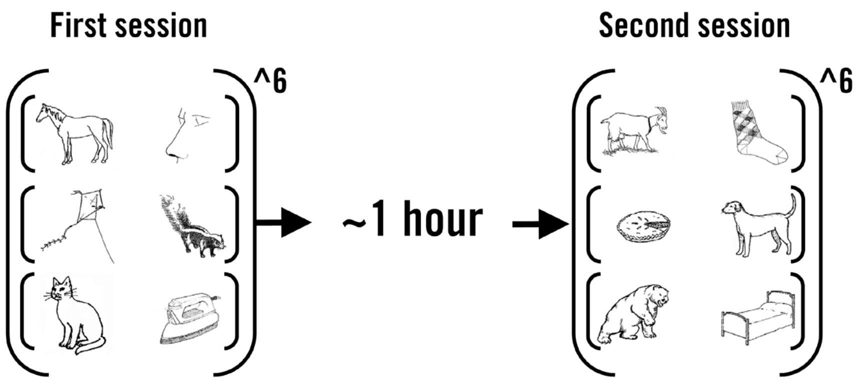

Figure 18. Experiment 2 structural overview. Whereas Experiment 1 presented Homogeneous and Heterogeneous sets in separate blocks, here each block contained both a Homogeneous set and a Heterogeneous set. Within each cycle stimuli were sampled randomly without replacement such that each sequence of two (e.g. 1-2, 3-4, 5-6) contained one item from the Homogeneous set and one from the Heterogeneous set. As in Experiment 1, the Homogeneous condition in Session 2 presented novel exemplars from the same

[18] Note that although the mean/modal lag between Homogeneous exemplars in Experiment 2 is 1 – as in Damian & Als' experiment – the lag here actually varies between 0 and 2, with 1 in 5 trials at Lag-0, and thus may not optimally serve Damian & Als' goal of avoiding a short-lived facilitation effect.

Homogeneous categories that appeared in Session 1, and the Heterogeneous condition in Session 2 presented novel exemplars from novel categories that had not appeared in Session 1.

Methods

Methods for Experiment 2 were essentially identical to those described for Experiment 1. The only major change was that the Homogeneous and Heterogeneous blocks were now directly interleaved.

Participants

Twenty-four University of Illinois undergraduate students participated in exchange for course credit. All were native monolingual US English speakers with normal or corrected-to-normal vision and no reported history of language impairments. None had participated in any of the other experiments or simulations described herein.

Materials

Materials were identical to those in Experiment 1, with the exception that, having recognized its regionally dominant name while running Experiment 1, the 'sofa' picture was renamed 'couch'.

Design

The design of Experiment 2 was identical to that of Experiment 1, except for changes resulting from the integration of Homogeneous and Heterogeneous blocks. Thus the experiment included two 144-trial sessions, separated by a one hour delay. Each session included four six-cycle blocks, with six pictures each appearing once in each cycle (for a total of 36 trials per block). Within each block, the six pictures always included three pictures from a single semantic category (e.g. horse, skunk, cat, i.e. Homogeneous items) and three from other, different, semantic categories (e.g. nose, kite, iron, i.e. Heterogeneous items). For each block in the second session, the Homogeneous items represented novel exemplars from one of the semantic categories that had appeared as a Homogeneous-item category in the first session (e.g. dog, bear, goat, to follow the horse et al. example), while the heterogeneous items always represented novel exemplars from semantic categories that had not appeared in the first session (e.g. sock, pie, bed). Thus, a participant always named the same categories in the homogeneous conditions in both sessions, and never named items from the same categories in the heterogeneous conditions in both sessions. Creation and assignment of item sets to structural positions in this experiment proceeded as in Experiment 1.

Procedure

The procedure for Experiment 2 was identical to that for Experiment 1 in all relevant respects. The time between the two sessions was mainly occupied by the inner speech tonguetwister experiment described in Oppenheim and Dell (2010; as in Experiment 1), and a minor variation that included an articulatory suppression component.

Equipment

As in Experiment 1, presentation was controlled via E-Prime. Responses were digitally recorded via a headmounted microphone, and transcribed offline. Naming latencies were collected via an E-Prime serial response box, and later verified via visual waveform analysis.

Data analysis

Since the experiments were essentially identical, so is the data analysis, including the regression models. Six participants were replaced because the previously specified per-trial exclusions left less than 80% of their responses for the naming latency analysis. As in Experiment 1, the present analysis includes a factor for ordinal position within a cycle. To keep it consistent with Experiment 1, and to facilitate comparison with Howard et al.'s (2006) results, I calculate this as a set-specific ordinal position with values of 1:3 within each cycle (i.e. a category-specific count, but including the equivalent coding for Heterogeneous trials).

Results and discussion

Errors

Participants named pictures correctly on the vast majority of trials, generating naming errors on just over one percent of total responses (Table 9). As in Experiment 1, the 30 semantic and within-block errors were roughly twice as likely on Homogeneous condition trials compared to Heterogeneous (21 Homogeneous, 9 Heterogeneous), but with relatively few errors, this trend achieved only marginal significance (OR=2.1, p=.09). Unfortunately, the small error counts in this experiment (i.e. 30 errors, versus 58 in Experiment 1) do not provide sufficient data for further statistical analyses.

Table 9. Errors from Experiment 2.

		Session1		Session 2	
Outcome	Example	Homog	Heterog	Homog	Heterog
Correct	*dog→dog*	1572	1630	1543	1566
Semantic error		10	2	7	2
Completed	*dog→goat*	5	1	3	0
Interrupted	*dog→go. . .*	5	1	4	2
Other within-block error		2	2	2	3
Completed	*sock→bed*	0	0	0	0
Interrupted	*sock→be. . .*	2	2	2	3
Total within-block errors		*12*	*4*	*9*	*5*
Omission	*dog→. . .*	0	0	2	0
Other error		8	13	9	14
Perspective	*dog→animal*	3	3	2	8
Phonological	*dog→tog*	1	0	0	1
Disfluency	*dog→d. . . uh. . . dog*	4	6	4	4
Other or uncodable	*dog→ti. . . @#%!*	1	3	5	3
Lipsmacks, *etc.*	*dog→[pop]. . . dog*	119	74	144	124
Equipment failure	*[static]*	16	8	19	16

Table 10. Logistic regression for semantic blocking effects with semantic and within-block errors (sum=30)

	Coef β	SE(β)	p	OR (exp(β))
(Intercept)	−6.18	0.29	**<.001**	0.0
Semantic context (Het, Hom)	0.74	0.44	**0.09**	2.1

Naming latencies

Participants named pictures correctly on 91.3% of trials. The exclusion guidelines from Experiment 1 identified no correct responses with latencies of less than 250ms or more than 1500ms, and hence none were excluded. Table 11 summarizes the latency data by Session, Cycle, and Semantic context, and Figure 19a conveys the same information graphically. Briefly, the change in procedure greatly reduced the apparent indications of task-specific strategies, while replicating the results of Experiment 1 in several important respects.

Table 11. Mean RTs from Experiment 2, estimated as in Table 6.

	Session 1									Session 2									Interaction		
	Heterogeneous			Homogeneous			Diff $(Hom_1 - Het_1)$			Heterogeneous			Homogeneous			Diff $(Hom_2 - Het_2)$			$Diff_2 - Diff_1$		
Cycle	RT (ms)	[95%	HPD]	RT	[95%	HPD]	Est.	[95%	HPD]	RT	[95%	HPD]	RT	[95%	HPD]	Est.	[95%	HPD]	Est.	[95%	HPD]
1	683.4	661.8	707.0	683.1	657.6	705.0	-0.7	-21.6	16.8	702.2	680.2	724.7	708.6	683.5	731.2	6.2	-13.9	23.7	6.9	-20.5	32.9
2	622.6	600.0	645.3	624.6	602.9	648.3	1.8	-15.9	20.3	640.5	617.8	662.4	642.5	621.0	666.9	1.9	-16.7	19.1	-0.9	-26.6	24.9
3	607.0	584.4	630.3	616.6	592.9	639.1	8.7	-10.6	24.9	624.7	601.4	648.5	637.5	616.7	660.6	13.5	-4.7	31.6	3.7	-20.7	27.8
4	607.6	586.0	630.4	614.8	591.3	636.3	6.8	-9.3	25.9	619.2	597.3	642.3	633.8	611.7	655.8	13.8	-3.3	31.4	7.0	-17.3	32.9
5	608.6	587.6	632.4	610.5	587.4	632.1	1.6	-15.4	20.5	624.0	600.8	645.6	636.4	611.9	657.9	12.0	-6.0	30.7	10.7	-18.7	33.8
6	606.0	581.7	626.8	615.1	594.1	638.1	9.0	-8.0	27.5	627.3	603.2	648.6	629.9	607.5	653.3	2.9	-18.4	19.5	-6.4	-29.6	23.7

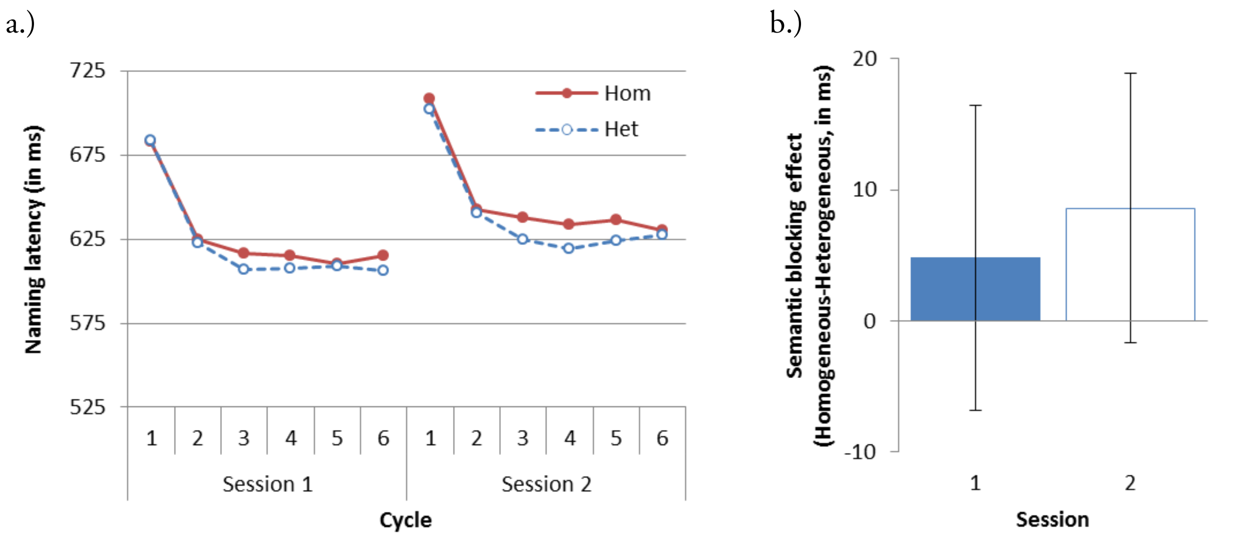

Figure 19. Mean naming latencies (a) and estimated semantic blocking effects (b) for Experiment 2. Estimates in (b) are simple main effects, drawn from a linear mixed effects model; error bars depict 95% highest posterior density intervals.

Table 12 summarizes the results of a linear mixed effects regression of the naming latency data. Mean naming latencies in the second session were substantially slower than those in first (First session: 626.24ms, Second session: 643.74ms; main effect of Session: Estimate=17.50ms, p<.001),[19] and sequential trials showed significant autocorrelation (main effect of Previous Trial RT: Estimate= 0.08, p<.001). Naming latencies were greatest in the first cycle, dropping ~70ms thereafter (main effect of the Cycle 1 Helmert contrast: p<.001), which is substantial but still far less dramatic than the 145ms change seen in Experiment 1. And unlike Experiment 1, which saw naming latencies fall dramatically within each cycle, here naming latencies *grew* by 3.6ms as a function of a picture's ordinal position within a cycle (main effect of Position in cycle (category-specific): Estimate=3.62ms, p=.02). So the changes in experiment structure appear to have reduced the potential for participants to prepare their responses before encountering a stimulus.

---

[19] This main effect of session may be problematic, depending on ones' assumptions about the additivity of the underlying effects. If we generally expect larger effects at longer base naming latencies (as in the second session), then finding even a significantly larger effect in the second session may not be so impressive. This issue is addressed in Experiment 3.

Table 12. Summary of the linear mixed effects regression for Experiment 2.

	Estimate	[95% HPD Int.]		p
(Intercept)	634.99	613.83	656.79	**<.001**
Previous trial RT	0.08	0.05	0.11	**<.001**
Session (1:2)	17.50	6.83	27.66	**<.001**
Cycle (1:6, Helmert)				
$\Delta Cycle_1 \rightarrow Cycle_{>1}$	−70.29	−86.26	−52.74	**<.001**
$\Delta Cycle_2 \rightarrow Cycle_{>2}$	−9.47	−20.84	2.37	**0.1**
$\Delta Cycle_3 \rightarrow Cycle_{>3}$	−1.97	−9.42	4.89	0.59
$\Delta Cycle_4 \rightarrow Cycle_{>4}$	0.88	−6.81	8.39	0.82
$\Delta Cycle_5 \rightarrow Cycle_{>5}$	−0.44	−9.22	8.10	0.92
Position in cycle (category-specific; 1:3)	3.60	0.49	6.71	**0.02**
× Cycle				
$\Delta Cycle_1 \rightarrow Cycle_{>1}$	−5.01	−14.33	3.80	0.28
$\Delta Cycle_2 \rightarrow Cycle_{>2}$	2.71	−5.76	11.33	0.54
$\Delta Cycle_3 \rightarrow Cycle_{>3}$	1.75	−7.27	10.61	0.69
$\Delta Cycle_4 \rightarrow Cycle_{>4}$	−3.59	−13.25	5.71	0.45
$\Delta Cycle_5 \rightarrow Cycle_{>5}$	−2.61	−13.30	8.24	0.63
Semantic context (Het,Hom)	6.96	−0.59	14.05	**0.06**
× Session	4.06	−10.80	19.27	0.58
× Cycle				
$\Delta Cycle_1 \rightarrow Cycle_{>1}$	3.99	−10.76	17.61	0.58
$\Delta Cycle_2 \rightarrow Cycle_{>2}$	7.33	−6.39	20.92	0.3
$\Delta Cycle_3 \rightarrow Cycle_{>3}$	−2.01	−16.52	11.78	0.78
$\Delta Cycle_4 \rightarrow Cycle_{>4}$	−5.63	−21.32	8.59	0.46
$\Delta Cycle_5 \rightarrow Cycle_{>5}$	−0.28	−18.10	16.77	0.97
× Position in cycle (cat.-spec.)	3.55	−3.61	10.36	0.32
× Cycle				
$\Delta Cycle_1 \rightarrow Cycle_{>1}$	−19.32	−37.94	−1.29	**0.04**
$\Delta Cycle_2 \rightarrow Cycle_{>2}$	6.38	−10.30	23.32	0.47
$\Delta Cycle_3 \rightarrow Cycle_{>3}$	−7.58	−25.68	9.56	0.41
$\Delta Cycle_4 \rightarrow Cycle_{>4}$	4.05	−15.22	22.10	0.67
$\Delta Cycle_5 \rightarrow Cycle_{>5}$	12.05	−8.45	33.63	0.27

Semantic interference effects

Pictures were named ~7ms slower in the Homogeneous condition than in the Heterogeneous (main effect of Semantic context: Estimate=6.96ms, p=.06), a marginally significant semantic blocking effect that was similar in direction and magnitude to the 7.0ms effect that was estimated for Experiment 1. Though the semantic blocking effect did not significantly increase across cycles (unlike Experiment 1), model estimates do suggest small increases after both the first and second cycles, which are numerically consistent with the patterns seen Experiment 1 (Helmert contrasts for Semantic context X Cycle interaction: Cycle 1: Estimate=3.99ms, $p=.58$; Cycle 2: Estimate=7.33ms, $p=.30$). As in Experiment 1, analyses restricted to the first cycle demonstrated that the semantic blocking effect grew by ~19ms with an item's ordinal position (1:3) within the cycle (i.e. the incremental interference effect; Semantic context X Position in cycle interaction, for a regression restricted to Cycle 1 data: Estimate=19.29, $p=.05$), though again this dropped off sharply thereafter (Cycle 1 Helmert contrast for the Semantic context X Position in cycle interaction: Estimate=-19.32, $p=.04$). This replicates a puzzle from Experiment 1: we see a large incremental accumulation of semantic interference within the first cycle (i.e. 19.29ms*2=38.58ms) that produces a comparatively small main effect of semantic blocking (main effect of Semantic context, for a regression restricted to Cycle 1 data: Estimate=3.14, $p=.72$), again suggesting that the characteristic absence of a semantic blocking effect from the first cycle of blocked-cyclic naming cannot be attributed simply to a temporary attenuation of the incremental interference effect.

Table 13. Summary of the linear mixed effects regression for Experiment 2, considering Cycle 1 and Cycles 2:6 separately.

	Estimate	[95% HPD Int.]		p
Cycle 1				
(Intercept)	696.63	673.00	721.71	**<.001**
Previous trial RT	0.02	−0.07	0.11	0.64
Session (1:2)	23.26	3.54	41.90	**0.02**
Position in cycle (category-specific: 1:3)	8.48	−1.71	19.22	0.11
Semantic context (Het,Hom)	3.14	−13.65	19.43	0.72
× Session	6.77	−24.66	38.91	0.68
× Position in cycle (category-specific)	19.29	−0.06	38.93	**0.05**
Cycles 2:6				
(Intercept)	625.35	604.88	645.87	**<.001**
Previous trial RT	0.10	0.06	0.14	**<.001**
Session (1:2)	16.48	5.79	26.90	**<.001**
Cycle (2:6)	−1.93	−5.77	1.99	0.31
Position in cycle (category-specific: 1:3)	2.91	−1.32	7.24	0.18
× Cycle	0.15	−2.15	2.43	0.9
Semantic context (Het,Hom)	7.13	−2.31	16.15	0.13
× Session	3.75	−9.59	16.60	0.55
× Cycle	0.88	−3.29	5.05	0.66
× Position in cycle (category-specific)	0.03	−8.11	8.27	1
× Cycle	0.97	−3.72	5.76	0.69

Although Figure 19a suggests a stronger semantic blocking effect in the second session than in the first, with patterns that visually resemble both Belke et al.'s no-lag results (Figure 14b) and the Dark Side simulation (Figure 14c), statistical support for this pattern was quite weak. Analyses suggest a ~4ms interaction of Semantic context with Session (p=.58), meaning that the semantic blocking effect was ~5ms in the first session and ~9ms in the second session (Figure 19b). This estimated magnitude for the interaction is reasonably concordant with that from Experiment 1 (5.63ms), and, as in Experiment 1, analyses of the Cycle 1 data suggest that the interaction was present from the beginning of the second session (Semantic context X Session interaction, for a regression restricted to Cycle 1 data: Estimate: 6.77ms, p=.68), which is consistent with the idea that the larger blocking effect in the second session reflects interference that carried over from the first. Again, these trends are what one would expect if some semantic interference persisted from the first session to affect naming latencies in the second, but the effects are still too small and possibly too variable to be statistically confirmed here. So while the results of this experiment and the previous are encouraging, they do not yet provide strong evidence that cumulative semantic interference lasts as long as incremental learning accounts would suggest.

Experiment 3

Experiments 1 and 2 each demonstrated significant incremental interference effects in a blocked-cyclic naming paradigm. Remarkably, these effects appeared full-strength in the first cycle of naming, in the absence of commensurate semantic blocking effects, suggesting that whatever is attenuating the semantic blocking effect in the early trials of blocked-cyclic naming is already present at the onset of a block. An obvious candidate would be semantic interference from previous Heterogeneous blocks that carries over to increase naming latencies in the Heterogeneous baseline condition.

But it is also remarkable that the incremental interference effects in Experiments 1 and 2 were not apparent after the first cycle. The Dark Side simulations predict that the incremental interference effect should shrink a bit with each cycle, but not so dramatically as in these experiments. There are many possible explanations for this absence (e.g. floor effects such that repeated retrievals generate minimal error to drive learning, strategic preparation afforded by the use of small numbers of items, direct stimulus-response mappings such that repeated retrievals depend less on semantic activation, type-based interference such that only the number of unique exemplars matters, or momentum in weight changes could smooth out effects), but most (with the possible exception of type-based interference) would predict that the incremental interference effect should return in subsequent cycles if same-category trials are spaced further apart. Navarrete et al.'s (2010) Experiment 1 results (Figure 20) provide some indication that it does return, showing a sustained incremental interference effect when adapting Howard et al.'s (2006) experiment to include several large (96-item) cycles.

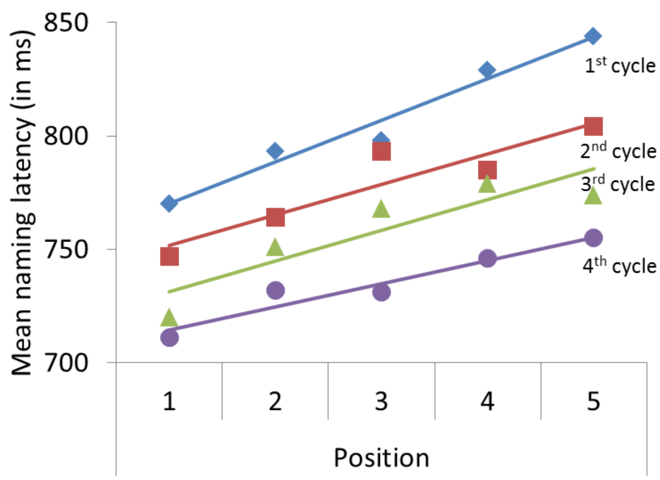

Figure 20. Incremental interference effects for mean naming latencies over four large cycles from Navarrete et al. (2011, Experiment 1).

Consistent with the hypothesis that cumulative semantic interference reflects the persistent restructuring of semantic-to-lexical mappings, Experiments 1 and 2 each showed trends in the expected direction of a stronger semantic blocking effect in the second session than in the first. But in each experiment that interaction failed to achieve statistical significance. Experiment 2 reduced a potential confound from Experiment 1 (i.e. predictability), but still seemed to suffer from a lack of power. For instance, the known-true semantic blocking effect in each experiment was relatively small compared to previous published studies, and only barely detectable. It seems reasonable to assume, as the data suggest, that the increase in the blocking effect across sessions should be smaller than the blocking effect itself, and with the observed variance, statistically detecting that increase would still require far more data than seems practical.

Another limitation is that Experiments 1 and 2 were only designed to evaluate whether cumulative semantic interference might produce *some* lasting effect; they had no way to evaluate whether that effect was smaller after an hour or the same size. Finally, it may also be problematic that Experiments 1 and 2 evaluated the persistence of cumulative semantic interference by comparing the blocking effect at one time point to that at another later time point. Larger blocking effects in the second session could thus merely reflect fatigue (e.g. overriding prepotent competitors in the Homogeneous condition of blocked-cyclic naming requires strong executive control, which is particularly impaired after an hour of tiresome psych experiments). Such confounds could be avoided by doing comparisons within the same session.

Experiment 3 seeks to address all of these concerns. As illustrated in Figure 21, it draws on the structure of Experiment 2 while incorporating a within-block design similar to that of Navarrete et al.'s Experiment 1. Exemplars from eight categories are now presented in 24-item 'super-blocks'. Within each cycle the resulting experiment structure resembles that of Howard et al. (2006) and within each block it resembles Navarrete et al.'s (2010) Experiment 1. Thus it forgoes the semantic context manipulation (Homogeneous versus Heterogeneous) entirely; but it uses pairwise ratings of semantic relations to investigate an expected homologue of the semantic blocking effect. To evaluate the persistence and decay of cumulative semantic interference, Experiment 3 compares naming latencies within the same super-block. In the Figure 21 example, we first compare naming latencies for mammals to those for body parts in the first block; there should be no difference here, so this contrast serves as a control. To the extent that cumulative semantic interference persists over an hour delay, people should be slower in the second block to name mammals than utensils. And, to the extent that cumulative semantic interference decays over that delay, people should be slower in the third block to name utensils than body parts.

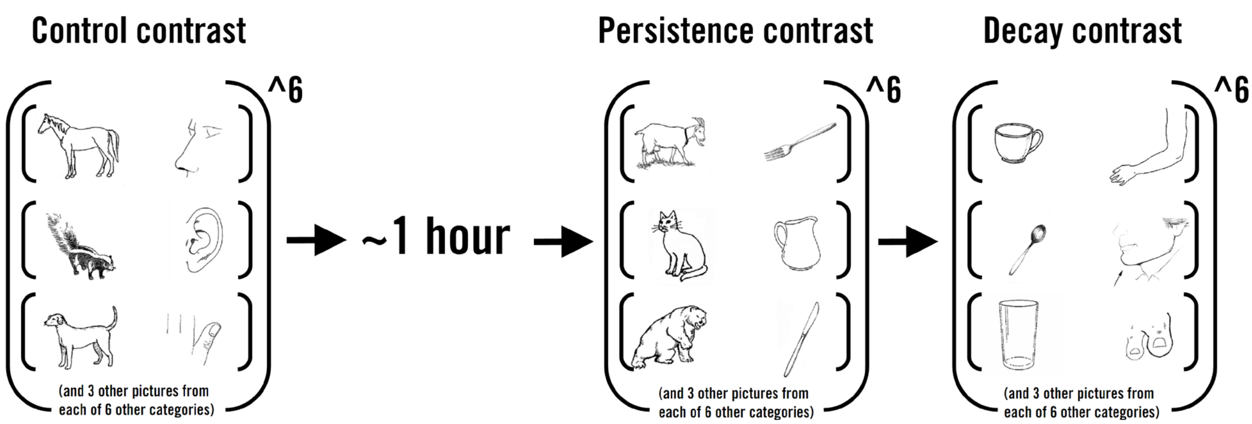

Figure 21. Experiment 3 structural overview. Each block contained 24 items from 8 semantic categories. Within each cycle stimuli were sampled randomly without replacement as in Experiment 2. In the second block, half of the items represent novel exemplars from 4 categories that appeared in the first block (e.g. mammals), and half represent novel exemplars from other categories (utensils); the difference in their naming latencies indexes the persistence of cumulative semantic interference. In the third block, half of the items

represent novel exemplars from the 4 categories introduced in the second block (utensils), and half represent novel exemplars from the remaining categories that were introduced in the first block (body parts); the difference in their naming latencies indexes the decay of cumulative semantic interference.

Methods

Subjects

Forty-eight University of Illinois undergraduate students participated in exchange for course credit. All were native monolingual US English speakers with normal or corrected-to-normal vision and no reported history of language impairments. None had participated in any other of these experiments. Two subjects were replaced due to equipment failure.

Materials

Pictures for this experiment were identical to those used in Experiments 1 and 2. They consisted of 72 black-and-white line drawings of object nouns from Schnur et al. (2006), including six exemplars from each of 12 semantic categories (e.g. mammals, shapes, appliances).

Design

Categories were divided into three sets, *a*, *b*, and *c*, with the goal of minimizing semantic overlap between sets. Thus one set consisted of Body Parts, Clothing, Roles, and Toys, while another consisted of Animals, Nature, Plants, and Shapes, and another consisted of Appliances, Furniture, Food, and Utensils. Each set was then divided in half, such that each half contained three exemplars from each category, yielding six 12-item subsets, a_1, a_2, b_1, b_2, c_1, and c_2.

Items were presented in three large blocks, manipulating the delay between the first and second half of each set (e.g. between a_1 and a_2,). Thus, subsets a_1 and b_1 might be presented in structural slots A_1 and B_1 of the first block, a_2 and c_1 in slots A_2 and C_1 of the second block, and c_2 and b_2 in slots C_2 and B_2 of the third. Assignment of subsets a_1: c_2 to the structural slots A_1:C_2 was fully counterbalanced across subjects, such that each subset appeared in each block with each other subset an equal number of times (with the exception that a_1 was never matched with a_2, b_1 with b_2, or c_1 with c_2.). As described earlier, A_1-B_1 represents a control contrast, A_2-C_1 gauges the persistence of cumulative semantic interference, and C_2-B_2 gauges its decay.

Procedure

Pictures were first presented in a familiarization phase and then in three large blocks. Approximately one hour intervened between the onset of the first block and that of the second. The familiarization phase roughly followed the structure of Howard et al.'s (2006) experiment, while each test block had a structure closer to that of Navarrete et al.'s (2010) Experiment 1.

The self-paced familiarization phase was identical to that of Experiments 1 and 2. Each picture was presented once in a structured random order, such that one item from each category appeared before the next item from any category appeared. After the participant named a picture (typically with the desired label), its desired label appeared on the screen for 1000ms. Participants were instructed to use this name for the rest of the experiment. Then a fixation point appeared before the next picture appeared. Since each image remained only until the voicekey was triggered, this familiarization protocol also helped train participants to use the voice key in the rest of the experiment (i.e. to speak loudly while avoiding lipsmacks and other non-speech sounds).

Then the experiment proper began. Participants were instructed to name each picture as quickly and accurately as possible without making any extraneous sounds. Each trial began with a fixation point ('+') that appeared in the middle of the screen for 500ms, followed by 500ms of a blank screen. Then a picture appeared in the center of the screen for 2000ms or until the voicekey detected a response. Then the display was cleared for 1000ms before the next trial began.

Pictures in the main experiment were distributed among three blocks, as described earlier. In each block, participants named 24 pictures for six cycles, for a total of 144 trials per block. Within each cycle, the items from the two slots were randomly interleaved as in Experiment 2, such that each pair of trials contained one item from each (e.g. A_1 or B_1), selected randomly without replacement, followed by one item from the other.

The first block of picture naming was followed by a long 'break'. During this time, participants were given multiple opportunities for personal breaks and completed an unrelated phonotactic learning experiment (i.e. Kittredge & Dell, 2011, which used stimuli and procedures similar to those of Warker & Dell, 2006). The stimuli in this secondary experiment were primarily nonwords in a nonword context (e.g. "hes fek neg kem"), and more importantly were held constant for all participants, so the secondary experiment should not have affected the contrasts of interest for the picture naming experiment.

The second session of the picture naming experiment, which included the second and third blocks, began approximately one hour after the first (mean: 62.7 minutes, range: 55.8-81.1 minutes). Participants were instructed that they would be naming more pictures, and reminded that they should name them as quickly and accurately as possible, without making any extraneous sounds. Until the onset of the second session, participants were given no indication that they would return to the picture naming task, as opposed to participating in a third unrelated experiment; they were only told that we had combined several experiments, and the whole package would take approximately 90 minutes.

Equipment

Experiment presentation was controlled via E-Prime v1.1. Responses were digitally recorded via a headmounted microphone, and transcribed offline. Since an unacceptably high proportion of responses were lost to voicekey malfunctions (e.g. lipsmacks, equipment errors) in Experiment 2, the onset of each picture display was now accompanied by an inaudible tone (as described by Jansen & Watter, 2008), recorded directly from the presentation computer to the audio recorder. Though naming latencies were still collected via voicekey activation through an E-Prime serial response box, the addition of the inaudible tone now allowed misidentified naming latencies to be detected and manually corrected offline.

Data analysis

Participants' utterances were digitally recorded, with the accuracy of responses and naming latencies verified offline. Naming latencies hand-corrected via visual waveform analysis in the case of 760 lipsmacks and other noises that triggered the voicekey too early, thus allowing those trials to be included in the naming latency analyses.

Cohort similarity ratings. Indices of cohort similarity were derived from subjective ratings collected by Catherine Hodgson (personal communication, 1/19/2007). Sixteen University of Manchester students were presented written pairs of words from Schnur et al.'s (2006) stimuli, including pairs from each semantic category[20] (e.g. horse-skunk). They were asked to rate how semantically related they thought each pair was, on a scale from 1 (least alike) to 5 (most alike). With this data, I calculated a Cohort similarity predictor as the mean pairwise relatedness rating for each item (e.g. horse) with the other two same-category items in its subset (e.g. horse-skunk, horse-dog). The resulting predictor is distributed with a mean of 3.22 and a range from 1.38 to 4.41. Note that while these subjective ratings likely reflect semantic associations as well as feature-based and essential similarity, there is reason to expect that they should act similarly with regard to cumulative semantic interference (e.g. Abdel Rahman & Melinger, 2007).

---

[20] Note that ratings were not collected for all 2556 possible word pairs. Thus they could not be used productively in Experiments 1 and 2, and do not consider possible interference from items in related categories (e.g. Alario & Moscoso Del Prado, 2010; Vigliocco et al., 2002).

Speech errors. Inaccurate responses were classified as in Experiments 1 and 2, except that removing the Homogeneous/Heterogeneous distinction obviated the need to consider non-semantic within-block errors. Therefore statistical consideration was limited to semantic errors alone, but otherwise used the same general logistic regression methods as in Experiments 1 and 2.

Naming latencies. Naming latency analyses followed strategies described in Baayen (2008) and Baayen and Milin (2010), since more closely related experiments (Alario & Moscoso Del Prado, 2010; Howard et al., 2006; Navarrete et al., 2010) have not converged on a particular analytical approach. As a first step, I minimally trimmed the data by removing erroneous responses and those with naming latencies that would be a priori implausible for legitimate responses. Since I had hand-verified all latencies, I only removed those where a participant began to speak after the picture had disappeared from the screen (i.e. >2000ms). Second, I transformed the data to approximate a normal distribution. Quantile-quantile plots suggested an inverse transform (1/RT), and each transformed naming latency was then multiplied by -1000 so that the model coefficients would a.) have the expected sign, and b.) not be too small.

Next, an initial linear mixed effects model was fitted to the data via lmer (Bates & Maechler, 2010). As in Experiments 1 and 2, this included both structural predictors (Block, Cycle, Position within cycle, Previous trial RT) and theoretically motivated predictors (Category-specific position within cycle, Cohort relatedness, and most importantly, the Control, Persistence, and Decay contrasts) as well as random intercepts for subjects and items. Block and Cycle were coded as centered Helmert contrasts. The Control, Persistence, and Decay contrasts were each coded as centered simple main effects. All other predictors were coded as centered linear or binary contrasts. Drawing on expectations form Experiments 1 and 2, I also included interactions of Category-specific position within cycle (i.e. the incremental interference effect) with Cycle and Cohort relatedness. Then 475 (2.32%) trials with absolute standardized residuals greater than 2.5 were excluded as influential points. Then, the model was refitted to the trimmed data, including forward-selected within-subject and -item random slopes (via ffRaneff.fnc from Tremblay, 2011, as in Experiments 1 & 2).

Regression tables, including p-values, were estimated via Markov Chain Monte Carlo simulations with 10,000 replications (pvals.fnc from Baayen, 2011). To ease interpretation, I also list back-transformed millisecond estimates of the effect of each predictor at the experiment grand mean, and use these when reporting effects in the text.

Results and discussion

Errors

Participants named pictures correctly in the vast majority of trials (20,426 trials, or 98.5%). The 310 trials (1.5%) that failed to elicit correct responses are summarized in Table 14. Analyses of the 161 semantic errors (Table 15) show that they decreased across cycles (main effect of Cycle: $OR=0.81$, $p<.001$) and increased within cycles as function of the trials within-category ordinal position (main effect of Position in cycle (category-specific): $OR=1.24$, $p=.036$). Both findings are consistent with Navarrete et al.'s (2010, Experiment 1) all-error analysis, and the ordinal position effect also agrees with nonsignificant trends in Howard et al.'s (2006) all-error analysis. Semantic errors were also more likely on trials where the target was more semantically similar to the same-category competitors in its block (main effect of Cohort similarity: OR=2.87, p<.001). This is a novel finding – Vigliocco et al. (2002) and Alario and Moscoso del Prado (2010) reported graded effects of semantic distance on naming latencies, but not errors – but is generally consistent with reports that blocked-cyclic naming errors are more likely in Homogeneous blocks than Heterogeneous (e.g. Belke, Brysbaert, et al., 2005; Belke, Meyer, et al., 2005; Ganushchak & Schiller, 2008; Jefferies et al., 2007; Schnur et al., 2006; Wilshire & McCarthy, 2002). However, there were no obvious error manifestations of either the Persistence or Decay contrasts.

Table 14. Error data from Experiment 3.

		Session1		Session 2			
		Block 1		Block 2		Block 3	
Outcome	Example	A_1	B_1	A_2	C_1	C_2	B_2
Correct	*dog↦dog*	3406	3399	3409	3396	3415	3400
Semantic error		36	27	23	26	22	27
Completed	*dog↦cat*	21	18	14	13	5	17
Interrupted	*dog↦ca...*	15	9	9	13	17	10
Omission	*dog↦...*	2	2	3	1	1	0
Other error		11	29	21	32	17	28
Perspective	*dog↦animal*	2	22	17	15	9	11
Phonological	*dog↦tog*	1	2	1	3	1	3
Disfluency	*dog↦d... uh... dog*	6	1	1	5	1	4
Other or uncodable	*dog↦ti... @#%!*	2	4	2	9	6	10
Equipment failure	*[static]*	1	0	0	0	0	1

Table 15. Logistic regression summary for Experiment 3 semantic errors (sum=161).

	Coef β	SE(β)	p	OR (exp(β))
(Intercept)	−6.01	0.22	**<.001**	0.00
Block	−0.11	0.10	0.27	0.89
Cycle	−0.21	0.05	**<.001**	0.81
Cohort similarity	1.05	0.29	**<.001**	2.87
Position in cycle (category-specific)	0.21	0.10	**0.036**	1.24
Control contrast	0.29	0.27	0.29	1.33
Persistence contrast	−0.06	0.29	0.83	0.94
Decay contrast	0.18	0.31	0.57	1.19

Naming latencies

Excluding errors and outliers left 19,935 (96.1%) observations for the naming latency analysis, summarized in Table 16. Results of the regression analysis are summarized in Table 17. As in Experiments 1 and 2, naming latencies showed significant autocorrelation (main effect of -1000/(Previous trial RT): Estimate= 54.63ms, p<.001). Responses generally slowed over the course of the experiment, such that mean naming latencies significantly increased after the both the first block and the second (First block Helmert contrast: Estimate= 18.49ms, p<.001; Second block Helmert contrast: Estimate= 8.64ms, p<.001). Such slowing is predicted if cumulative semantic interference spills over from one block to the next, though it could also reflect more generalized fatigue from the long experiment. Responses nonetheless showed substantial repetition priming, with naming latencies significantly decreasing after the first, second, and third cycles (First cycle Helmert contrast: Estimate= -55.81ms, p<.001; Second cycle Helmert contrast: Estimate= -7.82ms, p=.003; Third cycle Helmert contrast: Estimate= -5.12, p=.002)[21]. These repetition effects are a bit more gradual than those of Experiments 1 and 2, which could be attributed to the influence of the larger cycles (e.g. by disrupting strategic preparation or semantic priming). Finally, like the slowing across blocks, significant slowing within each cycle (main effect of Position in cycle (not category-specific): Estimate= 0.47ms, p<.001), could reflect incremental interference from weak semantic relations (e.g. Alario & Moscoso Del Prado, 2010; cf. Vigliocco et al., 2002), but could also reflect mere fatigue. Next we turn to more obvious manifestations of cumulative semantic interference, first considering its proximal effects, and then its more persistent consequences.

Table 16. Mean naming latencies for Experiment 3 (untransformed), estimated as in Table 6 and Table 11.

	Session 1									Session 2																	
	Block 1 (Control)									Block 2 (Persistence)									Block 3 (Decay)								
	A_1			B_1			Diff. $(A_1 - B_1)$			A_2 (60-min lag)			C_1			Diff. $(A_2 - C_1)$			C_2 (0-min lag)			B_2 (67-min lag)			Diff. $(C_2 - B_2)$		
Cycle	RT	[95% HPD]		RT	[95% HPD]		Est.	[95% HPD]		RT	[95% HPD]		RT	[95% HPD]		Est.	[95% HPD]		RT	[95% HPD]		RT	[95% HPD]		Est.	[95% HPD]	
1	682.4	662.7	703.4	688.5	669.0	710.1	-6.2	-19.9	7.3	716.1	695.0	736.2	707.5	686.6	727.4	8.5	-5.6	22.5	747.4	727.2	768.3	752.9	732.4	773.2	-5.4	-18.8	8.5
2	626.0	605.9	646.7	634.2	614.1	655.0	-8.3	-21.5	6.0	646.6	626.0	666.9	635.9	614.6	655.4	10.7	-3.7	24.0	659.7	638.6	679.5	666.6	646.5	687.0	-6.8	-19.8	7.5
3	621.9	601.7	642.6	627.1	606.0	646.6	-5.1	-18.8	8.3	645.7	625.1	665.9	638.6	618.1	659.2	7.2	-6.4	20.6	650.5	630.0	670.7	655.8	636.0	677.0	-5.3	-18.2	9.0
4	617.6	597.1	638.2	625.6	604.3	645.5	-8.0	-22.1	5.6	644.6	624.7	665.7	628.7	608.3	649.0	15.9	2.0	29.1	638.6	618.2	658.9	645.0	624.2	665.1	-6.4	-20.4	6.8
5	611.2	590.6	631.9	610.2	589.9	630.5	1.0	-12.8	14.8	638.6	618.0	658.4	636.5	616.1	657.0	2.0	-11.6	15.8	640.2	620.3	661.1	648.1	627.3	667.8	-7.9	-22.2	5.0
6	618.4	598.2	638.7	625.4	604.6	645.4	-6.7	-20.3	7.0	643.2	621.9	663.0	631.8	611.3	652.0	11.5	-2.3	25.2	650.1	630.1	670.5	657.0	636.5	677.4	-6.9	-20.6	6.9

[21] An increase in naming latencies after the fifth cycle violates my guiding assumption that changes should be continuous and monotonic, and seems likely an accident of multiple comparisons, but could be ascribed to either fatigue or a continued accumulation of semantic interference that eventually outstrips the repetition priming.

Table 17. Summary of the linear mixed effects regression for Experiment 3.

	Estimate(*1000)	[95% HPD Int.]		p	In ms
(Intercept)	−1.5995	−1.6351	−1.5622	**<.001**	625.20
-1000/(Previous trial RT)	0.1395	0.1196	0.1593	**<.001**	54.63
Block (1:3, Helmert)					
$\Delta Block_1 \rightarrow Block_{>1}$	0.0473	0.0264	0.0686	**<.001**	18.49
$\Delta Block_2 \rightarrow Block_{>2}$	0.0221	0.0062	0.0377	**0.009**	8.64
Cycle (1:6, Helmert)					
$\Delta Cycle_1 \rightarrow Cycle_{>1}$	−0.1425	−0.1678	−0.1178	**<.001**	−55.81
$\Delta Cycle_2 \rightarrow Cycle_{>2}$	−0.0200	−0.0333	−0.0067	**0.003**	−7.82
$\Delta Cycle_3 \rightarrow Cycle_{>3}$	−0.0131	−0.0211	−0.0044	**0.002**	−5.12
$\Delta Cycle_4 \rightarrow Cycle_{>4}$	−0.0001	−0.0086	0.0091	0.99	−0.04
$\Delta Cycle_5 \rightarrow Cycle_{>5}$	0.0130	0.0025	0.0233	**0.013**	5.08
Position in cycle (1:24)	0.0012	0.0005	0.0019	**<.001**	0.47
Position in cycle (category-specific; 1:3)	0.0160	0.0096	0.0224	**<.001**	6.25
× Cycle					
$\Delta Cycle_1 \rightarrow Cycle_{>1}$	−0.0352	−0.0448	−0.0247	**<.001**	−13.76
$\Delta Cycle_2 \rightarrow Cycle_{>2}$	0.0187	0.0088	0.0287	**<.001**	7.31
$\Delta Cycle_3 \rightarrow Cycle_{>3}$	−0.0103	−0.0208	−0.0004	**0.044**	−4.03
$\Delta Cycle_4 \rightarrow Cycle_{>4}$	−0.0011	−0.0123	0.0100	0.85	−0.43
$\Delta Cycle_5 \rightarrow Cycle_{>5}$	0.0001	−0.0125	0.0126	0.99	0.04
Cohort similarity (1-5)	0.0254	−0.0047	0.0561	**0.098**	9.93
× Position in cycle (cat.-spec.)	0.0106	0.0040	0.0172	**0.002**	4.14
× Cycle					
$\Delta Cycle_1 \rightarrow Cycle_{>1}$	0.0261	0.0012	0.0510	**0.041**	10.20
$\Delta Cycle_2 \rightarrow Cycle_{>2}$	0.0089	−0.0038	0.0222	0.18	3.48
$\Delta Cycle_3 \rightarrow Cycle_{>3}$	0.0068	−0.0055	0.0178	0.27	2.66
$\Delta Cycle_4 \rightarrow Cycle_{>4}$	0.0045	−0.0076	0.0172	0.47	1.76
$\Delta Cycle_5 \rightarrow Cycle_{>5}$	0.0086	−0.0054	0.0231	0.25	3.36
Control contrast ($A_1 - B_1$)	−0.0091	−0.0266	0.0076	0.28	−3.56
Persistence contrast ($A_2 - C_1$)	0.0196	0.0009	0.0383	**0.04**	7.66
Decay contrast ($C_2 - B_2$)	−0.0095	−0.0268	0.0085	0.28	−3.71

Proximal effects of cumulative semantic interference

Before considering whether cumulative semantic interference persists over long lags, we need to establish that the response times show the proximal manifestations of cumulative semantic interference that we have come to expect: the incremental interference effect and the semantic blocking effect. In fact, naming latencies do show the expected incremental interference effect across all cycles (Figure 22a; main effect of Position in cycle (category-specific): Estimate=6.25ms, p<.001).[22] This effect was strongest in the first cycle of naming (main effect of Position in cycle (category-specific) for a regression restricted to data from the first cycle: Estimate=20.64ms, p<.001), and though it significantly dropped thereafter (First cycle Helmert contrast X Position in cycle (category-specific): Estimate= -13.76ms, p<.001), it nonetheless remained significant in the subsequent five cycles (main effect of Position in cycle (category-specific) for a regression restricted to data from after the first cycle: Estimate=4.02ms, p=.004). The ~21ms estimated magnitude of the incremental interference effect in the first cycle in this experiment is well within the range that has been suggested for such effects in non-cyclic picture naming experiments (Alario & Moscoso Del Prado, 2010; Howard et al., 2006; Nickels et al., 2008), and is similar to the ~18ms[23] effect from the first cycle of Navarrete et al.'s (2011, Experiment 1) study. Remarkably, the estimate is also entirely consistent with the ~23ms and ~19ms estimates from Experiments 1 and 2, respectively, supporting

[22] Since this predictor is somewhat collinear with the non-category-specific Position predictor, it is worth noting that its addition improves model fit even after including the non-specific predictor, and the coefficient remains significant and approximately the same size if regressing the category-specific predictor against the non-specific predictor prior to inclusion.

[23] This 18ms estimate is derived from the Cycle X Ordinal position means listed in Navarrete et al.'s Table 1. Since they did not find a significant Cycle X Ordinal position interaction, they only reported statistical estimates for ordinal position effects as aggregated across all four cycles.

previous suggestions that the incremental interference effect is undiminished in the first cycle of blocked-cyclic naming experiments, and thus that the typical absence of a semantic blocking effect from that first cycle must have another source. It is also notable that the incremental interference effect in the current experiment continued to accumulate after the first cycle, though not as strongly as in the first, thus providing the missing link between Experiments 1 and 2 – where there were no indications of incremental interference effects after the first cycle – and Navarrete et al.'s Experiment 1 – which did not detect a significant decrement to the incremental interference effect over four cycles.

Participants were also marginally slower to name pictures that were more semantically related to other pictures in the same block (Figure 22b; main effect of Cohort similarity: Estimate=9.93ms, p=.098). Like the semantic blocking effect, this Cohort similarity effect grew across cycles (First cycle Helmert contrast X Cohort similarity: Estimate= 10.20ms, p=.041)[24]. Greater Cohort similarity was associated with stronger incremental interference effects (First cycle Helmert contrast

---

[24] Helmert contrasts suggest monotonic growth across all cycles, Curiously, Cohort similarity has little effect in the first cycle, but this seems to reflect factors distinct from those that led to a null or reverse blocking effect in the first cycle of blocked-cyclic naming. Here, the issue seems to be that the incremental interference effect for high-Cohort similarity items in the first cycle is not much greater than that for low-Cohort similarity items (Figure 22), and also appears to develop over successive blocks.

X Cohort similarity: Estimate= 4.14ms, p=.002). This echoes previous claims that both the semantic blocking effect (Vigliocco et al., 2002) and the incremental interference effect (Alario & Moscoso Del Prado, 2010) reflect gradations of semantic relations, as opposed to binary category membership, and demonstrates a new homologue of the semantic blocking effect that is typically seen in blocked-cyclic naming.

a.) b.)

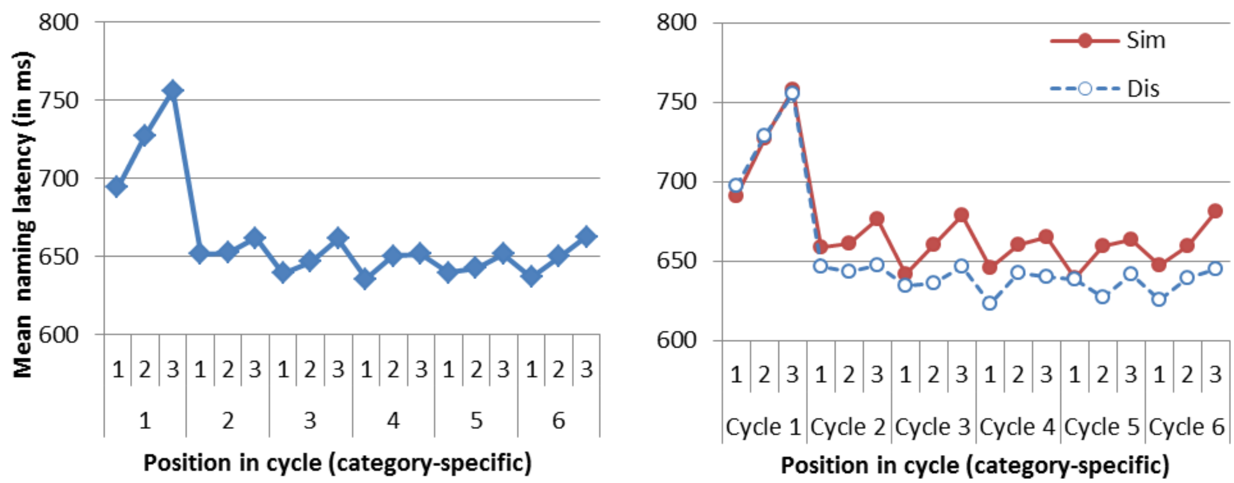

Figure 22. Incremental interference effects for Experiment 3, nested within cycle. a.) Overall incremental interference effects. b.) Incremental interference effects, by Cohort similarity. For this illustration, Cohort similarity is discretized into upper and lower 50th percentiles.

Persistent effects of cumulative semantic interference

To evaluate whether semantic interference persists over the long delay, and how much it might decay, the structure of this experiment provides three orthogonal contrasts, each comparing the mean naming latencies for two sets of pictures that were interleaved within a block (Figure 23a).

For comparison, I also present two simulations of the experiment, using the model described in Chapter 2 and the parameters specified in Table 1. The first simulation, depicted in Figure 23b, shows what we would expect if weight changes in the experiment were truly persistent with no decay. The second simulation, depicted in Figure 23c, shows what we would expect if weight changes in the experiment were subject to moderate time-based decay, specifically assuming that weight changes from a block decay with half-life of 15minutes after the block has ended. Clearly, the simulation that omits the decay function (i.e. Figure 23b) better matches the empirical results.

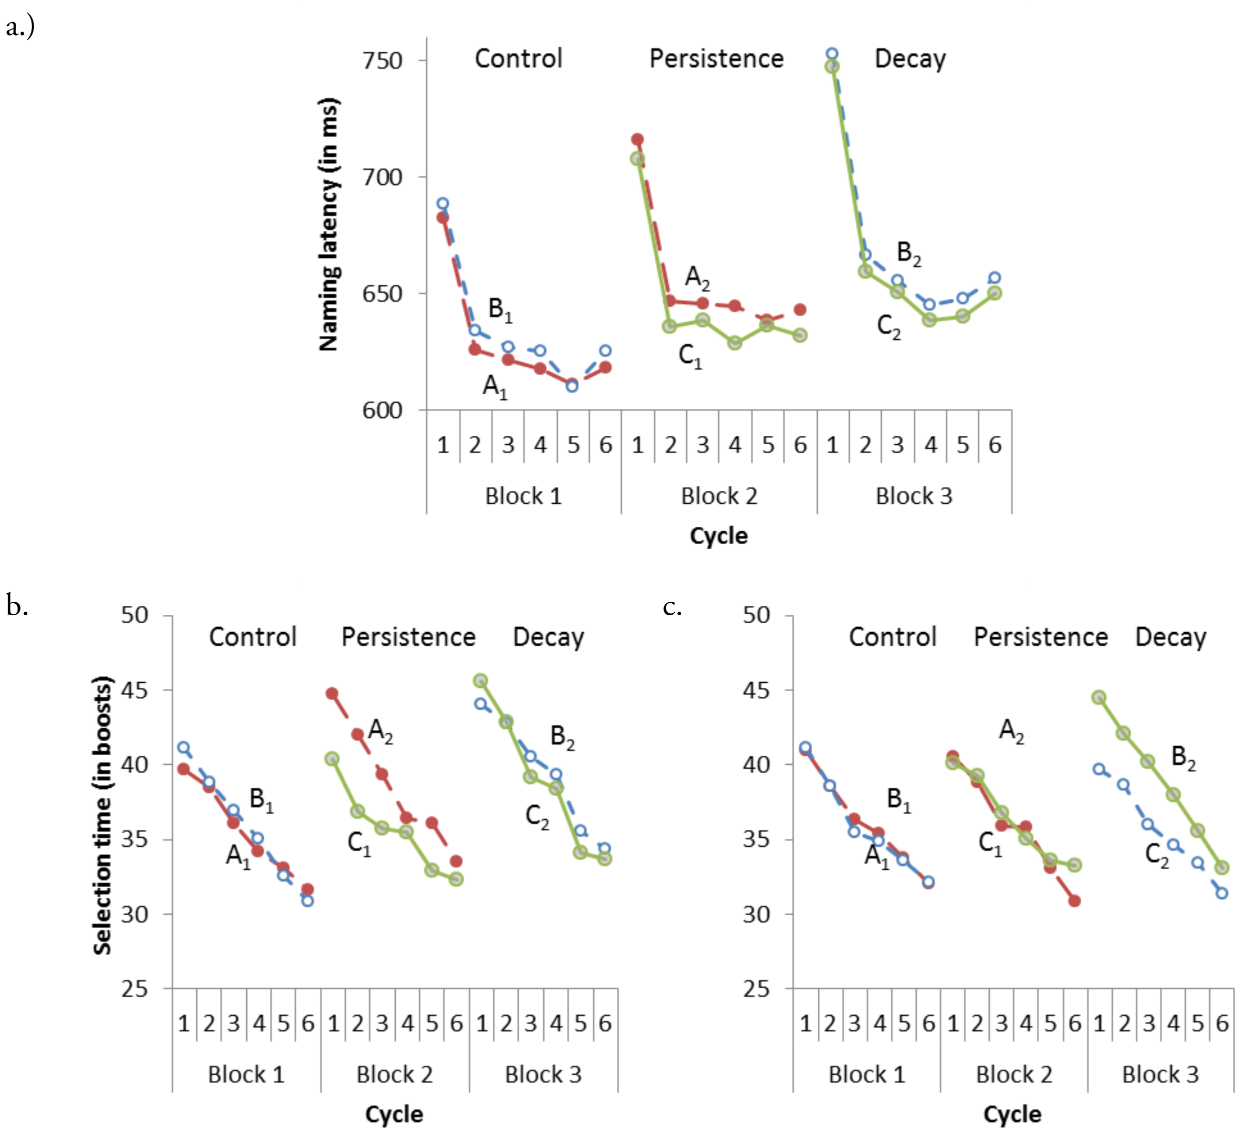

Figure 23. Empirical naming latencies (Panel a) and simulated selection times (Panels b and c) for Control, Persistence, and Decay contrasts, from Experiment 3. a.) A graphical depiction of the empirical naming latency data from Table 16. b.) A simulation of the experiment, assuming fully persistent connection changes with no decay. c.) A second simulation, assuming that weight changes from within a block decay with a 15-minute half-life after the block has ended. Both simulations use the model described in Chapter 2, with

the parameters specified in Table 1. Mean selection latencies for each are based on 48-replication runs of the full simulated experiment.

Control contrast

The first contrast serves as a control, comparing naming latencies for two interleaved slots in Block 1: A_1 and B_1. Since the sets are, at this point, balanced for the number of previous trials in which same-category (i.e. competing) pictures have been named, naming latencies for these sets should not significantly differ. Any difference in these naming latencies can also serve as an estimate of the amount of noise in the data. As expected, statistical analysis of this contrast indicates that naming latencies in this first block do not significantly differ (Control contrast (A_1-B_1): Estimate: -3.56ms, p=.2 8).

Persistence contrast

The second contrast considers the persistence of semantic interference by comparing naming latencies for two interleaved slots in Block 2: A_2 and C_1. A_2 items in this block represented novel exemplars from some of the same semantic categories accessed an hour before, in Block 1. C_1 items, by contrast, represented novel exemplars from categories that had not been accessed since the familiarization phase of the experiment, when each picture was named just once. Thus, to the extent that semantic interference persists and generalizes from A_1 items in Block 1 to novel same-category items in Block 2, an hour later, the A_2 items in Block 2 should be named more slowly than the C_1 items. Block 2 analyses confirm this prediction, estimating that the A_2 items are named 7.66ms slower than the C_1 (Persistence contrast (A_2-C_1):p=.040). As in Experiments 1 and 2, the effect of persistent semantic interference shown in this block is not large, especially when one considers the 20-ms incremental interference effects in the first cycle, but it does represent a significant naming latency difference in the predicted direction, statistically confirming the pattern suggested by the previous experiments and indicating that whatever phenomenon produces semantic interference in picture naming after a few seconds continues to affect picture naming much later. And the fact that the difference comes out here – without an obvious blocking manipulation – suggests it reflects relatively implicit aspects of word retrieval.

Decay contrast

If the eight-millisecond difference for the persistence contrast in Block 2 is smaller than might be expected for a fully persistent effect, then it would be tempting to assume that semantic interference had decayed in the time that passed between Block 1 and Block 2. The third contrast tests this possibility by comparing naming latencies for two interleaved slots in Block 3: B_2 and C_2. B_2 items in this block represented novel exemplars from the same semantic categories accessed in the B_1 slot of Block 1, an hour before, so the semantic interference for these items has had an hour to decay. C_2 items, by contrast, represented novel exemplars from categories accessed immediately prior, in the C_1 slot of Block 2, and should therefore show little if any decay. Thus, to the extent that semantic interference decayed in the hour between Block 1 and Block 3, the B_2 items in Block 3 should be named more quickly than the C_2 items. Surprisingly, the Block 3 data offer no evidence for such decay: if anything, the data show a trend in the opposite direction, with the B_2 items named slightly *slower* than the C_2 (Control contrast (C_2-B_2): Estimate=-3.71ms, p=.28). The difference for the Decay contrast is similar in magnitude to that estimated for the Block 1 Control contrast, so the most straightforward interpretation is simply that of a null effect: that cumulative semantic interference did not decay at all in the one-hour delay distinguishing B_2 from C_2 items.

General discussion

Cumulative semantic interference: persistence and decay

The most important finding to emerge from these experiments is that cumulative semantic interference lasts long enough to fit the kind of time course that we might expect of a learning-based phenomenon. Previous work had only shown that it lasted at least a few seconds or maybe up to a few minutes – longer than should be expected for a phenomenon based in residual activation, but still quite brief for a phenomenon that purportedly reflects the persistent restructuring of a person's lexical knowledge. Here we have the first evidence that consequences of cumulative semantic interference can remain detectable at least an hour later. This represents a major step in connecting millisecond-scale laboratory observations to the processes that speakers normally use to continually optimize their vocabularies throughout their lives.

This persistence was first suggested in two blocked-cyclic naming experiments and finally demonstrated in a third picture naming experiment that lent itself far less to strategic preparations. The demonstration in that third experiment is particularly important because it seems to preclude the possibility that the observed persistence might derive from the sort of explicit category-name associations (e.g. FRUIT:A___ --> APRICOT) that are typically studied in retrieval-induced forgetting paradigms and which could be claimed to underlie persistent interference in situations with more obvious blocking manipulations (e.g. in blocked-cyclic naming, interference when naming mammals in the second session could reflect memories of a cluster of mammals in the first session occluding access to the new mammals in the second). In other words, it suggests that the persisting interference came from processing individual semantically related items in the previous session, rather than from explicit considerations of a category that contains the items or from the success or failure of other task-specific strategies.

In fact, Experiment 3 may have succeeded in detecting persistent interference effects specifically because it limited the opportunity for participants to adopt strategies that would allow them to prepare their responses in advance of stimulus presentation. In Experiment 1, for instance naming latencies dramatically decreased within each cycle in both Homogeneous and Heterogeneous conditions, suggesting that participants may have been able to begin the naming process before a stimulus actually appeared (cf Griffin & Bock, 1998). Such preparation could be problematic for response time studies, which depend on time-locking a response to the onset of stimulus presentation. Comparable semantic blocking effects in Experiments 1 and 2 suggest that strategic preparation may not noticeably obscure the phenomenon (cf e.g. Roelofs, 2010), but preparation may also have the potential to reduce learning-based effects by reducing error. In the Dark Side model, for instance, lexical learning is driven by error in lexical activations. If strategy or other cues were able to reduce activation error during word retrieval (e.g. via input from a bias node or sequential prediction mechanism during blocked-cyclic naming), then the accumulation of semantic interference would also be attenuated. This is akin to an assumption that people do not learn much from very easy tasks. Such attenuation of error during lexical retrieval could explain the apparent lack of incremental interference effects after the first cycle in Experiments 1 and 2, on the assumption that repeatedly naming a small set of pictures quickly became easy enough that there was little left to learn. It could also explain why the blocking effect often appears quite stable after the first cycle in blocked-cyclic naming (e.g. Belke, Meyer, et al., 2005, Experiment 1): repeatedly naming the same items generates little error to drive error-based learning. Along the same lines, it is possible that the persistence of cumulative semantic interference finally came through in Experiment 3 because the larger cycles made naming more difficult, yielding an incremental interference effect that continued to build within all six cycles[25]; with more interference building within a block, there would be more interference left over an hour later.

So does cumulative semantic interference decay at all? A direct test in Experiment 3 revealed no indication that it did decay over the course of a one-hour delay (mirroring the simulation, in Figure 23b, which omitted a decay function), but this leaves us to wrestle with the fact that the persistent effects observed in the three experiments – 6ms, 4ms, and 8ms, respectively – are all quite

---

[25] Considering the three experiments, plus Navarrete et al.'s (2010) Experiment 1, the slope of the incremental interference effect after the first cycle seems to approximate a logarithmic function of the size of the cycle. However, with only four data points it might be overly optimistic to expect that this relationship would extend to other studies.

a bit smaller than the 20ms incremental interference effects seen in just the first cycle of each. How might we resolve this puzzle? One possibility is that the semantic interference that comes from incremental learning in accessing words is partly counteracted by semantic facilitation, e.g. from incremental learning in accessing meaning (Becker, Moscovitch, Behrmann, & Joordens, 1997) which is perhaps overwritten more easily than that at the lexical level (cf Wheeldon & Monsell's, 1994, notion of a short-lived semantic facilitation that gives way to longer-lasting semantic interference). If that were the case, semantic-level facilitation would lead us to underestimate the short-lag interference in the Experiment 3 Decay contrast (thus missing a reduction of semantic interference at the longer lag), but note that the long-lag interference in the persistence contrasts of all three experiments would likely be underestimated in the same way. Another possibility is that cumulative semantic interference does not decay with respect to time or unrelated experience, but much of the interference that accumulates within a specific episode (e.g. within a block of picture naming or a conversation with a particular interlocutor) is in some way bound to that particular context, and switching to a new context is sufficient to disengage that portion of the interference. This is related to the idea that, in a phonotactic learning experiment (e.g. Warker & Dell, 2006) people may acquire and retain constraints for an artificial language without necessarily integrating that learning with their normal language use. Completing Blocks 1 and 2 might then provide an equivalent release from cumulative semantic interference (i.e. language learning specific to the current block) without additional time-based decay.

Bridging the gap between incremental interference and the semantic blocking effect

The three experiments presented here also link the trial-by-trial incremental interference effect to the cycle-level semantic blocking effect. Experiments 1 and 2 demonstrated, for the first time, the presence of incremental interference effects in blocked-cyclic naming paradigms. Remarkably, the incremental interference effect was present in the first cycle of each, with magnitudes comparable to estimates from non-blocked procedures, both in previously published studies and for the same items in Experiment 3. This demonstration of a robust incremental interference effect in the first cycle of blocked-cyclic naming is important because it shows that, contrary to common characterizations of the semantic blocking effect, cumulative semantic interference is already present in the first cycle of naming. It does not require a gestalt semantic context to be established, and is not temporarily attenuated by semantic facilitation or post-selection inhibition. Whatever influences typically lead to an absent or reversed semantic blocking effect in the initial cycle(s) of blocked-cyclic naming are already present on the first trial of a block, before Homogeneous and Heterogeneous contexts have yet been differentiated (mean naming latencies for first trial of each block: Experiment 1: Homogeneous=756ms, Heterogeneous=794ms; Experiment 2: Homogeneous=681ms, Heterogeneous=699ms). While such a pattern defies characterization as a semantic *context* effect, it is easily resolved by recognizing that semantic interference persists across blocks, meaning that it affects items named in Heterogeneous blocks as well as those named in Homogeneous. For Heterogeneous items, semantic interference has already accumulated before the block has begun, and it simply takes a cycle or two for interference in the Homogeneous blocks to catch up.[26] Thus acknowledging the persistence of cumulative semantic interference solves a puzzle of how semantic facilitation turns into semantic interference by proposing that the apparent facilitation was in fact interference all along.

While Experiments 1 and 2 showed incremental interference effects in blocked-cyclic naming, Experiment 3 showed a homologue of the semantic blocking effect without including a semantic blocking manipulation. That is, it showed that naming latencies were slower, with steeper incremental interference effects, for items that were judged as more closely related to other items in the block. This finding recalls demonstrations of graded similarity effects in blocked-cyclic naming (Vigliocco et al., 2002) and hierarchical category relations in incremental interference effects (Alario & Moscoso Del Prado, 2010), but further shows that semantically graded incremental interference effects develop over cycles much like the semantic blocking effect. Thus we see that echoes of the semantic blocking effect can be seen without semantic blocking, and incremental interference effects can be found with semantic blocking, and that the major point of discrepancy between incremental interference effects and semantic blocking effects – the lack of a semantic blocking effect in the first cycle of blocked-cyclic naming – can in fact be explained as a consequence of persistent cumulative semantic interference.

[26] Curiously, the potential for such reversal is one reason that Damian and Als (2005) cited in claiming that cumulative semantic interference could not last longer than one block.

Conclusion

Cumulative semantic interference reflects persistent changes to the connections that people use to map from meanings to words. It manifests as both incremental interference effects and semantic blocking effects, and neither manifestation is restricted to a particular experimental paradigm.

CHAPTER 7: CONCLUSION

The average adult speaks 16,000 words per day (Mehl, Vazire, Ramírez-Esparza, Slatcher, & Pennebaker, 2007), retrieving each from an active vocabulary of 40,000 words at a rate of 2-3 words per second (Levelt, 1989). How is this possible? Part of the answer is that speakers benefit from an incredible amount of practice, and they use the recent past to predict the future. The Dark Side model starts with the idea that each act of word retrieval is also an act of word learning, and explores how such learning could account for an array of behavioral findings. It is, in fact, the first model of lexical retrieval to deal with both error and response time data in both healthy and impaired speakers.

Chapter 1 introduced the phenomenon of *cumulative semantic interference*, which chiefly manifests as slower and more error-prone word retrievals following the retrieval of semantically related words. This produces an incremental interference effect in the continuous paradigm, where words from many intermingled categories are retrieved just once, and it produces a semantic blocking effect in the blocked-cyclic naming paradigm, where words from the same category are retrieved repeatedly. The interference is robust to the inclusion of delays and filler-material, it generalizes to novel exemplars from the same category, and it manifests in semantic errors, omissions, and increased naming latencies in both healthy and impaired speakers.

A model presented in Chapter 2 posits that the process of lexical retrieval can be represented as activation, a winner-take-all selection mechanism, and an incremental learning algorithm, situated within a simple connectionist framework. Words are first activated via constellations of distributed semantic features. Many words are activated at this stage, and it is difficult, but necessary, to select a winner. So the network floods the lexical level with extra activation, repeatedly boosting each words activation until a clear winner emerges and is selected for production. Then the network tweaks its semantic-to-lexical mappings via an error-driven learning algorithm, strengthening connections from the activated semantic features to the desired response (typically the word that was just produced) and weakening connections from those features to all other words, in proportion to the erroneous activation. This framework is justified as a simple implementation of a few basic principles that are (nearly) universally accepted in the field.

Chapter 3 demonstrates that, by training the model on a small vocabulary and simulating the structure of several representative experiments, the model produces behaviors that are qualitatively similar to those reported for humans. Naming times increase with ordinal position in a continuous paradigm task, and as in humans this increase is based on the number of occurrences of semantically related items rather than the number of unique types. When naming items in a blocked-cyclic naming paradigm, the generates the expected naming latency manifestations of the semantic blocking effect, and adding noise to make selection less deterministic generates the semantic error and omission manifestations associated with patients with aphasia, including an item-lag-based recency effect for semantic errors. And as in humans, the semantic blocking effect generalizes to novel exemplars from the same semantic category.

Chapter 4 explores how the model accounts for cumulative semantic interference effects. First, it shows that the models account derives primarily from the competitive unlearning of undesired semantic-to-lexical associations. Surprisingly, facilitation of previous responses does little if anything to produce the expected interference effect and this leads to a realization that – contrary to years of published claims and even our own guiding assumptions – cumulative semantic interference does not require competition at the time of lexical selection (occlusion), because competition is already built into the learning process.

Chapter 5 reviews implications and predictions from the modeling work. One of the most important predictions, is the idea that if cumulative semantic interference reflects an incremental learning process, then it should adopt the timecourse of incremental learning, meaning that it should remain in effect for much longer than was currently supposed. Chapter 6 supported this prediction in three new experiments with human subjects, which showed that cumulative semantic interference generalizes to new exemplars from the same category after an hour delay, with no indication of time-based decay. These experiments also demonstrated, for the first time, the presence of incremental interference effects in blocked-cyclic naming and homologues of semantic blocking effects in a paradigm without semantic blocking, and showed that an apparent discrepancy between the incremental interference effect and the canonical semantic blocking effect may simply reflect the demonstrated persistence of cumulative semantic interference, thus supporting the models characterization of the incremental interference effect and semantic blocking effect as having a common source in the incremental adjustment of semantic-to-lexical mappings.

Future directions

Like the Dark Side model itself, future directions for research can be broadly classified as involving lexical activation, selection, or learning.

Activation. One possible line of future work involves the expected task-specificity of cumulative semantic interference. According to the Dark Side model, cumulative semantic interference should affect behavior in a particular task only to the extent that the task involves a mapping of activation from shared semantic features to divergent lexical responses. So the model may provide a basis for understanding particular production tasks, and findings from well-understood tasks could similarly prove useful for refining the model. For instance, recent studies (Damian, Vigliocco, & Levelt, 2001; Navarrete, Mahon, & Caramazza, 2010; Vitkovitch, Cooper-Pye, & Ali, 2010) in several Indo-European languages (German, Italian and English) have revealed some inconsistency in the extent to which orthographic word-naming (with or without syntactically specified determiners) elicits cumulative semantic interference. This inconsistency in results suggests some inconsistency in the extent to which mapping from orthography to phonology involves a mapping from meaning to phonology. Reading models (Plaut et al., 1996; Seidenberg & McClelland, 1989) typically provide multiple paths for mapping activation from orthography to phonology, and semantic effects have occasionally been reported in single-word naming (Strain, Patterson, & Seidenberg, 1995). So it would be interesting to see whether the factors that have been previously implicated as leading to stronger semantic effects in word naming (e.g. orth→ phon regularity and word frequency, as suggested by Strain et al., 1995, or impairment of the direct orth→ phon pathway) would also produce stronger cumulative semantic interference, possibly explaining some of the discrepancies in previous studies of cumulative semantic interference in word naming.

Selection. If the Dark Side model is correct, the simple fact of cumulative semantic interference does not prove one way or another whether lexical selection is competitive or, more generally, provide information about particular selection thresholds or decision rules. Thus, it leaves open the question of what kind of selection algorithm speakers actually implement. Considering other recent discussions in the literature (e.g. Abdel Rahman & Melinger, 2009a, 2009b; Mahon & Caramazza, 2009; Mahon et al., 2007; Navarrete et al., 2010), and the crucial role of competitive selection in developing the continental model of speech production (i.e. Levelt et al., 1999), it may be important for future work to constrain the possibilities. Beyond the current rhetorical motivations, constraining the possible algorithms for lexical selection would also have value in reducing the degrees of freedom for explanations of other lexical-access-related phenomena. But it will be difficult. In trying to establish whether lexical selection is a competitive process, a major challenge lies in finding a manipulation that should (according to multiple, often conflicting, theories) modulate the accessibility of a competitor without also directly affecting the accessibility of a target. Just as it is important to constrain the possible algorithms for lexical selection, it is essential that any accepted constraints be empirically and computationally supported, rather than accepted as articles of faith, so that other aspects of our theories are not falsely constrained.

The model's booster function opens another line of research by instantiating a role for executive control in lexical retrieval for speech production. That is, it conceptualizes lexical retrieval in speech production as a controlled process, rather than the automatic process suggested by the use of lateral inhibition in Howard et al.'s (2006) model. Previous work (Chang et al., 2006; Dell et al., 2008; Gordon & Dell, 2003) has established the idea that lexical retrieval is typically accomplished through the simultaneous satisfaction of multiple types of constraints (e.g. semantic and syntactic), and the Dark Side model has, to some extent, developed the idea that the imposition of these constraints may involve executive control processes. But any appeal to an executive process risks invoking a deus ex machina with the power to explain (or explain away) any data, so it is essential to better constrain the functions of this process, the mechanisms that it employs, and its abilities and limitations. In the Dark Side model, the core function of the control mechanism underlying the booster process is the conjunction of multiple types of constraints, thereby heuristically facilitating the retrieval of an appropriate word in a particular context. The model seeks to avoid the deus ex machina pitfall by assuming that the control mechanism knows neither the right nor wrong answers, but can employ relatively coarse constraints to bias response selection. Understanding the processes that are normally involved in lexical retrieval, as illustrated by cumulative semantic interference and related phenomena, will ultimately require a better understanding of the cognitive control mechanisms that modulate lexical retrieval and the types of constraints that they employ, and should therefore benefit from continued work with speakers with generalized and language-specific executive dysfunctions.

Learning. The incremental learning aspect of the model provides another focus for future research. For instance, I have argued that the same processes underlying cumulative semantic interference in the laboratory also govern lexical learning in daily life, beyond the lab. But one can also ask to what extent the specific experiment-induced learning is bound to the specific experimental context and to what extent it reflects a less context-dependent restructuring of one's vocabulary. Speakers regularly modulate their lexical access according to the demands of different speaking contexts (e.g. speaking different languages, or constraining and substituting vocabulary when producing infant-directed speech[27]), so if this were the case with cumulative semantic interference, then it opens up the question of how speakers are able to learn and employ modified

---

[27] There is some indication that infant-directed speech involves words that are specifically learned for that social context. For instance, although reduplication figures more prominently in infant-directed English (e.g. 'bye-bye'), the set of reduplicated utterances is relatively constrained and speakers do not appear to use the process productively (e.g. to contrast with 'bye-bye', note the absence of words such as 'hi-hi'), suggesting that these words, at least, are stored as distinct lexical items rather than being produced via the algorithmic modification of a word's phonology. As another example, the word 'potty' survives in infant-directed US English, even though chamber pots (from which the term is derived) have not been commonly used in the area in more than half a century, again suggesting the storage of a lexical item whose use is largely dictated by a specific communicative context.

means of lexical access, and how these context-specific modifications can build up to context-general modifications.

One can also ask about the time course of lexical learning. Several experiments have now shown that cumulative semantic interference lasts for a long time, but we know little about its initial development. The implemented Dark Side model assumes, by default, that incremental learning produces immediate full-strength changes in semantic-to-lexical connections. But it is also plausible that changes could start strong before decaying away (weight decay) or grow over the course of several minutes or seconds (akin to the use of a momentum parameter in some machine learning algorithms). Knowing precisely what to expect from a lexical learning-based phenomenon will therefore require a better understanding of how incremental learning is implemented in the brain and mind more generally.

Another question is how the incremental learning that I have described as the basis for cumulative semantic interference relates to the processes by which speakers acquire wholly novel words. Recent speech comprehension research suggests that novel words may not immediately be integrated into a listener's vocabulary. For instance, when listeners learned the novel word 'cathedruke' in Gaskell and Dumay's (2003, Experiment 2, et passim) study, it took several days before knowledge of this new word hindered access to the known-word 'cathedral'. This finding suggests that, at least in speech perception, new words may not immediately compete for selection. Given the current focus on competitive selection during lexical retrieval for production, it would be interesting to investigate whether cumulative semantic interference would immediately be obtained for retrieving novel words from established semantic categories. For instance, we might ask whether introducing speakers to a new mammal, the 'agouti' (a Central American rodent known for its particularly sharp incisors), would immediately interfere with retrieving the name of a familiar mammal, such as 'otter'. Depending on our assumptions about precisely how speakers retrieve newly learned words we might expect meaning-based production to differ from comprehension, showing that semantic interference both accumulates from and impairs the retrieval of these words.

Finally, the scope of the Dark Side model is quite limited, so future work should integrate the processes described in the model with other processes that must follow and precede lexical access. For instance, models of language production often allow activation to cascade from words to phonemes before a particular word is selected (Dell, 1986; Dell & O'Seaghdha, 1992; Dell et al., 1997; Harley, 1993; Rapp & Goldrick, 2000; Stemberger, 1985), suggesting that the delayed lexical selection that characterizes cumulative semantic interference may have interesting consequences for phonological encoding. On the other end of lexical access, single-word production more typically occurs in multi-word contexts (e.g. Dell et al., 2008), so it will ultimately be important to consider how cumulative semantic interference and incremental learning in lexical access plays out in sentence production. Previous work suggests that semantic interference may drive syntactic choices (Ferreira & Firato, 2002; Smith & Wheeldon, 2004), so it could be interesting to explore ways that speakers may negotiate cumulative semantic interference in their efforts to produce fluent speech.

Conclusion

Speech production requires quickly finding appropriate words to communicate specific ideas in a fluent stream of speech. The Dark Side model instantiates the idea that speakers prepare for this challenge by constantly adapting to their recent past. Thus cumulative semantic interference emerges from the processes that speakers use to continually learn, unlearn, and relearn their linguistic knowledge throughout their lives.

REFERENCES

Abdel Rahman, R., & Melinger, A. (2007). When bees hamper the production of honey: lexical interference from associates in speech production. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*(3), 604-14. doi:10.1037/0278-7393.33.3.604

Abdel Rahman, R., & Melinger, A. (2009). Semantic context effects in language production: A swinging lexical network proposal and a review. *Language and Cognitive Processes*, *24*(5), 713-734.

Abdel Rahman, R., & Melinger, A. (2009b). Dismissing lexical competition does not make speaking any easier: A rejoinder to Mahon and Caramazza (2009). *Language and Cognitive Processes, 24*(5), 749-760. doi:10.1080/01690960802648491

Alario, F.-X., & Moscoso Del Prado, F. (2010). On the origin of the "cumulative semantic inhibition" effect. *Memory & Cognition*, *38*(1), 57-66. doi:10.3758/MC.38.1.57

Anderson, J. D. (2008). Age of acquisition and repetition priming effects on picture naming of children who do and do not stutter. *Journal of Fluency Disorders, 33*(2), 135-155.

Anderson, M. C. (2003). Rethinking interference theory: Executive control and the mechanisms of forgetting. *Journal of Memory and Language, 49*(4), 415-445.

Anderson, M. C., & Neely, J. H. (1996). Interference and inhibition in memory retrieval. In E. L. Bjork, & R. A. Bjork (Eds.), *Memory* (pp. 237-313). San Diego, CA: Academic Press.

Anderson, M. C., & Spellman, B. A. (1995). On the status of inhibitory mechanisms in cognition: Memory retrieval as a model case. *Psychological Review, 102*(1), 68-100.

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*(5), 1063-1087.

Baayen, R. H. (2008). *Analyzing Linguistic Data: A Practical Introduction to Statistics Using R. Sociolinguistic Studies* (Vol. 2). doi:10.1558/sols.v2i3.471

Baayen, R. H. (2011). languageR: Data sets and functions with "Analyzing Linguistic Data: A practical introduction to statistics". R package version 1.2. http://CRAN.R-project.org/package=languageR.

Baayen, R. H., & Milin, P. (2010). Analyzing reaction times. *International Journal of Psychological Research*, *3*(2), 12-28.

Bates, D., Maechler, M., & Bolker, B. (2010). lme4: Linear mixed-effects models using S4 classes. R package version 0.999375-39. http://CRAN.R-project.org/package=lme4.

Becker, S., Moscovitch, M., Behrmann, M., & Joordens, S. (1997). Long-term semantic priming: a computational account and empirical evidence. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*(5), 1059-82.

Belke, E. (2008). Effects of working memory load on lexical-semantic encoding in language production. *Psychonomic Bulletin & Review*, *15*(2), 357-363. doi:10.3758/PBR.15.2.357

Belke, E., & Meyer, A. S. (2007). Single and multiple object naming in healthy ageing. *Language and Cognitive Processes*, *22*(8), 1178-1211. doi:10.1080/01690960701461541

Belke, E., Brysbaert, M., Meyer, A. S., & Ghyselinck, M. (2005). Age of acquisition effects in picture naming: evidence for a lexical-semantic competition hypothesis. *Cognition*, *96*(2), B45-54. doi:10.1016/j.cognition.2004.11.006

Belke, E., Meyer, A. S., & Damian, M. F. (2005). Refractory effects in picture naming as assessed in a semantic blocking paradigm. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, *58*(4), 667-692. doi:10.1080/02724980443000142

Biegler, K. a, Crowther, J. E., & Martin, R. C. (2008). Consequences of an inhibition deficit for word production and comprehension: Evidence from the semantic blocking paradigm. *Cognitive neuropsychology*, *25*(4), 1. doi:10.1080/02643290701862316

Blaxton, T. A., & Neely, J. H. (1983). Inhibition from semantically related primes: Evidence of a category-specific inhibition. *Memory & Cognition, 11*(5), 500-510.

Bock, K., & Griffin, Z. M. (2000). The persistence of structural priming: Transient activation or implicit learning? *Journal of Experimental Psychology: General, 129*(2), 177-192.

Brown, A. S. (1979). Priming effects in semantic memory retrieval processes. *Journal of Experimental Psychology: Human Learning and Memory, 5*(2), 65-77.

Brown, A. S. (1981). Inhibition in cued retrieval. *Journal of Experimental Psychology: Human Learning and Memory*, *7*(3), 204-215. doi:10.1037/0278-7393.7.3.204

Castro, Y. G., Strijkers, K., Costa, A. & Alario, X. (2008, July). *When cat competes with dog but not with perro: Evidence from the semantic competitor paradigm.* Poster presented at the Fifth International Workshop on Language Production, Annapolis, MD.

Chang, F., Dell, G. S., & Bock, K. (2006). Becoming syntactic. *Psychological Review, 113*(2), 234-272.

Cohen, L. & Dehaene S., (1998). Competition between past and present: Assessment and interpretation of verbal perseverations. *Brain, 121*, 1641–1659.

Collins, A. M. & Loftus, E. F. (1975). A spreading activation theory of semantic memory. *Psychological Review, 82*, 407-428.

Collins, A. M. & Quillian, E. R. (1969). Retrieval time from semantic memory. *Journal of verbal learning and verbal memory, 8*, 240-247.

Crowther, J. E., Martin, R. C., & Biegler, K. A. (2008, October). *The semantic blocking effect in naming: A computational model of inhibition failure.* Poster presented at the 46th Annual Meeting of the Academy of Aphasia, Turku, Finland.

Damian, M. F. (2003). Articulatory duration in single-word speech production. *Journal of experimental psychology. Learning, memory, and cognition*, *29*(3), 416-31.

Damian, M. F., & Als, L. C. (2005). Long-lasting semantic context effects in the spoken production of object names. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*(6), 1372-84. doi:10.1037/0278-7393.31.6.1372

Damian, M. F., Vigliocco, G., & Levelt, W. J. (2001). Effects of semantic context in the naming of pictures and words. *Cognition, 81*(3), B77-86.

Dell, G. S. (1986). A spreading-activation theory of retrieval in sentence production. *Psychological Review, 93*(3), 283-321.

Dell, G. S., Burger, L. K., & Svec, W. R. (1997). Language production and serial order: A functional analysis and a model. *Psychological Review, 104*(1), 123-147.

Dell, G. S., & O'Seaghdha, P. G. (1992). Stages of lexical access in language production. *Cognition, 42*(1-3), 287–314.

Dell, G. S., Oppenheim, G. M., & Kittredge, A. K. (2008). Saying the right word at the right time: Syntagmatic and paradigmatic interference in sentence production. *Language and Cognitive Processes, 23*, 583-608.

Dell, G. S., Schwartz, M. F., Martin, N., Saffran, E. M., & Gagnon, D. A. (1997). Lexical access in aphasic and nonaphasic speakers. *Psychological Review*, *104*(4), 801–838. [Washington, etc.] American Psychological Association [etc.].

Ferreira, V. S., & Firato, C. E. (2002). Proactive interference effects on sentence production. *Psychonomic bulletin & review, 9*(4), 795-800.

Ferreira, V. S., & Humphreys, K. R. (2001). Syntactic Influences on Lexical and Morphological Processing in Language Production. *Journal of Memory and Language, 44*(1), 52-80.

Ferreira, V. S., & Pashler, H. (2002). Central bottleneck influences on the processing stages of word production. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28*(6), 1187–1199.

Forde, E. M. E., & Humphreys, G. W. (1997). A semantic locus for refractory behaviour: Implications for access-storage distinctions and the nature of semantic memory. *Cognitive Neuropsychology, 14*(3), 367-402.

Ganushchak, L. Y., & Schiller, N. O. (2008). Motivation and semantic context affect brain error-monitoring activity: an event-related brain potentials study. *NeuroImage*, *39*(1), 395-405. doi:10.1016/j.neuroimage.2007.09.001

Garrett, M. F. (1975). The analysis of sentence production. In G. H. Bower (Ed.), *The psychology of learning and motivation: advances in research and theory* (pp. 133-177).

Gaskell, M. G., & Dumay, N. (2003). Lexical competition and the acquisition of novel words. *Cognition, 89*(2), 105-132. doi:10.1016/S0010-0277(03)00070-2

Goldinger, S. D. (1998). Echoes of echoes? an episodic theory of lexical access. *Psychological Review, 105*(2), 251-279.

Gordon, J. K., & Dell, G. S. (2003). Learning to divide the labor: An account of deficits in light and heavy verb production. *Cognitive Science, 27*, 1-40.

Griffin, Z. M., & Bock, J. K. (1998). Constraint, word frequency, and the relationship between lexical processing levels in spoken word production. *Journal of Memory and Language*, *38*(3), 313–338. Elsevier.

Gupta, P., & Cohen, N. J. (2002). Theoretical and computational analysis of skill learning, repetition priming, and procedural memory. *Psychological Review, 109*(2), 401-448.

Harley, T. A. (1993). Phonological activation of semantic competitors during lexical access in speech production. *Language and Cognitive Processes, 8*(3), 291–309.

Harm, M. W., & Seidenberg, M. S. (2004). Computing the meanings of words in reading: Cooperative division of labor between visual and phonological processes. *Psychological Review, 111*(3), 662-720.

Hocking, J., McMahon, K. L., & de Zubicaray, G. I. (2008). Semantic context and visual feature effects in object naming: An fMRI study using arterial spin labeling. *Journal of Cognitive Neuroscience,*

Houghton, G. (1990). The problem of serial order: A neural network model of sequence learning and recall. In R. Dale, C. Mellish & M. Zock (Eds.), *Current research in natural language generation* (pp. 287–319). London: Academic Press.

Howard, D., Nickels, L., Coltheart, M., & Cole-Virtue, J. (2006). Cumulative semantic inhibition in picture naming: experimental and computational studies. *Cognition*, *100*(3), 464-82. Elsevier. doi:10.1016/j.cognition.2005.02.006

Hsiao, E. Y., Schwartz, M. F., Schnur, T. T., & Dell, G. S. (2009). Temporal characteristics of semantic perseverations induced by blocked-cyclic picture naming. *Brain and language*, *108*(3), 133-44. Elsevier. doi:10.1016/j.bandl.2008.11.003

Jansen, P. A, & Watter, S. (2008). SayWhen: An automated method for high-accuracy speech onset detection. *Behavior Research Methods*, *40*(3), 744-751. doi:10.3758/BRM.40.3.744

Janssen, N., Schirm, W., Mahon, B. Z., & Caramazza, A. (2008). Semantic interference in a delayed naming task: Evidence for the response exclusion hypothesis. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 34*(1), 249-256.

Jefferies, E., Baker, S. S., Doran, M., & Ralph, M. a L. (2007). Refractory effects in stroke aphasia: a consequence of poor semantic control. *Neuropsychologia*, *45*(5), 1065-79. doi:10.1016/j.neuropsychologia.2006.09.009

Kan, I. P. & Thompson-Schill, S. L. (2004). Effect of name agreement on prefrontal activity during overt and covert picture naming. *Cognitive, Affective & Behavioral Neuroscience, 4*(1), 43-57.

Kittredge, A. K., & Dell, G. S. (2011). Implicit learning of phonotactic constraints: Transfer from perception to production. *Proceedings of the 33rd Annual Meeting of the Cognitive Science Society.* (pp. 2679-2684). Boston, MA.

Kraljic, T., & Samuel, A. G. (2005). Perceptual learning for speech: Is there a return to normal? *Cognitive Psychology, 51*(2), 141-178.

Kroll, J. F., & Stewart, E. (1994). Category interference in translation and picture naming: Evidence for asymmetric connections between bilingual memory representations. *Journal of Memory and Language, 33*, 149-174.

Levelt, W. J. M. (1989). *Speaking: From intention to articulation*. Cambridge, Mass.: The MIT Press.

Levelt, W. J., Roelofs, A., & Meyer, A. S. (1999). A theory of lexical access in speech production. *Behavioral and Brain Sciences, 22*(1), 1-38.

MacLeod, C. M., Dodd, M. D., Sheard, E. D., Wilson, D. E., & Bibi, U. (2003). In opposition to inhibition. In B. H. Ross (Ed.), *The psychology of learning and motivation* (pp. 163-214). San Diego, CA: Academic Press.

MacLeod, M. D., & Macrae, C. N. (2001). Gone but not forgotten: The transient nature of retrieval-induced forgetting. *Psychological Science : A Journal of the American Psychological Society / APS, 12*(2), 148-152.

Maess, B., Friederici, A. D., Damian, M., Meyer, A. S., & Levelt, W. J. M. (2002). Semantic category interference in overt picture naming: Sharpening current density localization by PCA. *Journal of Cognitive Neuroscience, 14*(3), 455-462.

Mahon, B. Z., & Caramazza, A. (2009). Why does lexical selection have to be so hard? comment on Abdel Rahman and Melinger's swinging lexical network proposal. *Language and Cognitive Processes, 24*(5), 735.

Mahon, B. Z., Costa, A., Peterson, R., Vargas, K. A., & Caramazza, A. (2007). Lexical selection is not by competition: A reinterpretation of semantic interference and facilitation effects in the picture-word interference paradigm. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 33*(3), 503-535.

Marsolek, C. J., Schnyer, D. M., Deason, R. G., Ritchey, M., & Verfaellie, M. (2006). Visual antipriming: Evidence for ongoing adjustments of superimposed visual object representations. *Cognitive, Affective & Behavioral Neuroscience, 6*(3), 163-174.

McCarthy, R. A., Kartsounis L. D. (2000).Wobbly words: refractory anomia with preserved semantics. *Neurocase, 6*, 487–97.

McClelland, J. L., & Rogers, T. T. (2003). The parallel distributed processing approach to semantic cognition. *Nature Reviews: Neuroscience, 4*(4), 310-322.

McClelland, J. L., & Rumelhart, D. E. (1985). Distributed memory and the representation of general and specific information. *Journal of experimental psychology. General*, *114*(2), 159-97.

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review, 102*(3), 419-457.

McGeoch, J. A. (1932). Forgetting and the law of disuse. *Psychological Review, 39*(4), 352-370.

Mehl, M. R., Vazire, S., Ramírez-Esparza, N., Slatcher, R. B., & Pennebaker, J. W. (2007). Are women really more talkative than men? *Science (New York, N.Y.)*, *317*(5834), 82. doi:10.1126/science.1139940

Melinger, A., & Koenig, J-P. (2007). Part-of-speech persistence: The influence of part-of-speech information on lexical processes. *Journal of Memory and Language, 56*(4), 472-489.

Melton, A. W., & Irwin, J. M. (1940). The influence of degree of interpolated learning on retroactive inhibition and the overt transfer of specific responses. *The American Journal of Psychology, 100*(3-4), 610-641.

Mensink, G., & Raaijmakers, J. G. (1988). A model for interference and forgetting. *Psychological Review, 95*(4), 434-455.

Mitchell, D. B., & Brown, A. S. (1988). Persistent repetition priming in picture naming and its dissociation from recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14*(2), 213-222.

Moss, H. E., Abdallah, S., Fletcher, P., Bright, P., Pilgrim, L., Acres, K., *et al.* (2005). Selecting among competing alternatives: Selection and retrieval in the left inferior frontal gyrus. *Cerebral Cortex, 15*, 1723-1735.

Navarrete, E., Mahon, B. Z., & Caramazza, A. (2008). The cumulative semantic cost in picture naming (Il Costo Cumulativo Intracategoriale in compiti di denominazio-ne di figure). *Paper presented at the XIV Congresso Nazionale Associazione Italiana di Psicologia, Padova, Italia, September.*

Navarrete, E., Mahon, B. Z., & Caramazza, A. (2010). The cumulative semantic cost does not reflect lexical selection by competition. *Acta psychologica*, *134*(3), 279-89. doi:10.1016/j.actpsy.2010.02.009

Nickels, L., Howard, D., Dodd, H., & Coltheart, M. (2008). Further investigations into cumulative semantic inhibition in picture naming: effects of long lag & repeated members of a category. *5th International Workshop on Language Production, July 27-30, Annapolis, MD.*

Nooteboom, S. G. (1969). The tongue slips into patterns. In A. G. Sciarone, A. J. van Essen, & A. A. van Raad (Eds.), *Leyden studies in linguistics and phonetics* (pp. 114-132). The Hague: Mouton.

Norman, K. A., Newman, E. L., & Detre, G. (2007). A neural network model of retrieval-induced forgetting. *Psychological Review, 114*(4), 887-953.

Oppenheim, G. M., & Dell, G. S. (2010). Motor movement matters: the flexible abstractness of inner speech. *Memory & cognition*, *38*(8), 1147-1160.

Oppenheim, G. M., Dell, G. S., & Schwartz, M. F. (2010). The dark side of incremental learning: A model of cumulative semantic interference during lexical access in speech production. *Cognition*, *114*, 227–252.

Oppenheim, G. M., Dell, G. S., & Schwartz, M. F.(2007). Cumulative semantic interference as learning. *Brain and Language, 103*, 175-176.

Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review, 103*(1), 56-115.

Postman, L., Stark, K., & Fraser, J. (1968). Temporal changes in interference. *Journal of Verbal Learning and Verbal Behavior, 7*, 672-694.

R Development Core Team. (2010). R: A Language and Environment for Statistical Computing. Vienna, Austria: R Foundation for Statistical Computing.

Rapp, B., & Goldrick, M. (2000). Discreteness and interactivity in spoken word production. *Psychological Review, 107*(3), 460-499.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black and W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64–99). New York: Appleton-Century-Crofts.

Roediger, H. L., Neely, J. H., & Blaxton, T. A. (1983). Inhibition from related primes in semantic memory retrieval: A reappraisal of brown's (1979) paradigm. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 9*(3), 478-485.

Roelofs, A. (1992). A spreading-activation theory of lemma retrieval in speaking. *Cognition, 42*(1-3), 107-142.

Roelofs, A. (2010). Attention, temporal predictability, and the time course of context effects in naming performance. *Acta psychologica*, *133*(2), 146-53. Elsevier B.V. doi:10.1016/j.actpsy.2009.11.003

Rogers, T. T., Lambon Ralph, M. A., Garrard, P., Bozeat, S., McClelland, J. L., Hodges, J. R., et al. (2004). Structure and deterioration of semantic memory: A neuropsychological and computational investigation. *Psychological Review, 111*(1), 205-235.

Rumelhart, D. E., McClelland, J. L., and the PDP research group. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition. Volume I.* Cambridge, MA: MIT Press.

Scarborough, D. L., Cortese, C., & Scarborough, H. S. (1977). Frequency and Repetition Effects in Lexical Memory. *Journal of experimental psychology. Human perception and performance*, *3*(1), 1-17.

Schnur, T. T., Lee, E., Coslett, H. B., Schwartz, M. F., & Thompson-Schill, S. L. (2005). When lexical selection gets tough, the LIFG gets going: A lesion analysis study of interference during word production. *Brain and Language, 95*(1), 12 – 13.

Schnur, T. T., Schwartz, M. F., Brecher, A., & Hodgson, C. (2006). Semantic interference during blocked-cyclic naming: Evidence from aphasia. *Journal of Memory and Language*, *54*(2), 199-227. doi:10.1016/j.jml.2005.10.002

Schnur, T. T., Schwartz, M. F., Kimberg, D. Y., Hirshorn, E. A, Coslett, H. B., & Thompson-Schill, S. L. (2009). Localizing interference during naming: Convergent neuroimaging and neuropsychological evidence for the function of Broca's area. *Proceedings of the National Academy of Sciences*, *106*(1), 322-327.

Schriefers, H, Meyer, A. S., & Levelt, W. J. M. (1990). Exploring the time course of lexical access in language production: Picture-word interference studies. *Journal of Memory and Language, 29*, 86-102.

Seidenberg, M. S., & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological review, 96*(4), 523–568.

Smith, M., & Wheeldon, L. R. (2004). Horizontal information flow in spoken sentence production. *Journal of experimental psychology. Learning, memory, and cognition, 30*(3), 675-86. doi:10.1037/0278-7393.30.3.675

Snodgrass, J. G., & Vanderwart, M. (1980). A standardized set of 260 pictures: norms for name agreement, image agreement, familiarity, and visual complexity. *Journal of experimental psychology. Human learning and memory*, *6*(2), 174-215.

Stemberger, J. P. (1985). An interactive activation model of language production. In A. W. Ellis (Ed.), *Progress in the psychology of language* (pp. 143–186). Hillsdale, NJ: Lawrence Erlbaum Associates.

Storm, B. C., Bjork, E. L., Bjork, R. A., & Nestojko, J. F. (2006). Is retrieval success a necessary condition for retrieval-induced forgetting? *Psychonomic Bulletin & Review, 13*(6), 1023-1027.

Strain, E., Patterson, K., & Seidenberg, M. S. (1995). Semantic effects in single-word naming. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*(5), 1140.

Thompson-Schill, S. L., D'Esposito, M., Aguirre, G. K., & Farah, M. J. (1997). Role of left inferior prefrontal cortex in retrieval of semantic knowledge: A reevaluation. *Proceedings of the National Academy of Sciences of the United States of America, 94*(26), 14792-14797.

Thompson-Schill, S. L., Schnur, T. T., Hirshorn, E. a, Schwartz, M. F., & Kimberg, D. Y. (2007). Regulatory functions of prefrontal cortex during single word production. *Brain and Language*, *103*(1-2), 171-172. doi:10.1016/j.bandl.2007.07.100

Tremblay, A. (2011). LMERConvenienceFunctions: A suite of functions to back-fit fixed effects and forward-fit random effects, as well as other miscellaneous functions.. R package version 1.5.3. http://CRAN.R-project.org/package=LMERConvenienceFunctions.

Tse, C. Y., Lee, C. L., Sullivan, J., Garnsey, S. M., Dell, G. S., Fabiani, M., et al. (2007). Imaging cortical dynamics of language processing with the event-related optical signal. *Proceedings of the National Academy of Sciences of the United States of America, 104*(43), 17157-17162.

Vigliocco, G., Vinson, D. P., Damian, M. F., & Levelt, W. J. M. (2002). Semantic distance effects on object and action naming. *Cognition*, *85*(3), B61-B69. Elsevier. doi:10.1016/S0010-0277(02)00107-5

Vitkovitch, M., Cooper-Pye, E., & Ali, L. (2010). The long and the short of it! Naming a set of prime words before a set of related picture targets at two different intertrial intervals. *European Journal of Cognitive Psychology, 22*(2), 161-171. doi:10.1080/09541440902743348

Vitkovitch, M., Rutter, C., & Read, A. (2001). Inhibitory effects during object name retrieval: The effect of interval between prime and target on picture naming responses. *British journal of psychology (London, England : 1953)*, *92 Part 3*, 483-506.

Wagenmakers, E.-J., Farrell, S., & Ratcliff, R. (2004). Estimation and interpretation of 1/falpha noise in human cognition. *Psychonomic bulletin & review*, *11*(4), 579-615.

Warker, J. A., & Dell, G. S. (2006). Speech errors reflect newly learned phonotactic constraints. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*(2), 387. APA AMERICAN PSYCHOLOGICAL ASSOCIATION.

Wheeldon, L. R., & Monsell, S. (1994). Inhibition of Spoken Word Production by Priming a Semantic Competitor. *Journal of Memory and Language*, *33*(3), 332-356. ACADEMIC PRESS INC. doi:10.1006/jmla.1994.1016

Widrow, B., & Hoff, M. E. (1960). Adaptive switching circuits. *IRE WESCON convention record* (pp. 96–104). New York: IRE. (Reprinted in Anderson, J. A., & Rosenfeld, E. (Eds.)

(1988). *Neurocomputing: Foundations of research* (pp. 123–134). Cambridge, MA: MIT Press.)

Wilshire, C. E., & McCarthy, R. A. (2002). Evidence for a context-sensitive word retrieval disorder in a case of nonfluent aphasia. *Cognitive Neuropsychology*, *19*(2), 165-186. doi:10.1080/02643290143000169

The Cycle 1 incremental interference effect in Experiments 1 and 2

Expt 1

OrdPos	Unrelated	Related	BlockEff
1	794	756	-37
2	698	697	-1
3	701	708	8

Expt 2

OrdPos	Unrelated	Related	BlockEff
1	699	681	-18
2	690	694	4
3	696	722	26

Experiments 1 and 2: combined analyses

Table 18. Combined error data for Experiments 1 and 2.

	Session1		Session 2	
Outcome	Homog	Heterog	Homog	Heterog
Correct	3212	3270	3165	3200
Semantic error	27	7	27	5
Completed	11	6	8	5
Interrupted	16	1	19	2
Other within-block error	2	11	2	7
Completed	0	1	0	0
Interrupted	2	10	2	7
Total within-block errors	*29*	*18*	*29*	*12*
Omission	1	0	2	0
Other error	24	24	35	38
Perspective	10	5	13	17
Phonological	4	2	4	2
Disfluency	7	8	7	7
Other or uncodable	4	8	13	13
Lipsmacks, *etc.*	171	134	202	184
Equipment failure	18	11	21	20

Table 19. Mean RT data for Experiments 1 and 2 (combined).

	Session 1						Session 2						Difference	
	Hom_{A1}		Het_{B1}		Hom_{A1}-Het_{B1}		Hom_{A2}		Het_{C1}		Hom_{A2}-Het_{C1}		$Diff_2$-$Diff_1$	
Cycle	Mean	SE	Mean	SE	Mean	SE	Mean	SE	Mean	SE	Mean	SE	Mean	SE
1	696.1	11.1	701.8	12.5	-5.7	8.3	721.3	11.3	721.4	11.4	-0.1	10.1	5.6	13.2
2	594.6	11.3	595.8	11.4	-1.2	7.2	605.5	12.7	602.0	11.4	3.5	7.7	4.7	9.7
3	593.7	11.1	581.9	11.5	11.8	5.1	605.9	11.8	590.7	11.7	15.2	7.6	3.4	7.6
4	597.2	10.8	586.8	11.1	10.4	6.9	607.3	12.8	591.3	11.6	16.0	6.5	5.6	9.3
5	597.1	9.7	588.1	11.2	9.0	7.1	609.3	11.3	592.8	11.4	16.5	5.3	7.5	8.2
6	592.6	10.2	586.1	11.9	6.5	6.9	607.0	13.7	597.6	11.8	9.4	7.2	2.8	10.5
Mean	611.9	4.9	606.7	5.3	5.1	2.9	626.0	5.6	616.0	5.4	10.1	3.1	4.9	4.0

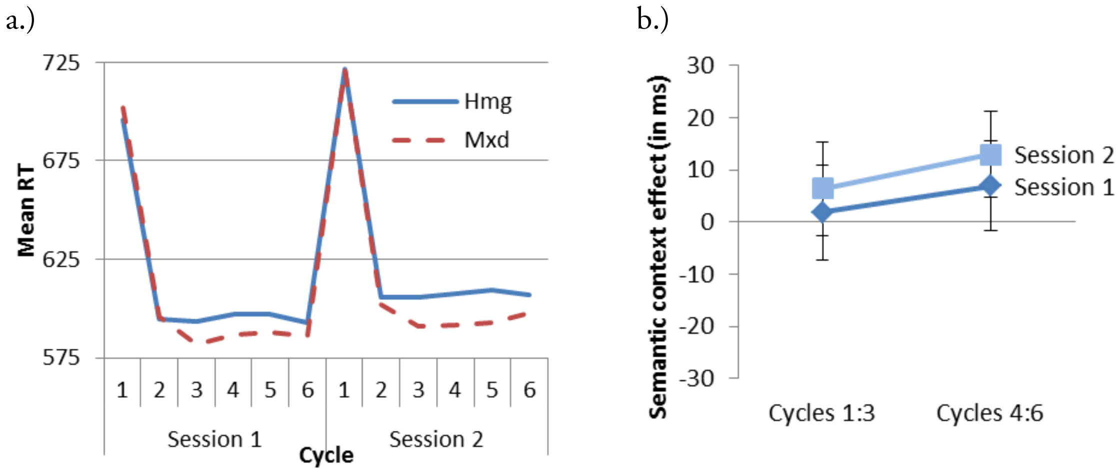

Figure 24. Mean naming latencies (a) and estimated semantic context effects (b) for Experiments 1 and 2, combined. Estimates in (b) are simple main effects, drawn from a linear mixed effects model; error bars depict MCMC-estimated 95% confidence intervals.

Table 20. Summary of the linear mixed effects regression for Experiments 1 and 2, combined. Since this is already a very large model, I represent Cycle as a linear effect.

	Estimate	[95% HPD Int.]		p
(Intercept)	614.88	599.71	630.45	**<.001**
Previous trial RT	0.14	0.12	0.17	**<.001**
Session (1:2)	8.87	0.32	17.73	**0.04**
Cycle (1:6)	−13.69	−16.57	−10.92	**<.001**
Position in cycle (1:3)	−11.06	−14.66	−7.39	**<.001**
× Cycle	1.42	0.09	2.85	**0.05**
Semantic context (Het,Hom)	7.06	1.64	12.45	**0.01**
× Session	5.00	−4.22	14.07	0.28
× Cycle	2.12	0.01	4.44	**0.06**
× Position in cycle	1.51	−2.96	6.41	0.53
× Cycle	−2.71	−6.27	0.79	0.13
Experiment (1a,1b)	33.98	4.72	63.62	**0.03**
× Session	14.41	−3.73	31.62	0.11
× Cycle	6.94	2.81	11.13	**<.001**
× Position in cycle	29.39	22.19	36.89	**<.001**
× Cycle	−4.16	−6.92	−1.42	**<.001**
× Semantic context	0.13	−9.28	10.10	0.98
× Session	−2.51	−22.00	15.15	0.8
× Cycle	−2.24	−6.61	2.26	0.32
× Position in cycle	2.87	−6.12	12.51	0.55
× Cycle	−0.10	−5.59	5.50	0.98

Experiment 3 supplementary information

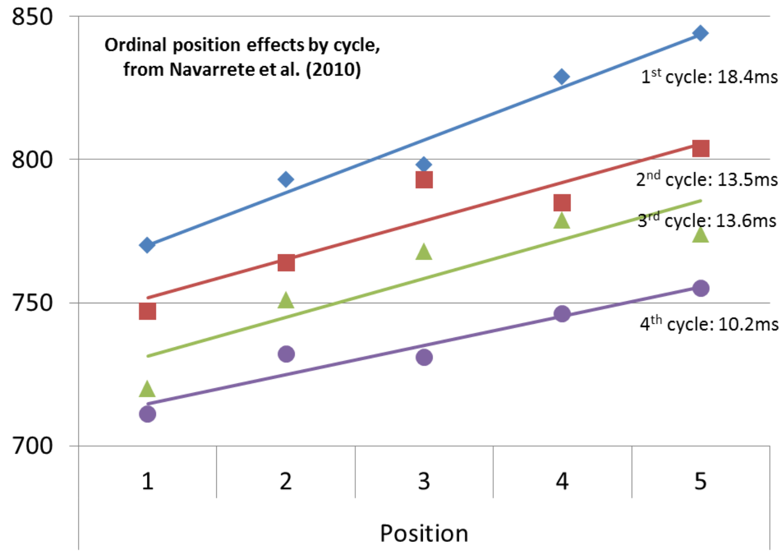

Figure 25. Navarrete et al.'s (2010) data suggest a decrease in the slope of the incremental interference effect, similar to that found in Experiment 3.

a.)

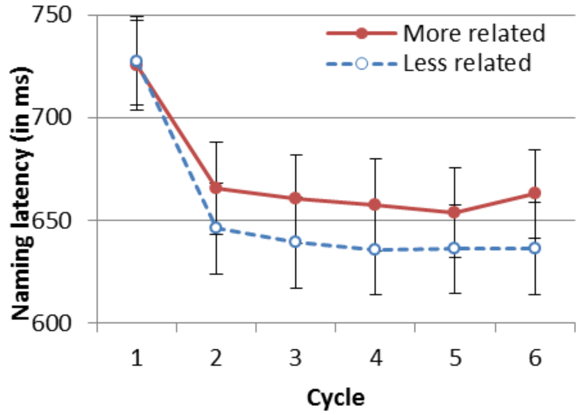

b.)

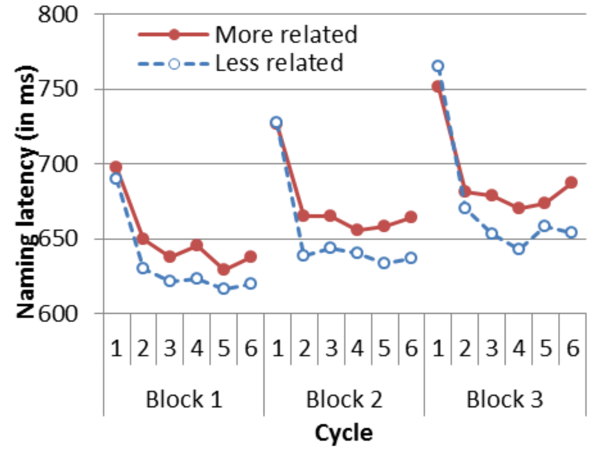

c.)

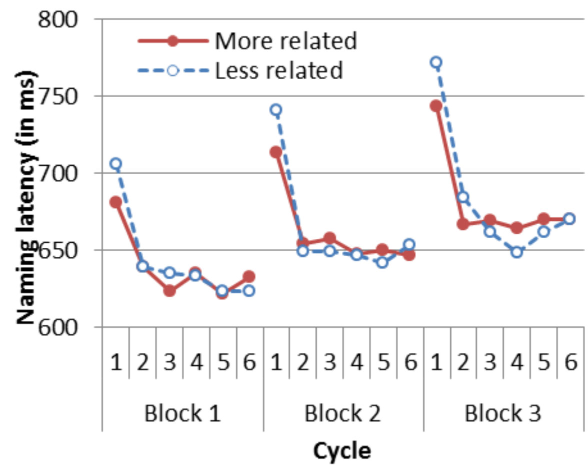

Figure 26. Apparent semantic blocking effects in Experiment 3. For these figures, Cohort similarity ratings are discretized, broken into two sets of ~10,000 trials, according to whether a target's pairwise relatedness with the other words in its set was above or below the median. (a) In-set relatedness measures (just comparing each item to the other two in its cohort) seem most informative, and appear to produce a semantic blocking effect on their own –though the lack of an effect in the first cycle is puzzling. (b) There is some hint that that first cycle issue develops over the course of the experiment. (c) One would think that a cross-set relatedness measure might correlate well with measures of persistent interference, but it

seems there is something strange and different going on there – there should not be any effect in Block 1.