

Using the letter decision task to examine semantic priming

Tom Heyman (tom.heyman@ppw.kuleuven.be)

Department of Experimental Psychology, 102 Tiensestraat
3000 Leuven, Belgium

Simon De Deyne (simon.dedeyne@ppw.kuleuven.be)

Department of Experimental Psychology, 102 Tiensestraat
3000 Leuven, Belgium

Gert Storms (gert.storms@ppw.kuleuven.be)

Department of Experimental Psychology, 102 Tiensestraat
3000 Leuven, Belgium

Abstract

The present research investigates semantic priming with an adapted version of the word fragment completion task. The letter decision task, as we will call it, holds some advantages over the traditionally used lexical decision task in that it eliminates retrospective semantic matching effects, it avoids the need to construct pseudowords, it is more engaging for participants and it enhances semantic processing, which in turn allows for a more fine-grained investigation of semantic activation. The letter decision task requires participants to complete words, from which one letter was omitted like lett_ce (*lettuce*), as fast as possible. The study found that words are completed faster when the preceding trial comprised a semantically related fragment like tom_to (*tomato*) than when it comprised an unrelated fragment like guit_r (*guitar*). Furthermore, the study provides insight in the nature of the priming effect. It demonstrates that priming effects are larger for strongly associated prime-target pairs.

Keywords: Semantic priming; Letter decision task; Associative strength.

Introduction

Semantic priming is the finding that the processing of targets (e.g., a picture, a word,...) preceded by a semantically related prime (also a picture, a word,...) is enhanced. For instance, the presentation of the word *cat* facilitates processing of the subsequently presented word *dog*. One of the debates in the semantic priming literature concerns the source of the priming effect (Hutchison, 2003; Lucas, 2000). The (unresolved) issue revolves around the type of relation between concepts that is necessary for priming to occur. That is to say, words can be associatively related, as evidenced by association norms (De Deyne, Navarro & Storms, 2012) or because both concepts share certain features. Returning to the *cat-dog* example, both cats and dogs have four legs, two eyes, are pets, etc. and thus they are related in terms of feature overlap (e.g., McRae & Boisvert, 1998). Moreover, the strongest associate of *cat* is *dog* hence both concepts are also associatively related. Whether priming is driven by word associations or feature overlap (or even something else) is an important question

since it has significant repercussions for theories about the organization of the mental lexicon. Consequently, a lot of research has been devoted to this topic.

The most frequently used paradigms to examine these issues are the lexical decision task, in which participants have to decide whether letter strings form existing words or not, and, to a lesser extent, the pronunciation task, in which participants read aloud words (see the reviews of Hutchison (2003), Lucas (2000) and Neely (1991)). The experimental designs further vary in the degree to which they allow automatic and controlled processes. These latter processes are conscious and strategic and they come into play when the prime-target coupling (e.g., *cat-dog*) is made explicit (Jones, 2010). This is for instance the case in the standard lexical decision task where participants are required to respond only to the second item of the pair (i.e., the target *dog*) and not to the first (i.e., the prime *cat*). Strategic effects are volatile and vary over subjects, whereas automatic processes are ubiquitous. Thus, automatic processes are thought to reliably reflect the structure of the mental lexicon (Lucas, 2000). Hence, considerable effort has been put into developing methodologies that prevent controlled processes. One method to reduce strategic effects is the continuous lexical decision task (McNamara & Altarriba, 1988; Shelton & Martin, 1992). Here, prime-target pairs are decoupled by asking participants to respond not only to the target but also to the prime.

In the present study, we took a different approach. It was (partly) motivated by the fact that there is little consensus regarding the nature of semantic priming. A possible explanation for the divergent and sometimes unreplicated findings (see Hutchison (2003) and Lucas (2000)) is that the experimental paradigms are not sensitive enough to detect or tease apart subtle effects. The widely used lexical decision task may rely more on superficial processing of words, whereas deeper semantic processing may be necessary to fully uncover the structure of the mental lexicon. Hence, in this study, we used a different method to examine semantic priming. It is an adaptation of the word fragment completion task, a task that has mainly been used in implicit memory studies (i.a., Bassili, Smith & MacLeod,

1989; Challis & Brodbeck, 1992; McDermott, 1997; Roediger & Challis, 1992; Weldon, 1993). There are several variants of the word fragment completion task, but the general idea is that participants are presented with words from which one or more letters are omitted (e.g., *r_d* or *_orn_d*). Participants then are assigned to fill in the gap(s). In some experiments, the dependent variable of interest is the actual answer participants give. Put differently, the question is whether participants complete *r_d* as *red* or as *rod*. In other experiments, there is only one correct answer and the crucial dependent variable is the proportion correct responses within a certain time interval or alternatively, the time required to give the correct solution. Concretely, how many participants accurately identify *_orn_d* as *tornado* and/or what is the average reaction time? In this study, we examined semantic priming using a modification of the latter type. But instead of difficult words with many blank spaces, we opted for relatively simple stimuli with only one blank space. Furthermore, participants were told that the missing letter was always a vowel. The task conceptually resembled a continuous lexical decision task in that participants had to complete both prime and target words (and also unrelated filler items). For instance, on trial *n* participants got the fragment *tom_to* (it should be completed as *tomato*) and on trial *n+1* they got *lett_ce* (it should be completed as *lettuce*). For the sake of clarity, we will therefore coin the term continuous letter decision task to refer to the experimental paradigm in this study. As in a (continuous) lexical decision task, the main dependent variable is reaction time since accuracy will be near perfect. Hence, it is expected that *lett_ce* is completed faster when it is preceded by a semantically related stimulus like *tom_to* than when it is preceded by an unrelated stimulus like *guit_r* (it should be completed as *guitar*).

We believe that there are some advantages of the continuous letter decision task over the continuous lexical decision task. First of all, in the lexical decision task participants may endorse a retrospective semantic matching strategy. Neely and Keefe (1989) argued that participants might use information about whether the considered letter string is semantically related to the preceding letter string to reduce their response time. Concretely, when there is a semantic relation between two consecutively presented letter strings, the correct answer for the latter letter string is always “word”. If there is no such relation, the second letter string is a word *or* a non-word. In fact, when the proportion of non-words in the experiment is high then the absence of a relation between two consecutive letter strings indicates that the second letter string is more likely to be a non-word. It is possible that participants notice these contingencies, which in turn yields strategic priming effects that are inseparable from (interesting) automatic priming effects. However, the continuous letter decision task introduced here does not suffer from a semantic matching strategy. That is to say, a semantic relation between two words on consecutive trials is not predictive for the correct response to the latter word fragment. The fact that *tomato* and *lettuce* are related does

not give information about which vowel is missing in the fragment *lett_ce*.

A second advantage of the letter decision task with respect to the lexical decision task is that it obviates the need to construct pseudowords. Besides practical convenience, it has also theoretical implications since previous research suggested that the nature of the pseudowords and their similarity to real words modifies priming (Shulman & Davison, 1977) and also the word frequency effect (Stone & Van Orden, 1993). Such issues are avoided in the letter decision task.

Thirdly, it is not far-fetched to argue that the letter decision task is more challenging, without becoming burdensome, than the lexical decision task. Although participants may not exactly be filled with joy when performing the experiment, the task is more engaging, which in turn enhances the intrinsic motivation of participants (Deci & Ryan, 1985).

Finally and perhaps most importantly, the letter decision task presumably involves a deeper semantic processing. In the lexical decision task, shallow processing of letter strings may be sufficient to discriminate words from non-words (Rogers, Lambon Ralph, Hodges & Patterson, 2004), thereby limiting the facilitatory effect of a related prime. Because the letter decision task is more effortful, a related prime has more potential to exert its influence.

Taken together, it may be fruitful to use the letter decision task to examine semantic priming. Hence, the first goal of the present study was to establish whether a priming effect could be obtained with this task.

A second goal was to examine the nature of the priming effect. Every crucial target like *lett_ce* (*lettuce*) was either preceded by a related prime (*tom_to*, *tomato*) or an unrelated prime (*guit_r*, *guitar*). As is traditionally the case in priming research, one could consider relatedness as a dichotomy (i.e., *tomato-lettuce* are related whereas *guitar-lettuce* are not). However, one could argue that relatedness is not an all or none matter, but rather that there is variability in the strength with which two words are related (for a similar proposal, see Hutchison, Balota, Cortese & Watson, 2008). For instance, *thunder-lightening* has a stronger forward association than *tomato-lettuce*, meaning that more people give *lightning* as an association for *thunder* than *lettuce* as an association for *tomato* (based on the large scale Dutch Word Association Database from De Deyne et al., 2012). Thus, one might hypothesize that the priming effect for *thunder-lightening* is stronger than the effect for *tomato-lettuce*. The second goal of this study was to examine this prediction.

Method

Participants

Participants were 40 first-year psychology students of the University of Leuven (7 men, 33 women, mean age 18 years), who participated in return for course credit. All participants were native Dutch speakers.

Materials

A total of 76 related prime-target pairs like *tom_to-lett_ce* (*tomato-lettuce*) were constructed. All stimuli were Dutch word fragments. Primes and targets were always category coordinates. Categories ranged from fruits and music instruments to mammals, tools, professions, etc. Moreover, prime-target pairs had a forward association strength that ranged from 3% to 30%. These and other measures of association strength were derived from the Dutch Word Association Database (De Deyne et al., 2012). In addition, 76 unrelated filler pairs were constructed.

All word fragments were generated by omitting one vowel from a Dutch noun. Only word fragments that had a unique correct response were used. Of the 76 crucial targets, 16 required an “a” response, 22 an “e” response, 18 an “i” response, 13 an “o” response and 7 a “u” response.

Two lists were created such that a random half of the 76 crucial targets were preceded by their related prime in List A, whereas in List B they were preceded by an unrelated word, and vice versa. The 38 unrelated pairs for each list were constructed by randomly recombining primes and targets, with two limitations. The first is of course that the resulting prime-target pairs were no category coordinates and indeed unrelated, as evidenced by a lack of a forward and backward association between prime and target. Second, a fraction of the related prime-target pairs were response congruent, meaning that the same vowel is missing in both the prime and the target. The unrelated pairs were created in a way that they match in terms of response congruency. When a related pair is response congruent so is the corresponding unrelated pair and the other way around. So for example, there where *pa_rd* (to be completed as *paard*, Dutch for *horse*) was preceded by *zebr_* (to be completed as *zebra*) in List A, it was preceded by *t_rwe* (to be completed as *tarwe*, Dutch for *wheat*) in List B, which was actually the prime for *me_l* (to be completed as *meel*, Dutch for *flour*) in List A. Hence, each list consists of 76 critical prime-target pairs (38 related pairs and 38 unrelated pairs) and an additional 76 unrelated filler pairs.

Procedure

Participants were randomly assigned to one of the two lists. Twenty participants received List A and 20 List B. The task itself was a continuous letter decision task. The continuous nature of the task breaks the 152 pairs down to 304 trials. On each trial, participants were presented with one word fragment. Primes were always shown on odd-numbered trials and targets on even-numbered trials. The order of the pairs within the experiment was random and varied over participants.

On every trial, participants saw a word from which one letter was omitted. They were informed that the missing letter was always a vowel. Participants had to complete the word by pressing either “a”, “e”, “u”, “i”, or “o” on an AZERTY keyboard. The instructions stressed both speed and accuracy. Every word fragment was displayed in the center of the screen and remained present until a response

was made. The inter-trial interval was 500 ms. Before the experimental phase, participants did 20 practice trials. The practice trials were identical to the experimental trials except that 20 new semantically unrelated word fragments were utilized. The experiment was run on a Dell Pentium 4 with a 17.3-inch CRT monitor using Psychopy (Peirce, 2007). It was part of a series of unrelated experiments and took approximately 15 minutes.

Results

First, the split-half reliability of the response times to the 76 crucial targets was calculated using the Spearman-Brown formula. Split-half correlations for List A and List B separately were obtained for 10,000 different randomizations of the participants. The resulting reliabilities, averaged over the 10,000 randomizations, were .92 for List A and .88 for List B, which is rather high for response times. Note that all analyses were performed only on the 76 crucial target trials.

Erroneously completed targets (3.3% of the data) and targets preceded by an incorrectly completed prime were not included in the analysis (5.3% of the data). Furthermore, responses faster than 250 ms and slower than 4000 ms were removed after which an individual cut-off value for each participant was computed as the mean response time plus 3 standard deviations. Response times exceeding this criterion were also excluded (another 3.9% of the data was discarded). The exclusion criteria are similar to regular priming studies using the standard lexical decision task, except for the exclusion of target trials following incorrect prime completion. This has to do with the continuous nature of the task: post-error slowing and/or subpar prime processing conceivably obscure target response times and/or priming effects. It should be noted though that the results were qualitatively the same if different exclusion criteria were used.

The log-transformed response times were then fitted using a mixed effects model with a random intercept for participants and items (i.e., the 76 crucial targets). The response times were regressed on 4 predictors: one critical predictor called Relatedness, which is a binary variable indicating whether the target (*lett_ce*, *lettuce*) was preceded by a related prime (*tom_to*, *tomato*) or an unrelated prime (*guit_r*, *guitar*), and three covariates, namely, Contextual Diversity of the target (CD Target¹, acquired from Keuleers, Brysbaert & New, 2010), Word Length of the target in number of characters (Length Target) and the log-transformed response time to the prime (RT Prime). To facilitate the interpretation of the effects, CD Target, Length Target and RT Prime were z-transformed. Furthermore, Relatedness was coded such that targets preceded by a related prime served as a baseline. Thus the intercept should be interpreted as the expected response time to a target with

¹ Contextual diversity is the log-transformed number of contexts in which a certain word occurs. This variable has been shown to be more informative than word frequency (Brysbaert & New, 2009).

an average length (≈ 6 characters) and an average contextual diversity (≈ 2.4) that was preceded by a related prime with an average response time (≈ 1103 ms). The analyses were carried out in R (version 2.15.2) (R development core team, 2011), employing the lme4 package (Bates & Sarkar, 2007). Markov Chain Monte Carlo p-values (pMCMC) and 95% highest posterior density intervals (HPD95) were obtained with the pvals.fnc() function of the languageR package, with 10,000 iterations (Baayen, 2008).

The results are summarized in Figure 1, which depicts the 95% highest posterior density interval for the fixed effects. Note that the HPD95 of the intercept, which ranged from 6.76 to 6.85, is not presented because it would have distorted the x-axis. Figure 1 shows that all predictors have a HPD95 that excludes zero. Hence, there is a significant priming effect (pMCMC $< .001$). To grasp the magnitude of the effect, one can derive model predictions based on the point estimates of the fixed effects (i.e., the dots in Figure 1; the estimate of the intercept was 6.8). The expected response time for the average participant and the average target following an average related prime equals 904 ms. This response time increases to 944 ms when the target is preceded by an unrelated prime. In other words, there is a priming effect of 40 ms.

In the previous analysis, Relatedness was a binary predictor. However, a continuous variable is needed to examine whether a stronger relation between word pairs yields a larger priming effect. To this end, five predictors that capture the associative strength between two words were derived from the Dutch Word Association Database (De Deyne et al., 2012). The five predictors are Forward Association Strength (i.e., how often is the target given as an associate to the prime; FS), Backward Association Strength (i.e., how often is the prime given as an associate to the target; BS) and three *semantic* relatedness measures. Semantic relatedness was calculated by computing the distributional overlap of the vector of association response counts between a pair of words as the cosine between these vectors (S raw). In addition, two variations were included, where (a) the counts were logarithmically transformed (S log) or (b) weighted using point-wise mutual information which is often used in semantic vector models (S pmi) (Church & Hanks, 1989; Turney & Pantel, 2010). Both related and unrelated prime-target pairs get a score for all five variables. For unrelated pairs, FS and BS values are all zero, but the presence of shared associates results in cosine values for S raw, S log and S pmi that are often somewhat larger than zero.

A model comparison approach was adopted to assess the merits of these continuous predictors with respect to the binary predictor. In a first step, the same mixed-effects model from the previous analysis was used, but now the binary predictor Relatedness was replaced by one of the five continuous variables. This results in six models of which the fit indices are reported in Table 1. The AIC and BIC scores reported in Table 1 evaluate the goodness of fit against the number of parameters of the model (Akaike, 1974; Schwarz,

1978). Lower values are indicative of a better fit. Since the models compared here are non-nested, AIC and BIC scores were used to assess which model, and thus which predictor, best fits the data. The results show that all continuous measures were better than the binary predictor. The best continuous predictor was S log.

In a second step, we started from the model with S log and added the other continuous variables to investigate whether they can explain the remaining variance. It turned out that only BS was a significant predictor (pMCMC = .011) besides S log (pMCMC = .006).




Figure 1: 95% highest posterior density intervals of the four regression weights. The dots represent the point estimates of the weights.

Table 1: AIC and BIC scores for the six mixed effects models. Models only differ in the predictor that captures the nature of the prime-target relations (the first column).

Predictor	AIC	BIC
S log	138.8	185.9
S raw	145.1	192.2
S pmi	141.8	188.8
FS	150.8	197.9
BS	140.1	187.1
Relatedness (binary)	152.8	199.9

Discussion

The present research proposes a different method, that is, the letter decision task, to examine semantic priming. In this task, participants are shown words from which one letter (i.e., a vowel) is omitted. Participants have to fill in the missing letter as fast as possible. Word fragments were selected such that there was only one correct completion possible, thereby making the task conceptually comparable to the lexical decision task. As argued in the introduction,

there are several advantages over the lexical decision task. Concretely, the letter decision task eliminates retrospective matching effects, it does not require experimenters to construct pseudowords, it is more engaging than the lexical decision task and it involves deeper semantic processing. Crucially, this study shows that the continuous letter decision task can capture semantic priming effects. Hence, the present task is a viable alternative to examine semantic priming in future research. The employed methodology greatly reduces strategic priming effects, although it is theoretically possible that (some) participants engaged in expectancy generation despite the low relatedness proportion². To completely disentangle automatic and strategic processes one might use a standard letter decision task with a short stimulus onset asynchrony. In this paradigm a briefly presented complete prime word is quickly replaced by a to-be-completed target. The short interval prevents expectancy generation (but not retrospective matching in a lexical decision task, see e.g., Shelton and Martin, 1992), while the letter decision task eliminates retrospective matching. In addition, one could manipulate the relatedness proportion in the continuous letter decision task to check whether expectancy generation plays a role. Our lab is currently investigating these issues.

Furthermore, this study provides evidence for the hypothesis that priming effects are greater for strongly related prime-target pairs. Models that regard relatedness as a continuous rather than a binary variable fitted the data better. More specifically, semantic relatedness and backward association strength were shown to predict the response times to the target word fragments the best. Thus, the stronger prime and target words are associated, the faster participants completed the target word. The fact that backward association strength plays a role seems to indicate that the benefit is larger for reciprocally associated prime-target pairs. These findings also highlight the value of the letter decision task. Because this task enhances semantic processing, it allows for a more detailed analysis of semantic activation, which may not be possible with a classic lexical decision task.

The method to assess the merits of continuous predictors over a binary predictor may seem a bit odd. Here, a model comparison approach was used, whereas it might be intuitively compelling to average over participants to obtain a priming effect for each separate item. Indeed, one could look at the average response time of the participants who got the related pair (e.g., *tom_to-lett_ce*, *tomato-lettuce*) and subtract it from the average response time of the participants who got the unrelated pair (e.g., *guit_r-lett_ce*, *guitar-lettuce*) and this for all 76 crucial targets. The resulting 76 priming effects could be regressed on continuous measures like forward association strength, backward association

² There were 304 trials in the experiment resulting in 303 pairs because of its continuous nature. Thus, the relatedness proportion is only 12.5% (i.e., 38/303). Note that this number may be a little higher for some participants due to the random ordering of pairs (e.g., *shower-chocolate* followed by *cake-vault*).

strength,... (see Hutchison et al., 2008 for such an approach). However, several researchers have argued against averaging over participants because it inflates type 1 error (Baayen, Davidson & Bates, 2008; Lorch & Myers, 1990; Quené & van den Bergh, 2008). Nevertheless, the results from this study are largely consistent with those from Hutchison and colleagues (2008).

It should be noted that the present research only considers associative strength of prime-target pairs. As described in the introduction, it is debated whether semantic priming is primarily driven by associations between words or by similarity in terms of feature overlap between prime and target. Although this research did not directly address this issue, it does hint at the importance of associations. But we immediately hasten to point out that all related pairs in the experiment were category coordinates, hence there will be considerable feature overlap between related primes and targets as well. Future research incorporating a continuous measure for feature overlap can provide further insight on this matter.

Acknowledgments

We wish to thank Pieter Moors for his comments and suggestions. Tom Heyman is a research assistant of the Research Foundation-Flanders (FWO-Vlaanderen). Correspondence should be addressed to Tom Heyman, Department of Experimental Psychology, University of Leuven, Tiensestraat 102, 3000 Leuven, Belgium. E-mail: tom.heyman@ppw.kuleuven.be

References

- Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19, 716-723.
- Baayen, R. H. (2008). *Analyzing Linguistic Data. A Practical Introduction to Statistics Using R*. Cambridge University Press.
- Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, 59, 390-412.
- Bassili, J. N., Smith, M. C., & MacLeod, C. M. (1989). Auditory and visual word-stem completion: Separating data-driven and conceptually driven processes. *Quarterly Journal of Experimental Psychology*, 41A, 439-453.
- Bates, D. M., & Sarkar, D. (2007). *lme4: Linear mixed-effects models using S4 classes*. R package version 0.999375-42.
- Brysbaert, M., & New, B. (2009). Moving beyond Kucera and Francis: A critical evaluation of current word frequency norms and the introduction of a new and improved word frequency measure for American English. *Behavior Research Methods*, 41, 977-990.
- Challis, B. H., & Brodbeck, D. R. (1992). Level of processing affects priming in word fragment completion. *Journal of Experimental Psychology: Learning, Memory & Cognition*, 18, 595-607.

Church, K., & Hanks, P. (1989). Word association norms, mutual information, and lexicography. In *Proceedings of the 27th Annual Conference of the Association of Computational Linguistics*, pp. 76-83, Vancouver, British Columbia.

Deci, E. L., & Ryan, R. M. (1985). *Intrinsic motivation and self-determination in human behavior*. New York: Plenum.

De Deyne, S., Navarro, D., & Storms, G. (2012). Better explanations of lexical and semantic cognition using networks derived from continued rather than single word associations. *Behavior Research Methods*.

Hutchison, K. A. (2003). Is semantic priming due to association strength or feature overlap? A microanalytic review. *Psychonomic Bulletin & Review*, 10, 785-813.

Hutchison, K.A., Balota, D.A., Cortese, M.J., & Watson, J.M. (2008). Predicting semantic priming at the item level. *The Quarterly Journal of Experimental Psychology*, 61, 1036-1066.

Jones, L. L. (2010). Pure mediated priming: A retrospective semantic matching model. *Journal of Experimental Psychology: Learning, Memory & Cognition*, 36, 135-146.

Keuleers, E., Brysbaert, M., & New, B. (2010). SUBTLEX-NL: A new measure for Dutch word frequency based on film subtitles. *Behavior Research Methods*, 42, 643-650.

Lorch, R. F., & Myers, J. L. (1990). Regression analyses of repeated measures data in cognitive research. *Journal of Experimental Psychology: Learning, Memory & Cognition*, 16, 149-157.

Lucas, M. (2000). Semantic priming without association: A meta-analytic review. *Psychonomic Bulletin & Review*, 7, 618-630.

McDermott, K. B. (1997). Priming on perceptual implicit memory tests can be achieved through presentation of associates. *Psychonomic Bulletin & Review*, 4, 582-586.

McNamara, T. P., & Altarriba, J. (1988). Depth of spreading activation revisited: Semantic mediated priming occurs in lexical decisions. *Journal of Memory and Language*, 27, 545-559.

McRae, K., & Boisvert, S. (1998). Automatic semantic similarity priming. *Journal of Experimental Psychology: Learning, Memory & Cognition*, 24, 558-572.

Neely, J. H. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. In D. Besner & G. W. Humphreys (Eds.), *Basic processes in reading: Visual word recognition*. Hillsdale, NJ: Erlbaum.

Neely, J. H., & Keefe, D. E. (1989). Semantic context effects in visual word processing: A hybrid prospective/retrospective processing theory. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory*. New York: Academic Press.

Peirce, J. (2007). PsychoPy—psychophysics software in Python. *Journal of Neuroscience Methods*, 162, 8-13.

Quené, H., & van den Bergh, H. (2008). Examples of mixed-effects modeling with crossed random effects and with binomial data. *Journal of Memory and Language*, 59, 413-425.

R development core team (2011). *R: A language and environment for statistical computing*. Vienna: R Foundation for Statistical Computing, <http://www.R-project.org>.

Roediger, H. L., III, & Challis, B. H. (1992). Effects of exact repetition and conceptual repetition on free recall and primed word-fragment completion. *Journal of Experimental Psychology: Learning, Memory & Cognition*, 18, 3-14.

Rogers, T. T., Lambon Ralph, M. A., Hodges, J. R., & Patterson, K. (2004). Natural selection: The impact of semantic impairment on lexical and object decision. *Cognitive Neuropsychology*, 21, 331-352.

Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, 6, 461-464.

Shelton, J. R., & Martin, R. C. (1992). How semantic is automatic semantic priming? *Journal of Experimental Psychology: Learning, Memory & Cognition*, 18, 1191-1210.

Shulman, H. G., & Davison, T. C. B. (1977). Control properties of semantic coding in a lexical decision task. *Journal of Verbal Learning and Verbal Behavior*, 16, 91-98.

Stone, G. O., & Van Orden, G. C. (1993). Strategic control of processing in word recognition. *Journal of Experimental Psychology: Human Perception and Performance*, 19, 744-774.

Turney, P. D., & Pantel, P. (2010). From Frequency to Meaning: Vector Space Models of Semantics. *Journal of Artificial Intelligence Research*, 37, 141-188.

Weldon, M. S. (1993). The time course of perceptual and conceptual contributions to word fragment completion priming. *Journal of Experimental Psychology: Learning, Memory & Cognition*, 19, 1010-1023.