Psychological Review
2001, Vol. 108, No. 1, 149–182

0033-295X/01/$5.00 DOI: 10.1037//0033-295X.108.1.149

Perception and Preference in Short-Term Word Priming

David E. Huber, Richard M. Shiffrin, Keith B. Lyle, and Kirsten I. Ruys
Indiana University at Bloomington

Responding optimally with unknown sources of evidence (ROUSE) is a theory of short-term priming applied to associative, orthographic–phonemic, and repetition priming. In our studies, perceptual identification is measured with two-alternative forced-choice testing. ROUSE assumes features activated by primes are confused with those activated by the target. A near-optimal decision discounts evidence arising from such shared features. Too little discounting explains the finding that primed words were preferred after passive viewing of primes. Too much discounting explains the findings of reverse preference after active processing of primes. These preference changes highlight the need to use paradigms (like the present ones) capable of separating preferential and perceptual components of priming. Evidence of enhanced perception was found only with associative priming and was very small in magnitude compared with preference effects.

This article presents a new theory of short-term priming termed *ROUSE,* standing for responding optimally with unknown sources of evidence. Short-term priming refers to paradigms in which "irrelevant" primes are presented immediately prior to a target presentation to which a response must be given; typically, the task requires a lexical decision or naming response (measured by response time, used when the target is above threshold) or identification (measured by accuracy, when the target is presented at threshold). Associative, orthographic–phonemic, and repetition priming are considered. The new theory is closely tied to the results from a new set of studies that considerably expand the set of conditions tested in such paradigms. We believe the results would appear inexplicable without the associated theory. Conversely, the theory would be hard to justify without reference to the results. These considerations lead us to delay presentation of the theory until the results of the first study are presented.

A central theme of this article is the attempt to understand the effect of a prime on performance. In particular, we are interested in distinguishing effects that alter the perceptual response to the target during and shortly after its presentation from preference effects that alter other aspects of the priming situation. These are subtle distinctions (e.g., both perceptual and preference effects can affect bias and sensitivity in signal-detection terms); their understanding requires a review of empirical and theoretical research as well as detailed analysis of our present results. Such considerations led us to organize the article in the following way. The introduction reviews the most pertinent prior empirical findings and theoretical interpretations and relates our notions of perception and preference to the notions of sensitivity and bias that are found in signal-detection theory. The first study is then presented; its results are used to motivate the ROUSE theory, which is presented next. The remaining studies test various aspects of the theory and explore additional issues.

Meyer and Schvaneveldt (1971) observed that lexical decisions were made more quickly to pairs of associated words than to pairs of unassociated words. Meyer, Schvaneveldt, and Ruddy (1974) modified the task by presenting a single prime word prior to lexical decision for a target word. In contemporary versions of this task, a prime word is presented for a duration ranging from 20 ms to several seconds and followed by a target word to which a response must be given. *Facilitation* is defined as faster or more accurate responses to targets preceded by related primes than preceded by unrelated primes. Facilitation has been observed for a number of prime–target relations, including but not limited to associations (Evett & Humphreys, 1981; Marcel, 1983; McNamara, 1994; Meyer & Schvaneveldt, 1971; Perea & Gotor, 1997), mediated associations (McKoon & Ratcliff, 1992; McNamara, 1992), semantic similarity (McRae & Boisvert, 1998; Perea & Gotor, 1997), orthographic similarity (Evett & Humphreys, 1981), phonemic similarity (Meyer et al., 1974), and repetitions (Evett & Humphreys, 1981; Humphreys, Besner, & Quinlan, 1988). These effects are what we refer to as *short-term word priming.*

In a lexical decision task, participants are asked to determine as quickly and accurately as possible whether the target string of letters is a valid word; in naming, participants simply pronounce the visually presented words. In both, response time is the measure of interest. In perceptual identification, target words are presented for tens of milliseconds and immediately postmasked. Participants attempt to identify the briefly flashed target word, and accuracy is

David E. Huber, Richard M. Shiffrin, Keith B. Lyle, and Kirsten I. Ruys, Department of Psychology, Indiana University at Bloomington.

Keith B. Lyle is now at Department of Psychology, Yale University; Kirsten I. Ruys is now at Department of Psychology, University of Amsterdam, the Netherlands.

This research was supported by Public Health Service–National Institute of Mental Health (NIMH) Grant 12717, National Science Foundation (NSF) Grant SBR 9512089, and NIMH Training Grant MH19879. David Huber was supported by an Indiana University College of Arts and Sciences Dissertation Year Research Fellowship and NSF Graduate Research Fellowship DGE-9253867 006. This article is based on David Huber's doctoral dissertation.

We thank William K. Estes and Eric-Jan Wagenmakers for helpful comments on drafts of this article.

Correspondence concerning this article should be addressed to David E. Huber, who is now at the Department of Psychology, University of Colorado, Boulder, Campus Box 345, Boulder, Colorado 80309-0345. Electronic mail may be sent to dhuber@psych.colorado.edu.

the measure of interest. In these paradigms, much experimentation and concern have been directed to the possibility that decision strategies (e.g., a tendency to respond with a word related to a prime) may affect the results.

Performance in perceptual identification is typically assessed by accuracy of naming the briefly flashed target word. In our work, a forced-choice variant of perceptual identification is used to help control decision strategies. This forced-choice paradigm was first used by Ratcliff, McKoon, and Verwoerd (1989) and later by Ratcliff and McKoon (1997) to test long-term repetition priming (in long-term repetition priming, the prime that is identical to the target is presented many trials prior to the test phase). In their technique, two choice words (always consisting of the correct target word and an incorrect foil word) were presented soon after the brief flash of the target word. This two-alternative forced-choice (2-AFC) procedure proved very useful for separating perceptual and preferential aspects of long-term repetition priming. We have borrowed this technique for the sequence of short-term priming studies reported here, and we used 2-AFC to study short-term associative, orthographic–phonemic, and repetition priming.

Within the large research effort directed toward short-term priming (e.g., see Neely, 1991, for a review of associative–semantic priming results), one major focus has been the determination of conditions leading to different amounts of facilitation (e.g., McKoon & Ratcliff, 1992; McNamara, 1992); many other studies have used short-term priming as a tool to explore various aspects of cognition. In this article, we explore conditions producing both facilitation and decrements in performance and ask how each should be interpreted. For example, does facilitation imply that more information has been extracted from the presentation of the target? Can strategies or other influences account for the results? With traditional word-identification tasks, it is difficult to determine whether changes in performance are due to an enhanced perceptual response to the target versus other factors.

Throughout this article, we make a distinction between priming that produces effects independent of the perceptual response to the target presentation (termed *preferential*) and priming that produces effects by altering the perceptual response to the target presentation (termed *perceptual*). We empirically validate this distinction through the finding that preferential effects are ubiquitous and readily change in magnitude and even direction, whereas perceptual effects are small or missing. Our terms are similar to those of Masson and Borowsky (1998) in which "contextual information" is considered separately from prime effects resulting in "perceptual encoding." More specifically, we label a prime-induced change that interacts with the extraction of information from the target presentation a *perceptual effect,* whereas other changes such as guessing biases occurring during decision making we label a *preferential effect.* In any real setting, the continuous stages of information transfer through the system from sensory processing to overt response ensures that such a distinction will be less than precise in the limit (and additional precision requires detailed modeling), but we have found the distinction useful for a variety of descriptive purposes. As a result of our definitions, in a study that equally primes both targets and foils, perceptual effects selectively enhance the choice of the target, whereas preferential effects alter processing of both target and foil in such a way that performance overall is not enhanced.

Preference factors could play a role in a variety of ways whether it be explicitly or implicitly. In lexical decision, there could be an explicit preference to respond "word" to words related to the prime. Likewise, in the naming version of perceptual identification, there could be an explicit preference to produce words related to the prime. Preference factors need not be explicit and might, for example, consist of an implicitly generated preactivation for all prime-related words. Preference effects are defined by their independence from the perceptual response to the target flash; thus, preactivation is defined to be preferential if it does not alter the extraction of (high or low level) features in the perceptual response to the target presentation. In some models, such independence would be evidenced as an additive component of preactivation. Alternatively, a prime might alter perceptual processing of the target and, if so, could do so in a way that either improves or harms perception. Improved perceptual processing might arise, for example, through increased top-down support or excitation between high-level features.

In our studies using 2-AFC testing in perceptual identification, we gained insight into the distinction between preferential versus perceptual effects of priming by manipulating the posttrial choice words. In the critical condition, both choice words were equally related to the prime. If performance in this condition was higher than that in the condition in which neither choice word was related to the prime, we assumed that the effect of priming was to enhance the perceptual extraction of target information (but this is not a mandatory conclusion; see Wagenmakers, Zeelenberg, Schooler, & Raaijmakers, 2000, for an explanation of such "both-primed benefits," which blurs the distinction between perceptual and preferential effects). In most experiments, we also included preference conditions in which only the target or only the foil was related to the prime. These conditions allowed assessment of the direction and magnitude of preference effects.

2-AFC Testing: Preference and Perception Versus Bias and Sensitivity

Suppose that a prime is presented, followed by a brief flash of the target (e.g., SAUCE), followed by two choices, one of which is the target, and the other the foil (e.g., SAUCE vs. TRAIN). The four conditions of interest are: (a) neither-primed: both target and foil unrelated to the prime (e.g., prime = SHELF); (b) both-primed: both target and foil related to the prime (e.g., prime = GRAVY); (c) target primed: target related to the prime but foil unrelated (e.g., prime = APPLE); and (d) foil-primed: foil related to the prime but target unrelated (e.g., prime = FREIGHT). In the *Results* section of Experiment 1, in the parenthetical note, we describe how these four conditions are used in combination to assess preferential and perceptual priming effects.

The signal-detection approach (e.g., MacMillan & Creelman, 1991) assumes that at the moment of decision, there are evidence values for the choices that are selections from two evidence distributions. Forced-choice performance (e.g., p[c]) is inversely, monotonically related to the overlap of these evidence distributions; as such, performance provides a measure of sensitivity with 2-AFC testing. Bias can be thought of as the placement of a criterion for making a response to a single probe item. In the case of 2-AFC testing, it is typically assumed that the alternatives are

directly compared and the better chosen (though a criterion could be assumed in this case as well).

The critical point is that sensitivity and bias are defined in terms of the evidence distributions accumulated over the whole task. Perceptual and preference factors, however, are defined in terms of task components; changes in evidence distributions or criteria that are separate from evidence arising from the flash of the target itself are termed preferential, whereas changes in evidence accumulation when the target is processed are termed perceptual.

To illustrate the difference, suppose that priming of both choices increases the evidence equally for target and foil and has no other effect. We would consider this to be a preference effect. In signal-detection terms, both evidence distributions would shift upward, but the overlap would not change and performance (i.e., sensitivity) would not change. In addition, suppose that along with the increases in evidence, priming increases the variability of both the target and foil evidence distributions. We would still regard this to be a preference effect (i.e., no selective change for the target). Yet the increase in variability would increase the overlap between the distributions and reduce performance (i.e., sensitivity). In general, in cases where both choices are primed and the effect of priming equally changes target and foil distributions, we assume that any difference from baseline is preferential. Perceptual effects should, in signal-detection terms, produce a different effect on the evidence distribution for targets than on the evidence distribution for foils. These ideas echo those of Norris (1995) and Masson and Borowsky (1998), who argued against equating changes in sensitivity with changes in perception.

Similar arguments that are based on the evidence distributions can be used to make inferences in conditions in which only the target or only the foil is primed. In particular, a difference between foil-primed and neither-primed, or between both-primed and target-primed, must signal a preferential effect, because until the choices appear, the conditions being compared are identical.

It is not easy to find studies in the literature that distinguish preference from perception. The great majority of short-term priming studies use lexical decision or naming and obtain response-time measures; these studies neither lend themselves to analyses in terms of signal detection nor unambiguously allow the separation of preference from perception. However, some perceptual identification studies with same–different responses to a single-choice word used primed foil words and can be interpreted in terms of preference and perception: Johnston and Hale (1984) used a same–different procedure to study short-term repetition priming. For comparison, they used a baseline condition that contained no prime word. In a follow up study, Hochhaus and Johnston (1996) repeated the experiment with a neutral prime word baseline and obtained similar results. In both sets of experiments, an analysis with repetition-primed targets providing hits and repetition-primed foils providing false alarms yielded reduced sensitivity compared with the unprimed situation. In addition, there was a bias in favor of repeated words. However, the sensitivity drop could have been due to a decrease in perceptual encoding of targets, to an increase in variability of preferences, or both.

Masson and Borowsky (1998, Experiments 2 and 3) used a same–different perceptual identification task to examine semantic priming. In their studies, the same (targets) and different (foils) choices presented after the target flash were equally related to the prime. They found a (modest) increase in sensitivity caused by related primes and no change in bias. Interestingly, these results held for both word primes and picture primes, even though the target presentation and the subsequent choice were words in both cases. This increase in sensitivity can be interpreted as an improvement in target perception (certain studies presented in this article provide a replication of these results as well as Johnston and Hale's (1984) results and place them in a larger context). Masson and Borowsky predicted the sensitivity increase with an attractor model of priming (Masson, 1991, 1995). The theory does not assume enhanced perceptual encoding at an early stage of encoding; however, as encoding continues, prime and target presentations interact in a manner consistent with our definition of a perceptual factor.

Experiment 1: Repetition and Associative Priming

Repetition and orthographic–phonemic priming occasionally reveal deficits (Dominguez & de Vega, 1997; Hochhaus & Johnston, 1996; Humpreys et al., 1988; Lukatela & Turvey, 1996; Lupker & Colombo, 1994; O'Seaghdha & Marin, 1997), but associative–semantic priming universally seems to produce facilitation, even when preference factors are controlled (Masson & Borowsky, 1998). We contrasted repetition and associative priming in Experiment 1, in a paradigm using perceptual identification and including for each type of priming relation the four 2-AFC conditions: neither-primed, both-primed, target-primed, and foil-primed. To provide a both-primed condition with repetition priming, two primes are necessary, and, therefore, two primes were presented on every trial in all conditions. In two additional conditions, both choice words were primed in mixed fashion: In one, the target was a repetition of one prime and the foil an associate of the other prime; in the other, the target was an associate of one prime and the foil a repetition of the other prime. Two versions of this study were run with separate participants, one in which participants actively processed the prime words and another in which participants passively viewed the prime words. The task for active priming required participants to determine whether the two prime words matched in animacy.

Method

Participants. There were 55 participants in the passive priming condition and 52 in the active priming condition. All participants in all experiments were native English-speaking Indiana University undergraduates receiving introductory psychology course credit.

Materials. The word-association norms of Nelson, McEvoy, and Schreiber (1994) were used to construct the stimulus set. These norms are based on one associative response to a prime word by each participant. One hundred and twenty prime–associate pairs were used with average association strength of .378 (meaning that with this probability, the associate was given as the first response to the prime). Prime and associate words were three to five letters in length and could be of different lengths. All prime words could be judged as animate or inanimate (i.e., they could serve as reasonably concrete nouns). The target and foil words were drawn from the same pool of associates for all conditions, but a separate pool of words, also three to five letters in length, served as the primes for the neither-primed condition and as the unrelated primes in the target-primed and foil-primed conditions. A different pool of four- and five-letter words was used for the practice sessions and the threshold determination block of trials.

Randomly generated letter-like pattern masks were used to avoid pattern-mask habituation. These were created by randomly selecting a position for a vertical bar within the width of a character. Then, two randomly determined vertical positions were chosen on each side of the character width. These were connected to each other and to the top or bottom of the vertical bar by separate line segments. The resulting appearance was a butterfly shape tilted to the right or the left (see Figure 1 for examples of the pattern masks). All words were displayed in capitalized, Times Roman, 22-point font size.

Equipment. Stimulus materials were displayed on PC monitors with presentation times synchronized to the vertical refresh. The refresh rate was 120 Hz providing display increments of 8.33 ms. To avoid phosphor decay, stimuli were displayed as black against a gray background. Participant booths were enclosed and the lighting dim to avoid eye strain. The resulting visual contrast was close to 100%. Chin rests were used to control monitor distance. Monitor distance and font size were chosen such that target words encompassed less than 3° of horizontal visual angle. All responses were collected through response boxes with four keys.

Procedure. Besides the neither-primed condition, the eight priming conditions were as follows: both repetition primed, target repetition primed, foil repetition primed, both associatively primed, target associatively primed, foil associatively primed, target repetition primed and foil associatively primed, and target associatively primed and foil repetition primed. These conditions are illustrated with examples using particular words in Table 1. Each participant received 12 trials on each primed condition and 24 trials on the neither-primed condition scattered across 120 experimental trials.

Figure 1 shows the sequence of events on each experimental trial. (Figure 1 serves as a general guideline for the sequence of events within a trial in most of the experiments in this article.) The portion of Figure 1 within the dashed box only appeared for the active priming group. Each presentation sequence for perceptual identification consisted of a fixation point for 500 ms (not shown in Figure 1), a blank screen for 500 ms (not shown in Figure 1), two prime words until response (only for active priming), two prime words for 500 ms, a briefly flashed target word (of a duration determined individually for each participant), a pattern mask of duration such that the total duration of target plus mask was 500 ms, and a final display of two choice words (for 2-AFC). Setting the total target plus mask time to 500 ms equated the duration from onset of the target word (as well as offset of the prime word) until presentation of the choice words. To reduce word-length effects, words with fewer than the maximum number of letters (five letters being the maximum in Experiment 1) were flanked on either side by pattern-mask characters. Although prime and associate could be of differing lengths, the target and foil always contained the same number of letters on a given trial as did the two prime words.

The two primes appeared above and below the center position with slightly less than 3° visual angle separating them. The 2-AFC words appeared to the left and right of the center position separated by 2° of visual angle. These choice words remained onscreen until participants responded. Participants were instructed (correctly) that one of the choices would always be the flashed target word. Following their response, feedback was given before moving to the next trial.

Choice words were randomly assigned to a priming condition and randomly assigned as targets or foils. Associate length and animate versus inanimate prime were equally assigned to the nine different conditions. Left–right target position was counterbalanced across trials randomly. If only one of the two prime words was related to the choice words (i.e., target-primed or foil-primed), its top–down position was counterbalanced across trials randomly. For the both-primed conditions, the positions of the target-related and foil-related primes were similarly counterbalanced. Presentation order of the conditions was counterbalanced randomly. The conditions and number of trials per condition were such that a prime-related choice word was equally likely to be target or foil. It was hoped that this would reduce any "explicit" strategy for choosing for or against prime-related words. With the exception of the active priming versus passive priming manipulation, all variables were within subject. To avoid contamination from long-term repetition priming, a given word appeared only once within the experiment.

Participants received 16 trials of practice on perceptual identification. In the active priming condition, perceptual identification practice was preceded by 16 trials of practice on the animacy matching task in isolation; all subsequent trials included both animacy matching and perceptual identification. For the animacy matching task, participants were instructed to press a key labeled *match* if the two prime words matched in animacy and otherwise press a key labeled *mismatch.* During practice, they immediately received feedback on their animacy judgments. Thereafter, every 8th trial

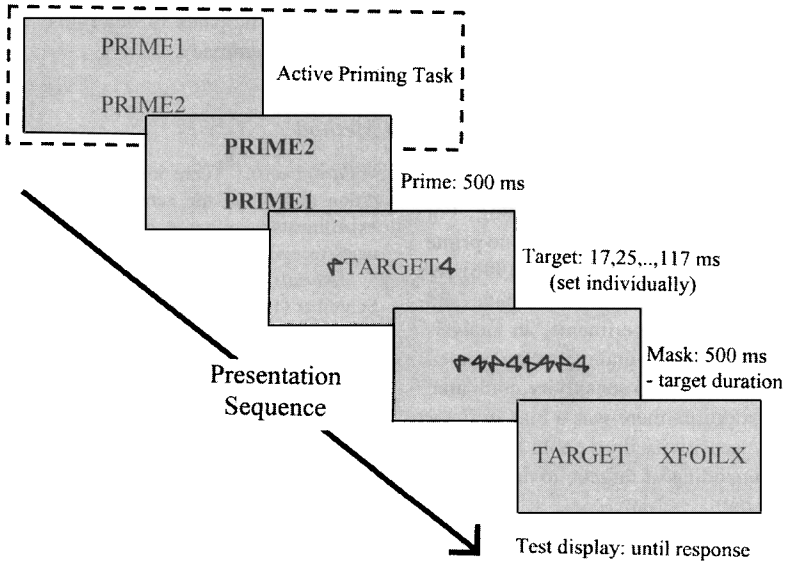

Figure 1. The sequence of displays for trials in most experiments. The display contained within the dashed box was only presented to active priming groups of participants. Prior to this sequence, a fixation point followed by a blank screen were each displayed for 500 ms.

Table 1
Experiment 1: Examples and Results

Type of priming and priming condition	Primes	Target	Choice words	Passive p(c)	Active p(c)
Neither	CHEF + ACRE			.692	.760
Associative					
both	SOCK + TOAD			.671	.793
target	SOCK + ACRE		SHOE	.733	.790
foil	CHEF + TOAD	SHOE	or	.670	.776
Repetition			FROG		
both	SHOE + FROG			.626	.647
target	SHOE + ACRE			.770	.716
foil	CHEF + FROG			.567	.781
Repetition (alt. assoc.)					
target	SHOE + TOAD			.712	.686
foil	SOCK + FROG			.594	.720

Note. p(c) = forced-choice performance; alt. assoc. = the alternative choice word was associatively primed.

they received cumulative feedback for the last 8 trials. Following their animacy decision, prime word(s) remained on the screen for an additional 500 ms before being replaced by the flash of the target word (see Figure 1). This was done so participants would not miss the target flash while responding to the prime word(s). Compared with the display for animacy matching, the prime words switched locations and were displayed in bold face during this 500 ms. In the passive priming condition, prime words appeared in bold face for 500 ms (hence, the 500 ms prior to target presentation were identical in both conditions). For passive priming, participants were instructed that prime words were a warning to prepare for the flash of the target word.

Following initial practice trials, participants were presented with a block of 72 perceptual identification trials. The purpose of this block was to find the duration of target flash at which performance was 75%. Appropriate durations averaged about 50 ms, although there were large individual differences, with times ranging from 25 ms to 117 ms. A staircase method was used to find the appropriate target duration during this threshold determination block. Participants were fully informed about the procedure. The words for the threshold determination block and practice trials were randomly selected (i.e., neither-primed). Prime relatedness was not introduced until the experimental trials.

Results

A note on analyses. One source of evidence concerning possible effects of priming on target perception is obtained from subtraction of neither-primed probability correct, p(c), from both-primed p(c), assessed statistically by an appropriate *t* test. Improved performance in the both-primed condition provides evidence of a perceptual enhancement. The idea is that, on average, preference effects are equated if both choices are primed and that priming might increase variability of evidence but would be unlikely to decrease variability. An increase in performance therefore provides relatively unambiguous evidence for a beneficial effect of priming on target perception. If a deficit is observed, it is evidence for the existence of increased variability with priming (preferential variability) or perceptual inhibition; such a result leaves open the possibility that perceptual enhancement exists, as long as the (negative) effect of preferential variability is large enough to overcome the perceptual enhancement.

A comparison of neither-primed to foil-primed involves equal effects on target perception, because in both cases the prime is unrelated to the target (i.e., in both cases the displays are identical up until the 2-AFC). Therefore, the difference between performance in these conditions is an indicator of preference effects. Similarly, the comparison of target primed to both primed involves equal effects on target perception because in both cases the prime is related to the target. Therefore, the difference between these is again an indicator of preference effects. One of our analyses combines these, assessing the effect of priming the foil with an *F* test holding target priming (or lack of target priming) constant (i.e., an ANOVA with two main effects: whether the target is primed and whether the foil is primed). A positive result of this test indicates the presence of preference effects, but the failure of the test does not disprove preference because the average direction of preference and preferential variability can potentially counteract one another. For example, if there is a slight preference against primed foils, the corresponding increase in performance can be offset by the decrease in performance associated with increased variability. Failing the combined *F* test, the separate comparisons (neither-primed vs. foil-primed and target-primed vs. both-primed) are individually checked with *t* tests to see if one of these comparisons in isolation suggests a preference effect.

A final analysis assesses whether, on average, preference is in favor of or against prime-related words (i.e., the direction of preference): The foil-primed p(c) is subtracted from the target-primed p(c), and an appropriate *t* test is used for statistical analysis. Although this difference includes any perceptual effects of target priming, it is useful in combination with the first 2 tests and, in particular, in comparisons of the active and passive priming conditions.

Passive priming results. The upper panel of Figure 2 and the passive probability-correct column of Table 1 show the accuracy results for the various conditions. (The predictions shown in Figure 2 are discussed in the next section of this article.) There was a difference between repetition priming and associative priming, $F(1, 54) = 6.89$, $p < .025$ that interacted with the four priming conditions, $F(3, 52) = 7.39$, $p < .001$. This difference reflects the fact that there was a deficit in the both-primed condition for repetition priming only. In addition, participants tended to choose

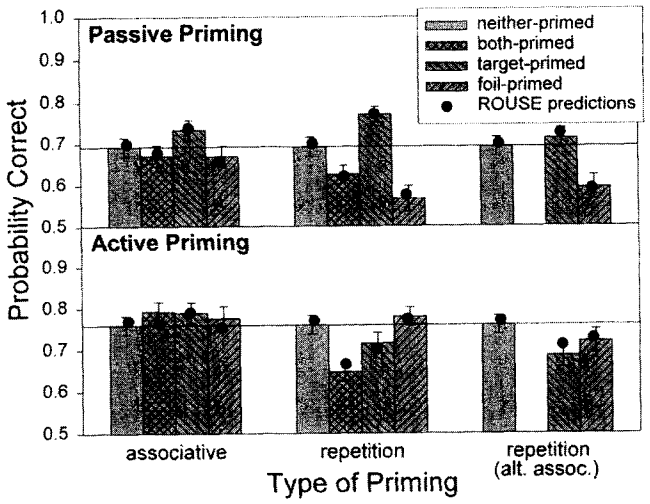

Figure 2. Accuracy results and predictions for Experiment 1. Error bars are ± *SEM*. Passive versus active priming was a between-subjects manipulation. For comparison with the three types of priming, the single neither-primed condition is displayed three times, with a level indicated by the horizontal line. ROUSE = responding optimally with unknown sources of evidence; alt. assoc. = the alternative choice word was associatively primed.

the related choice word, whether the relation was associative or repetition, although this effect was larger for repetitions.

For repetition priming, preferential variability, perceptual deficit, or both, was a factor: Performance in the both-primed condition was worse than the target-primed condition, $t(54) = 2.92, p < .005$ (see, *A note on analyses* above for an explanation of the statistics used to assess preference). In addition, there was a large overall preference effect, $F(1, 54) = 43.18, p < .001$; this consisted of a preference to choose the word that repeated a prime, $t(54) = 6.00, p < .001$.

For associative priming, performance in the both-primed condition was not different than in the neither-primed condition, $t(54) = 0.97, p = .34$. There was an overall preference effect, $F(1, 54) = 7.73, p < .01$, to choose the word that was an associate of the prime, $t(54) = 2.85, p < .005$.

When the target was repetition primed and the foil associatively primed, performance was higher than when the foil was repetition primed and the target was associatively primed, $t(54) = 3.16, p < .0025$. In other words, participants tended to choose the repeated word even if the alternative word was associatively primed.

Active priming results. Participants took an average of 3,173 ms with a p(c) of .732 in the animacy matching task, suggesting this was a difficult task involving considerable processing of the primes.

The lower panel of Figure 2 and the active probability-correct column of Table 1 show the accuracy results for the various conditions. There were differences between associative and repetition priming, $F(1, 51) = 19.95, p < .001$, that interacted with the four basic priming conditions, $F(3, 49) = 10.99, p < .001$. The difference was due to a deficit in both-primed performance (i.e., preferential variability or perceptual deficit) and a difference between the target-primed and foil-primed conditions, each of which only occurred for repetition priming. Surprisingly, participants tended to choose the choice word that was not a repetition of a prime. Thus, the direction of preference was opposite to that seen with passive priming.

Within repetition-priming conditions, performance in the both-primed condition was worse than the neither-primed condition, $t(51) = 4.62, p < .001$, suggesting that preferential variability outweighed any improvement in target perception caused by prime repetition. Although there was no preference effect according to the combined measure, $F(1, 51) = 2.41, p = .13$, the individual comparison of the both-primed condition with the target-primed condition, $t(51) = 2.97, p < .0025$, revealed evidence of a preference effect. In addition, the comparison of target-primed with foil-primed shows that the preference was against repeated words (target-primed lower than foil-primed), $t(51) = 2.14, p < .025$. Within the associative priming conditions, there were no differences across the four basic conditions, $F(3, 51) = 1.04, p = .39$.

When the target was repetition primed and the foil associatively primed, performance was lower than when the foil was repetition primed and target associatively primed, although this difference did not reach significance, $t(51) = 0.98, p = .17$. In other words, there may have been a slight preference to choose the nonrepeated word even if that word had been associatively primed.

Discussion

The first noteworthy result is found in the repetition conditions, for both active and passive prime processing: There was a substantial deficit in the both-primed condition compared with the neither-primed condition. There are two obvious hypotheses, either or both of which could be true: (a) The primes produce a deficit in perceptual processing of the target and (b) the increase in variability of evidence induced by the prime (i.e., preferential variability) outweighs any improvement in perceptual processing of the target (if there is any such gain). The ROUSE model presented in the next section explains the results in terms of Explanation b.

The second noteworthy result is the switch from a preference for repeated words with passive priming to a preference against repeated words with active priming. Using traditional priming tasks, which presumably include both preferential and perceptual effects, decreased performance with repetition (Humphreys et al., 1988) and orthographic–phonemic priming is occasionally observed (Dominguez & de Vega, 1997; Lukatela & Turvey, 1996; Lupker & Colombo, 1994; O'Seaghdha & Marin, 1997). Presumably, whatever is responsible for these deficits should apply at least as strongly to the case of repetition priming. Typical explanations for these deficits appeal to lexical suppression (e.g., Lupker & Colombo, 1994) or phonological competition (e.g., O'Seaghdha & Marin, 1997). The ROUSE theory presented next provides a unique perspective on these occasionally observed deficits by proposing that they are the product of a preference against prime-related words. Conditional on a key assumption concerning one parameter setting, the theory specifies the conditions in which negative preferences can be found. These turn out to be a subset of the cases in which participants more fully (i.e., actively) process the primes. In other situations, including those in which primes are processed to a lesser degree (i.e., passively), the theory predicts facilitation that is due to priming, as is more commonly observed.

Repetition priming is rarely studied in short-term word-identification priming, except with subthreshold prime presentations, because of concerns of strategic responding. However,

within a rapid serial visual presentation (RSVP) sequence of words, the effect of presenting a word upon its later re-presentation has been studied. This paradigm led to the observation known as "repetition blindness" (Kanwisher, 1987). In a typical repetition blindness experiment, participants fail to report the second presentation of a word. Sometimes cited as an example of repetition blindness, Johnston and Hale (1984) and Hochhaus and Johnston (1996) also found repetition deficits. These experiments are unique in that, similar to our Experiment 1, preference factors were controlled, resulting in the observation of a deficit in sensitivity. In these studies, participants were not required to respond to the prime word, and similar to our passive priming results, they observed a bias in favor of repeated words.[1] Any simple interpretation of such bias, however, must explain the reversal of this tendency in our active priming condition.

Unlike the results of Masson and Borowsky (1998), we did not find an associative–semantic enhancement when preference factors were controlled. However, associative effects in our studies might have been weakened by the use of two primes, a hypothesis shown to be correct in Experiment 4. More generally, our repetition and associative priming results provide strong evidence that priming produces preference effects but no direct evidence that priming produces perceptual changes in target processing. Thus, the ROUSE model presented in the next section is a model of preference effects.

With passive priming, the tendency to choose prime-related words occurs in associative priming as well as repetition priming (i.e., a positive preference). With active priming, this tendency is reversed for repetition priming. Yet for both active and passive priming, there is a deficit in the both repetition primed condition. This pattern of results presents a complex set of interactions that seems at first glance to defy simple explanation. We see next that this assessment is incorrect, because a rather simple Bayesian model of decision making for this task provides a coherent account of the pattern of results. The theory accounts for the data assuming only preferential factors are involved in priming.

ROUSE

Our theory can be broken into two parts. In the first part, we explain how a preference for prime-related words can arise through source confusion (the *Unknown Sources of Evidence* section). In the second part, we show how an optimal decision process can remove or even reverse this preference through the discounting of features known to have been in a prime word (the *Responding Optimally* section). We assume that with active priming, preference removal is adequate or even excessive, whereas with passive priming preference removal is insufficient.

Unknown Sources of Evidence

Preference effects occur because features shared between choice words and prime words can be activated by either source, and it is unknown which source(s) provided activation. A choice word containing prime features thereby tends to be favored. One can think of this activation of choice word features as a kind of preactivation. However, unlike the preactivation in spreading activation theories (Anderson, 1983; Collins & Loftus, 1975), we assume preactivation only applies to shared features between prime and choice word, and as such only part of the representation might be affected (to the degree that features are shared). This preactivation does not affect the process of target perception and is, therefore, preferential in nature (i.e., preactivation does not make it any more likely that the target presentation will activate features). To be more precise, the theory holds that the primes, the flash of the target, and general visual noise are all independent sources of feature activation.[2] A preference for prime-related words arises because of a failure to distinguish the various sources of activation (see Johnson, Hashtroudi, & Lindsay, 1993, for a review of source-monitoring phenomena).

We show that variability in prime activation is the basis for the deficits observed in the both-primed repetition conditions (as well as the deficits in the average of the target-primed and foil-primed conditions). On average, the target and foil choices contain an equal number of prime-activated features. However, because of the probabilistic nature of activation, either the target or the foil may contain more prime-activated features on a given trial, producing decision noise that reduces performance.

Similar to the REM-Implicit (REMI) model that Schooler, Shiffrin, and Raaijmakers (2001) developed for long-term repetition priming, the ROUSE model assumes each word consists of a vector of lexical-semantic features.[3] With 2-AFC testing, the participant must choose between two words, and, therefore, only the features contained in these two words need to be considered. Furthermore, because a feature common to both choices is assumed to be either jointly activated or inactive in both, the evidence for both is equal. Therefore, a feature shared by both alternatives is completely nondiagnostic and is not considered. Considering only diagnostic features, the result is separate vectors

[1] In the same–different paradigm, bias and sensitivity are assessed using *same* responses to same words as a hit rate and same responses to different words as a false-alarm rate. This is done separately for the case of primed same–different test words and unprimed same–different test words. Therefore, a change in sensitivity between primed and unprimed words is analogous to a change in performance between neither-primed and both-primed using the 2-AFC testing procedure. Likewise, a change in bias between primed and unprimed words is analogous to assessing the presence and direction of preference in 2-AFC testing.

[2] Because the ROUSE theory supposes that activation from the primes combines with activation from the target flash, it is tempting to place ROUSE in the class of compound-cue theories. However, instead of providing extra cues beyond those contained within the choice words, preactivation by the primes only matters in so far as it overlaps with the representations of the choice words. In a more open-ended paradigm such as lexical decision or the naming version of perceptual identification, the analogy to compound-cue theories is more sensible although we have not fully worked out the application of ROUSE to these paradigms at this time. In truth, the ROUSE theory combines aspects of both spreading-activation and compound-cue theories but does not clearly fall into either category.

[3] In the REMI model of Schooler et al. (2001), a match between a perceived feature and the lexical–semantic feature of a choice word provides evidence in favor of that choice word, and a mismatch is evidence against that choice word. Similarly, in ROUSE, an active ON lexical–semantic feature provides evidence in favor, and an inactive OFF feature is evidence against. The only noteworthy difference between the models arises from considerations of temporal order. In ROUSE, it is assumed that sources of evidence operate independently, whereas in REMI a single perceived vector is loaded from the combined sources and this results in dependencies. It is unclear at this time if this difference produces testably different results.

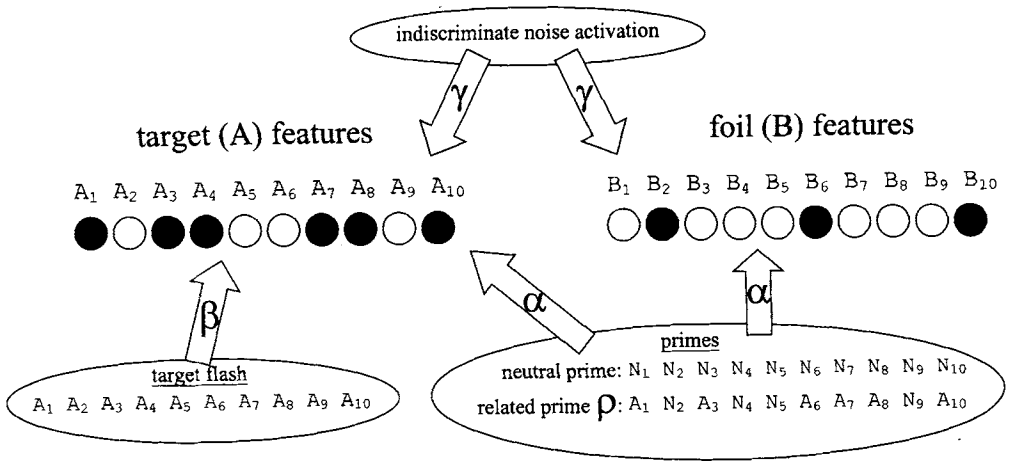

Figure 3. The responding optimally with unknown sources of evidence (ROUSE) theory assumes three sources independently activate target and foil features. With the source(s) of activation unknown, the target presentation may be the source of any feature activation; therefore, any active feature is taken as evidence in favor of the word to which it belongs. The values β (target flash), γ (noise), and α (primes) are the probabilities that each source has activated a feature and that the feature remains active (and indistinguishable from target activation) until the decision process is initiated. Noise is assumed to arise from various sources, including mistaken perception of letters that is due to the pattern mask. Similarity between prime and choice word (depending on the condition) is assumed to be 0 for unrelated primes, the parameter ρ for related primes, and 1 for identical primes (i.e., repetition priming). In the simulations, lexical entries are represented by 20 features. These features might contain orthographic, phonemic, or semantic information.

of unique features for the target and foil; increasing similarity between the choice words correspondingly reduces the number of diagnostic features. We used a vector length of 20 in our simulations (although in Figure 3 we used 10 features to reduce clutter). We made the simplifying assumption that each feature is binary, existing in an ON or OFF state (ON features in a choice may be thought of as apparently perceived features that match a corresponding feature in that choice). Perfect perceptual processing would result in the turning on of all target features (although at the same time, some foil features could be turned on by the primes or by visual noise). Thus, ON features provide evidence in favor of the choice word to which they belong and OFF features provide evidence against the choice word to which they belong. If the same number of features exists in both choice words, and all ON or OFF features provide the same evidence, an optimal decision rule would be to choose the word with the largest number of ON features. At the start of a trial, all features are OFF. The features are turned ON by three sources of evidence (see Figure 3):

1. *The target flash.* With Probability β, target features are turned ON (β depends on flash duration).
2. *Visual noise.* With Probability γ, any feature in either choice word is turned on.
3. *The primes.* With Probability α, any feature that is shared between a prime and either choice word is turned on.[4] (A feature can be turned on by more than one source on a given trial; it is on if at least one source turns it on.)

Each of the three activation probabilities (α, β, and γ) represents the joint probability that the source has turned on a feature and that the feature has remained on until the time of the decision process. A more detailed account could replace these probabilities with activation and deactivation parameters. Because the features produced by prime activation eventually decay with time or intervening events, they cease to cause confusions. ROUSE is a model of short-term priming in that priming arises through activation and, therefore, is not expected to exist over extended durations.

Because ROUSE is a feature-based model, the similarity between any two words can be manipulated in a principled fashion through the proportion of shared features. In some conditions, two words are allowed or assumed to share some but not all features. When this is the case, a parameter, ρ, determines the probability that a feature is shared by both words. In different studies, prime words may be similar to choice words, and, in others, choice words may be similar to each other. For example, in Experiment 1 we assumed that the associative priming involves the sharing of semantic features between primes and primed choices, and, therefore, we let ρ be the probability that a feature of an associatively primed choice word is shared with the prime. For simplicity, we assumed that unrelated words share no features ($\rho = 0$) and repetitions share all features ($\rho = 1$). Associative priming (and in later studies, orthographic–phonemic priming) involves some proportion of shared features that must be determined on each simulated trial according to the probability ρ. It should be emphasized that for studies like those in the present article in which the features of the primes and choices are readily available to the participant, the participant need not pay heed to the value of ρ

[4] If our model was more thoroughly developed, we would lay out more precisely the factors that ought to affect the values of these three parameters. For example, β might vary with aspects of target presentation and duration, γ might vary with mask and presentation characteristics, and α might vary with time between prime and target flash. Such details are not needed at the present stage of development of the ROUSE model.

because it is clear to the participant which features overlap and which do not. (Note: A feature common to the two choice words is ignored in the decision process. This assumption follows from an optimal decision process because a shared feature, whether ON or OFF, provides the same degree of evidence toward both choices.)

Responding Optimally

If not somehow countered, confusion concerning the source of features activated by a prime would lead to a strong preference for a choice word containing those features. We believe participants implicitly or explicitly make a choice decision in something approaching optimal fashion. To do this, evidence in favor of a choice word that is due to that word having an activated feature must be lowered (i.e., discounted) when that feature had also been in a prime. Such discounting is appropriate because a prime rather than the target might have been the source of activation. For example, suppose that the features are letters, that a *T* is perceived in the first letter position, that the choice words are TOWN and SEAM, and that neither prime word has a *T* in first position. The perceived *T* provides good evidence that the flashed word was TOWN because the *T* could only have come from the flash or visual noise. On the other hand, suppose that TENT was presented as a prime word; in this case the perceived *T* in the first position could have come from the prime, and the evidence in favor of the choice TOWN should count somewhat less. In an optimal setting, such discounting ought to apply to all features that are shared between prime and choice word; in practice, discounting should apply only to features the participant knows are also present in the primes. In the present tasks, the primes are presented well above threshold, even in the passive priming condition, so we assume that the system, participant, or both, knows which features are shared between primes and choice words, and appropriate discounting can be carried out on all of these. It becomes evident that this discounting of features can remove and even reverse the preference to choose prime-related words. (In *Theories of Discounting,* we discuss similar notions of discounting that have been proposed for priming within the social cognition field.)

At the time of decision, the ON–OFF state of each feature is assessed to calculate the odds that a given choice is the target. The word with odds that are greater than one is chosen. For this calculation to be carried out, the activation probabilities (α, β, and γ) must be used. These values are not readily available to participants and must be estimated (ρ need not be estimated because it reflects the known properties of the words). The different preference results for active and passive priming are predicted by ROUSE on the basis of the following assumption: the estimate of prime feature activation, labeled α' is low ($\alpha' < \alpha$) with passive priming and too high ($\alpha' > \alpha$) with active priming.

It is presumed that estimation of all three parameters takes place during the many practice trials (in most experiments there are at least 80 practice trials prior to collecting data) and in the initial phase of the experimental trials. Our use of feedback on every trial may have been a crucial factor in allowing participants to arrive rapidly at estimates of these parameters. Future investigations of ROUSE may more precisely address the time course of learning for these feature-activation estimates. In any event, a series of simulations of ROUSE revealed that misestimates of β and γ do not change the pattern of results across the priming conditions (they only change the overall level of performance). Therefore, we set the estimates of β and γ to their true values.

The appropriate level of evidence follows from a Bayesian calculation. The odds for the target over the foil are given in Equation 1 (*A* refers to the target word and *B* refers to the foil word). Assuming equal priors, as is appropriate for our experimental design, the normative decision is to choose the target *A* if the odds are greater than one and to choose the foil *B* if the odds are less than one. If the odds are equal to one, the target is chosen with a probability of .5.

$$\Phi\left(\frac{\mathrm{A}}{\mathrm{B}}\right) = \frac{p(\mathrm{A,\ B\ activation\ pattern}|\mathrm{A\ is\ target,\ B\ is\ foil})}{p(\mathrm{A,\ B\ activation\ pattern}|\mathrm{B\ is\ target,\ A\ is\ foil})}. \quad (1)$$

In an optimal treatment of feature activation, an active feature provides evidence in favor of a choice word as target, but an inactive feature provides evidence against a choice word as target. Because both active and inactive features provide useful evidence, all features in both choice words contribute to the likelihood calculation, except features that are common to both choices. If one assumes that each feature yields an independent source of evidence and breaks Equation 1 into two separate products for the features of each choice word, the following equation is generated:

$$\Phi\left(\frac{\mathrm{A}}{\mathrm{B}}\right) = \frac{\prod_{i=1}^{N} \frac{p(\mathrm{V}\{\mathrm{A}_i\}|\mathrm{A\ is\ target})}{p(\mathrm{V}\{\mathrm{A}_i\}|\mathrm{A\ is\ foil})}}{\prod_{j=1}^{N} \frac{p(\mathrm{V}\{\mathrm{B}_j\}|\mathrm{B\ is\ target})}{p(\mathrm{V}\{\mathrm{B}_j\}|\mathrm{B\ is\ foil})}}, \quad (2)$$

where V (A_i) represents the value of the *i*-th feature of A and takes on one of two values, which denote the ON and OFF states of activation; this is similar for V(B_i). The product in the numerator is termed the likelihood ratio for the target word and is based on evidence coming from the *N* features of the target *A;* the product in the denominator is termed the likelihood ratio for the foil word and is based on evidence coming from the *N* features of the foil *B.* Thus, the choice of *A* if the odds in Equation 1 are greater than one is equivalent to the choice of *A* if its likelihood ratio in Equation 2 is greater than that for *B.*

There are only four possible evidence ratios that could appear in the product terms found in Equation 2 (as shown in Figure 4), depending on whether a feature is ON or OFF and whether a feature is known to have appeared in the prime(s) (i.e., whether a prime is a potential source of activation). For example, consider an OFF feature that did not appear in the prime(s). Assuming that the feature is part of the target, the target flash or noise could have been a source of activation. Because the feature is OFF, both of these sources must have failed so the probability is $(1 - \gamma)(1 - \beta)$. Assuming that the feature is not part of the target, only noise is a potential source of activation, so the feature is OFF with probability $(1 - \gamma)$. This leads to the ratio seen in the upper left panel of Figure 4 (which is equal to $1 - \beta$). A related calculation produces the same result, $1 - \beta$, in the lower left panel of Figure 4. In other words, prime and noise activation are irrelevant to the evidence provided by OFF features; only target activation matters.

The lower right panel in Figure 4 gives the evidence for an ON feature that appeared in a prime; such a feature is termed *dis-*

Appeared in Prime(s)	State of Activation: OFF	State of Activation: ON
NO	$\frac{(1-\gamma)(1-\beta)}{(1-\gamma)} = (1-\beta)$	$\frac{1-(1-\gamma)(1-\beta)}{1-(1-\gamma)}$
YES	$\frac{(1-\gamma)(1-\alpha')(1-\beta)}{(1-\gamma)(1-\alpha')} = (1-\beta)$	$\frac{1-(1-\gamma)(1-\alpha')(1-\beta)}{1-(1-\gamma)(1-\alpha')}$

Figure 4. The feature-likelihood ratios that might appear in the numerator or denominator product terms of Equation 2. This is a 2 × 2 contingency that depends on whether a feature is active or inactive and in the primes or not in the primes. It is assumed that the primes (and their features) are known. The numerator of each ratio is conditional on the feature existing within the target (and not the foil), and the denominator is conditional on the feature existing within the foil (and not the target). Features that appear in both the target and the foil provide no discriminating information and are not considered in the decision process. As in Figure 3, α, β, and γ, respectively refer to the probabilities of feature activation by the primes, target, and noise. In theory, estimates of these probabilities should appear in the feature likelihood ratios. In practice, misestimates of β and γ do not change the pattern of results across priming conditions and are therefore set to their true values. Only the estimate of α (labeled α') differs from its true value.

counted because its evidence ratio is less (i.e., closer to one) than if it had not appeared in a prime (the term in the upper right panel). It is the estimate of prime activation, α', that determines the level of discounting.

The relative size of prime activation, α, compared with the estimate of prime activation, α', produces the direction of the preference: for example, whether the target-primed condition is better than the foil-primed condition. If participants choose optimally (i.e., $\alpha' = \alpha$) and the number of diagnostic features turned on by each of the sources of activation is sufficiently large, feature evidence from primes is discounted properly, and on average there is not a tendency to choose or not choose words related to a prime. (Note, however, that the overall performance of these preference conditions taken together is still predicted to be lower than performance in the neither-primed conditions, because variability exists in the number of prime-activated features; see the explanation in the second paragraph below.)

If participants are conservative and overestimate the effect of the prime (i.e., $\alpha' > \alpha$) and if the number of diagnostic features turned on by each of the sources of activation is sufficiently large, words that are not related to a prime tend to be chosen. Such a situation is what we assume exists with active priming. If participants are less aware of the primes as a potential source of activation, they may underestimate the effect of the primes (i.e., $\alpha' < \alpha$). If they do so, they tend to choose prime-related words. This is the situation we assume holds with passive priming. A pictorial explication of these arguments concerning preference effects is given in the *ROUSE and Discounting* section.

Next, consider predicted performance in the both-primed conditions. On average, the number of ON features that are shared with the primes is the same for the two choice words (i.e., preference is controlled). However, there is variability in these numbers, so that sometimes one and sometimes the other choice word is favored, purely by chance. This chance process adds noise to the decision, decreasing predicted performance compared with the neither-primed condition. The size of this variability effect depends on the values of α and ρ, with increasing variability for larger values of either parameter. This means that the deficit is largest for repetition priming, for which all features are shared.

This effect of variability of evidence arising from prime activation is best described as preferential because it occurs even when there is no change in the perception of the target. However, primes have a second effect on predicted performance, an effect that might be described either as perceptual or preferential, depending on one's perspective: Features that are turned on by the prime are unavailable to be turned on by the target. This is an example of performance being harmed by a prime, because of blocking. Considering the situation from the point of view of the features that are eventually ON, one can describe the harm as perceptual, because perception is blocked from producing distinguishing evidence. For example, if the parameter α equals 1.0, then in the both-primed condition, every feature of both choice words is turned on by the primes; now the flash of the target provides no additional information, and performance is at chance. One could argue that no perception has occurred. On the other hand, one could argue that perception is unaffected by the primes but that the evidence provided by perception is overwritten by the prime features; in this case, one might prefer to describe the blocking effect as preferential. Fortunately, this debate in the context of the present studies is of little consequence: It turns out that the fit of the model to the data resulted in very low estimated values of α, β, and γ. With low values for these parameters, a feature turned on by a prime is rarely also turned on by the target flash, so the effect of blocking turns out in practice to be of negligible importance; this being so, it is not critical to decide whether blocking should be thought of as a perceptual or preferential effect. In particular, the both-primed deficit predicted with ROUSE is almost entirely due to preferential variability rather than blocking.

The associative case is similar to the repetition case, differing only in having relatively few features shared between prime and target, or between prime and foil ($\rho < 1$). This change lessens the effects of priming generally. It should be noted again that although we need to estimate the value of ρ to fit the model, the participant need not estimate the value of ρ: Both the prime features and choice word features are available to the participant on each trial (assuming the participant pays attention to the above threshold primes), so the shared features are identifiable and do not have to be estimated by the participant. Procedures for producing simulations with the ROUSE model appear in Appendix A.

Theories of Discounting

The idea of discounting evidence is far from unique to the present treatment. A particularly relevant example arises in the area of evaluative priming. Similar to the preference reversal observed in Experiment 1, social psychologists have observed priming reversals in evaluative judgments. Lombardi, Higgins, and Bargh (1987) instructed participants to construct sentences using synonyms of "persistent" versus "stubborn." After this task, par-

ticipants gave a one-word label to a neutral description of a target person. Participants who later remembered their constructed sentences tended to use a label that was the opposite of the synonyms they received (i.e., if their sentence included "determined," they rated the target person as "stubborn"). Conversely, participants who could not remember their sentences tended to use a label that was similar to the synonyms.

Experiments in which such *contrast* (i.e., a preference against prime-related words) and *assimilation* (i.e., a preference for prime-related words) priming occur have been explained by proposing that participants may or may not attempt to remove the influence of prime items (Martin, 1986; Schwarz & Bless, 1992) or more generally to remove what is believed to be a contaminating influence of certain noticed mental events (e.g., Wegener & Petty, 1995). In assimilation priming, the effect of the prime is realized in full, whereas in contrast priming the effect of the prime is removed, which results in the preference against prime-related words. One theory of discounting holds that source similarity between prime and target is of particular importance (Mussweiler & Neumann, 2000). With externally presented targets, source-similarity theory predicts that externally presented primes lend themselves to discounting, whereas internally generated primes are less likely to be discounted. Indeed, Mussweiler and Neumann found that internally generated primes (e.g., antonym generation task) led to assimilative priming, whereas a simple presentation of these same primes led to contrast priming.

We believe, however, that these theories of discounting tend to produce very explicit and broadband effects. We show in the remaining studies that the discounting mechanism at work in the present situation is much more subtle and perhaps less available to awareness than that entailed by such theories.

ROUSE and Discounting

To explicate the discounting mechanism in ROUSE that removes or even reverses preference, we present Venn diagrams in Figure 5 for the target-primed and foil-primed conditions.[5] In this figure, we illustrate how the average direction of preference (or lack thereof) arises from two offsetting factors: When priming is used, source confusion and discounting can combine to produce either evidence gains or losses. We do not use this figure to explain the both-primed deficits in performance that are caused by priming because these are due to variability and the figure depicts averages. Figure 5 illustrates the separate evidence situations for the target and foil when similarity between prime and primed choice is intermediate (e.g., $0 < \rho < 1$, as would be the case for associative or orthographic–phonemic priming). The left-hand set of panels shows the situation without discounting (i.e., $\alpha' = 0$), and the right-hand set of panels shows the situation with optimal discounting (i.e., $\alpha' = \alpha$).

In general, features can provide one of three levels of evidence. Features that are OFF provide evidence against a choice word regardless of prime matching (as seen in the OFF column of Figure 4); these are represented as the black regions in Figure 5. Features that are ON provide evidence in favor of the word to which they belong if they do not also appear in a prime (as seen in the upper right cell of Figure 4); the white regions in Figure 5 represent this situation. Last, ON features that also appeared in a prime provide a discounted level of evidence in favor of a choice word (as seen in the lower right cell of Figure 4); this situation is represented by the gray regions in Figure 5: The higher the level of discounting is (i.e., the higher is α'), the darker the shade of gray is, and the lower the evidence is provided by such features.

In Figure 5, each panel shows three circles representing features activated from each of the three sources (as labeled in the upper left-hand panel). Technically, these circles should overlap, to represent that a given feature may be activated by more than one source. However, the estimated values of the three parameters for our various studies were quite low, so that the regions of overlap would be quite small in practice. Therefore, we separated the circles because it makes the figure less cluttered and makes the associated logical inferences easier to follow; for the present parameter values, errors of inference caused by this simplification are negligible.

As can be inferred from Equation 2, performance corresponds to the product of the evidence from each feature for the target compared with the same product for the foil; in Figure 5 this may be thought of as the average "lightness" of the panel for the target evidence (the upper Venn diagram for a given condition) compared with the average lightness of the panel for the corresponding foil evidence (the lower Venn diagram for a given condition). To determine the degree of preference in the target-primed and foil-primed conditions, one must compare the lightness advantage of the top cell over the bottom cell in one column with the lightness advantage of the top cell over the bottom cell in an adjacent column (admittedly, this is not easy to do).

In the left-hand four panels of Figure 5, the situation is illustrated for the case when there is no discounting (i.e., $\alpha' = 0$); therefore, all active features are displayed in white. These panels make it clear how, without discounting, the prime-activated ($\rho\alpha$) regions result in a preference for prime-related words. This result would be caused by unknown sources of evidence, in the absence of discounting: In the target-primed condition, there is additional evidence in favor of the target due to the prime-activated features (more whiteness in the upper panel), and in the foil-primed condition, there is additional evidence in favor of the foil due to the prime-activated features (more whiteness in the lower panel). It is clear that this results in a preference for the prime-related choice, which in turn ensures that performance for the target-primed condition will be higher than performance for the foil-primed condition.

Next, we consider the action of responding optimally, which is exemplified in ROUSE by the discounting of features known to exist in the primes. This is depicted in the optimal discounting (i.e., $\alpha' = \alpha$) set of four diagrams at the right in Figure 5. Discounting only pertains to activated features that also appear in the primes. In the case of prime-activated features, it is known that all of these features appeared in the primes and, as such, they are all discounted (i.e., the $\rho\alpha$ regions are now gray). Target- and noise-activated features are only discounted if they are also shared with the primes as determined by ρ (i.e., $\rho\beta$ and $\rho\gamma$ represent subsets of target- and noise-activated features that are discounted); therefore,

[5] The concept of discounting may seem clear enough that the reader sees no need for additional explanation. However, the model's predictions for some of the subsequent studies are far from intuitively clear, and the next two sections help the reader to follow the subsequent developments.

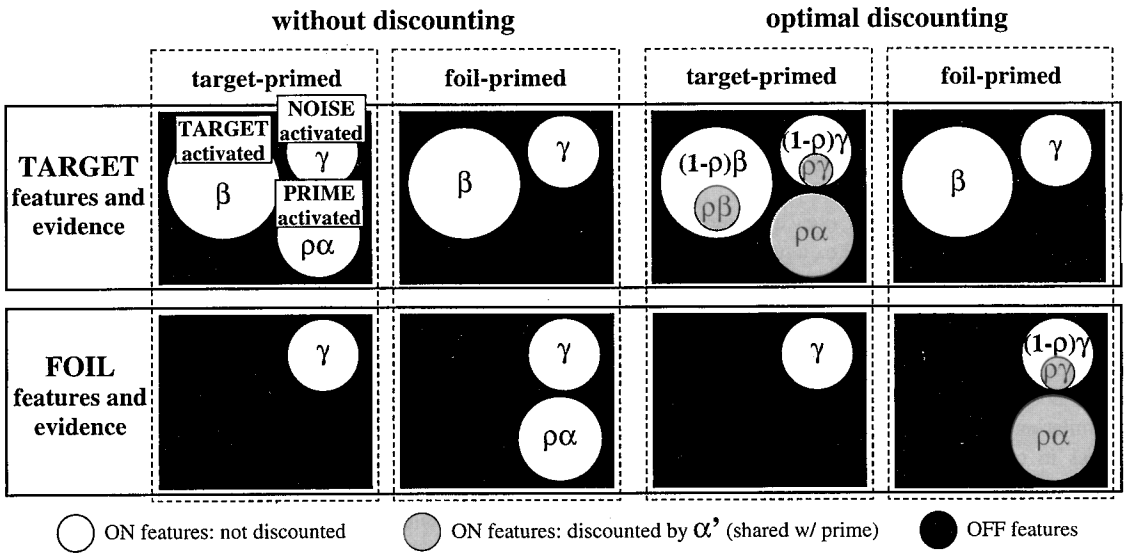

Figure 5. Venn diagrams comparing target and foil evidence without discounting (left) with the situation with optimal discounting (right) for an intermediate level of prime similarity ($0 < \rho < 1$). With optimal discounting, the average evidence is equated to the situation without priming (not shown) because of the offsetting factors of evidence gains from prime-activated features and evidence losses from discounted target- and noise-activated features; it follows that the target-primed and foil-primed conditions are equated. These diagrams portray average evidence and therefore do not show performance deficits due to increases in variability with priming (see Figure 6 and the attendant discussion). The target-, noise-, and prime-activated regions contain features that are active as a result of the labeled source. To simplify the situation, features activated by more than one source are not shown, and the circles, therefore, do not overlap (a fairly accurate approximation when the feature activation probabilities are low). As seen in Figure 4, only three levels of evidence exist for each feature. OFF features (black background) provide evidence against the choice word to which they belong (ratio less than 1), whereas ON features not contained in a prime (white regions) provide strong evidence in favor of the choice word to which they belong (ratio greater than 1). Discounted features (gray regions) are active features that appeared in a prime and therefore provide weak evidence in favor of the choice word to which they belong. The parameters and products of parameters contained within each region correspond to the proportion of total features that can be expected to exist within that region. For example, prime-activated features must exist in a prime, with probability ρ, and be activated by the prime, with probability α, resulting in the joint probability $\rho\alpha$. In the case of optimal discounting, the target- and noise-activated regions are broken into those features that do not exist in a prime and therefore are not discounted (probability $1 - \rho$) versus those features that exist in a prime and are discounted (probability ρ). Performance is determined by choosing the word with the greater product of evidence, as dictated by Equation 2. This approximately corresponds to choosing the word that has a panel in a column that is "lighter" in color.

these discounted features are represented as gray regions within the β and γ circles.

Close inspection of Figure 5 reveals that discounted prime-activated features ($\rho\alpha$) influence preference in a qualitatively different manner than discounted target- ($\rho\beta$) and noise-activated ($\rho\gamma$) features. As shown in Figure 4, the evidence provided by discounted features is bounded below by 1.0, and as such even heavily discounted features (i.e., $\alpha' = 1$) provide more evidence in favor of an alternative than they would if these features had remained OFF. This means that in isolation, discounted prime-activated features would result in a preference for prime-related words despite discounting (i.e., a gain in evidence). A preference removal or reversal can only be obtained from the discounting of target- and noise-activated features, which results in a real loss of evidence (i.e., part of the β and γ circles have turned from white to gray). Thus, the gain in evidence (albeit discounted) provided by prime-activated ($\rho\alpha$) features is offset by the loss of evidence provided by discounting target- ($\rho\beta$) and noise-activated ($\rho\gamma$) features.

This reasoning motivates the prediction of equal performance in the target-primed and foil-primed conditions when discounting is set to its optimal (i.e., $\alpha' = \alpha$) level (performance will still be lowered relative to the neither-primed case because of extra variability with priming). If discounting is less than optimal (i.e., $\alpha' < \alpha$), a condition we assume holds in the case of passive priming, the gain of evidence that is due to prime-activated features outweighs the loss that is due to discounting target- and noise-activated features, which results in a preference for prime-related words. If discounting is excessive (i.e., $\alpha' > \alpha$), a condition we assume holds in the case of active priming, the loss of evidence that is due to discounting target- and noise-activated features outweighs the gain in evidence that is due to prime-activated features, which results in a preference against prime-related words (i.e., active priming).

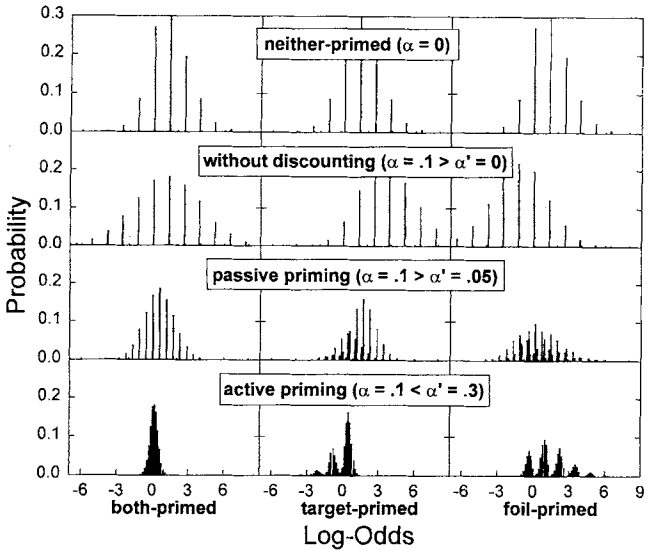

Figure 6. The distribution of predicted log-odds (log of Equation 1) in the four priming conditions with repetition priming (ρ = 1) for three different levels of discounting (with α = .1, setting α' = 0, α' = .05, and α' = .3). In addition to changes in average evidence, this figure portrays variability and performance deficits due to increases in variability with priming. The neither-primed condition is repeated three times in the top row for convenience of comparison (this condition is unaffected by discounting). The other parameters are as follows: N = 20, γ = .02, β = .05.

These intuitive predictions, however, turn out to be accurate only when sufficient numbers of features are available for discounting—numbers that were approached in our model applied to the data only in the case of repetition priming with dissimilar choice words. For lower values of prime similarity or tests with similar choice words, the direction of preference in the active priming case was not what this reasoning suggests, for reasons taken up in the sections *ROUSE and Prime Similarity* and *ROUSE and Choice-Word Similarity.*

Distributions of Odds

The Venn diagram represents an attempt to provide insight into the trade-off of factors that govern performance by depicting mean numbers of different types of activated features and their associated evidence values. An alternative and more precise way to illustrate the working of the model is through graphs of the distributions of odds (in favor of the target). These distributions depend on the probabilities of obtaining various numbers of ON and OFF features in the different conditions. A detailed description of how to produce these distributions can be found in Appendix A.

The distribution of the odds is highly skewed (because of the multiplication of probabilities), so for clarity we plot the distribution of the log-odds. In viewing the graphs in Figure 6, recall that a correct decision is made if the odds of Equation 1 are greater than 1.0, equivalent to the log-odds being greater than zero. The distributions are also quite discrete because only certain combinations of ON and OFF features are at all likely. Figure 6 shows distributions for the four basic conditions (neither primed, both primed, target primed, and foil primed) for repetition priming (ρ = 1) with typical activation parameters (β = .05, γ = .02, and α = .1), for three values of α': 0, .05, and .3 (corresponding to no discounting, discounting that is too small in light of the actual value of α, and discounting that is too large in light of the actual value of α). The top row of panels repeats the neither-primed distribution three times, for ease of comparison with the panels below. Otherwise, the distributions in the first column are for both primed, in the second column for target primed, and in the third column for foil primed. Table 2 provides some summary statistics concerning these distributions.

Looking first at the case α' = 0 (row 2), one can see that priming both alternatives (column 1) adds noise and causes the distribution to spread compared with the neither-primed condition; although the mode remains at the same position, the extra variability causes more of the distribution to fall below zero, reducing performance. Priming only the target or only the foil (columns 2 and 3) adds some variability, but the primary effect is to shift the log-odds in favor of the primed choice. The case of α' = .05 (row 3) includes discounting, which lessens the evidence value for the features that are in the primes, squeezing the log-odds toward zero in all three primed conditions. Overall performance is as expected, with a drop for both primed and a preference for the primed alternative (compared with α' = 0, the preference is diminished). The case α' = .3 (row 4) discounts much more strongly, shrinking the distributions severely, but still produces a both-primed deficit. The fact that too much discounting occurs in this case reverses the direction of preference; the unprimed choice tends to be chosen, which improves performance when the foil is primed and harms performance when the target is primed.

For the target-primed and foil-primed conditions with discounting, the ON features of the unprimed alternative provide strong evidence, whereas the ON features of the primed alternative are discounted; this results in the extra variation producing more discrete points in the distributions compared with the situation without discounting. For the both-primed condition, all ON features are discounted in both alternatives, which results in the same degree of variation (i.e., number of spikes in the distribution) compared with the situation without discounting. Furthermore, when both alternatives are primed, performance is unaffected by

Table 2
ROUSE Log-Odds Summary Statistics

	Priming condition			
α'	Neither primed	Both primed	Target primed	Foil primed
0				
M	1.26	1.14	3.69	−1.27
Mdn	1.29	1.29	3.87	−1.29
SD	1.67	2.81	2.27	2.36
.05				
M	1.26	0.50	1.33	0.43
Mdn	1.29	0.57	1.13	0.16
SD	1.67	1.24	1.24	1.66
.3				
M	1.26	0.14	−0.02	1.41
Mdn	1.29	0.16	0.31	1.13
SD	1.67	0.34	0.86	1.47

Note. Statistics are for Figure 6 distributions.

the direction of preference that is induced by misestimates of α (i.e., p[c] independent of α'). Table 2 shows that the mean and median log-odds of the both-primed condition decreases as α' increases but that the standard deviation decreases correspondingly, which results in no change in performance. In other words, because every ON feature is discounted in the both-primed condition, changes in α' simply result in a rescaling of the log-odds distribution centered on the normative decision criterion of 0.

Setting the Value of γ

We discovered when fitting ROUSE to the data from the various studies in this article that the probability of noise-activated features, γ, was usually estimated to be very low and that the fits were seldom harmed if the value was set to some fixed value close to zero. ROUSE is capable of producing both-primed deficits and changes in the direction of preference for much higher values of γ (which require higher values for the other parameters), but unless γ is set to a small value, ROUSE has trouble also predicting changes in preference direction as similarity changes (the explanation is related to factors discussed in the later sections *ROUSE and Prime Similarity* and *ROUSE and Choice-Word Similarity*). For these reasons, the value of γ was not estimated but was fixed at .02 throughout this article.

Setting the Vector Length

Generally speaking, we found that the vector length (above a certain minimum below which preference removal breaks down—see the *ROUSE and Choice-Word Similarity* section for a discussion of this breakdown) was a scaling factor primarily determining the length of time needed to carry out simulations rather than determining the pattern of predictions: Longer lengths (such as 100) produced similar predictions to those for shorter lengths, once suitable modifications were made to the values of the other parameters. A length of 20 was long enough to enable preference removal but was short enough to allow parameter fitting to be carried out in reasonable time, and this value was used throughout this article.

The ROUSE Model Applied to Experiment 1

Fitting the ROUSE model to the active and passive priming results in Experiment 1 required estimation of the four parameters α, α', β, and ρ (γ was set to .02; not estimated). Because the same words were used for both the active and passive groups, we required the parameter estimation program to use a single associative prime similarity parameter, ρ, for both groups, whereas the other parameters were separately estimated for the two groups. Averaging across 20,000 simulations per condition to obtain predictions for a given set of parameters, the parameters were assigned values that minimized the error between predictions and observations. The error measure used was the sum of the chi-squares from each condition, in which chi-square was calculated using the normal theory maximum likelihood method (see Curran, West, & Finch, 1996, for a comparison of various methods for calculating χ^2). For each condition i, containing N_i observations, this chi-square calculation method compares the log-likelihood of the constrained model, L_C, to the log-likelihood of the unconstrained estimate, L_U, using the observed probability, o_i, and predicted probability, p_i.

$$L_C = [o_i \log(p_i)] - [(1 - o_i) \log(1 - p_i)] \quad (3)$$

$$L_U = [o_i \log(o_i)] - [(1 - o_i) \log(1 - o_i)] \quad (4)$$

$$\chi_i^2 = -2N_i(L_C - L_U). \quad (5)$$

Table 3 gives the parameter values that produced a best fit to the data from Experiment 1. The predictions are given in Figure 2 as the dots on each bar. The best-fitting values as well as the observed means, the standard deviations of the observed means, and the numbers of observed data points can be found for all experiments in Appendix B.

Table 3
Best-Fitting ROUSE Parameters

Probability parameter	Experiment 1		Experiment 2		Experiment 3		Experiment 4	
	Passive	Active	Passive	Active	Passive	Active	Passive	Active
α (prime actual)	.073	.085	.105	.110	.112	.090	.379	.167
α' (prime estimate)	.054	.152	.075	.152	.097	.125	.290	.999[a]
β (target flash)	.034	.054	.053 / .077[b]	.055 / .056[b]	.046 / .074[b]	.062 / .083[b]	.048 / .037[b]	.062 / .056[b]
ρ (associative)	.296		—				.078 / 0.0024[c]	
ρ (orthographic)	—		.700				.854 / .027[c]	
$\Sigma\chi^2$ (error)	11.05		96.59				12.30	

Note. Dashes indicate inapplicable parameters.

[a] Simulations revealed that setting the estimate of alpha, α', to values approaching the actual α while holding the other parameters constant produced little change in the fit to the observed data. Nevertheless, the parameter estimation routine was able to find miniscule improvements in the fit by setting α' to its maximum value of 1.0.

[b] Experiments containing separate pools of target words for different conditions were allowed separate target encoding parameters, βs, for each pool of target words. In Experiment 2, βs refer to the orthographic and repetition priming conditions (on the left and right of the forward slash, correspondingly). In Experiment 3, the βs refer to the dissimilar and similar choice word conditions. In Experiment 4, the βs refer to the associative and orthographic priming conditions.

[c] In Experiment 4, the neither-primed and both-primed conditions necessarily introduced some degree of similarity between the choice words. Therefore, in Experiment 4, the similarity parameter, ρ, to the left of the forward slash, refers to prime similarity, and the second ρ after the forward slash is choice-word similarity.

It is clear that this rather simple model manages to capture the essentials of the data. There are a few things that can be said about the parameter estimates. The estimate of α (i.e., α') was lower than α for passive priming and higher than α for active priming, as was needed to reverse the direction of preference for repetition priming (the sharp-eyed reader will note that, for the active group, the associative predictions show a slight preference for targets, whereas the repetition predictions show a preference for foils; this is an example of a change in preference direction as ρ changes and is discussed following Experiment 2). The best-fit estimate of prime similarity, ρ, was .296. In other words, compared with repetition priming in which all the features were shared, the associatively related words were estimated to share about 30% of their features. We assume that these shared features are semantic in content, although associatively related words might share other, perhaps contextual, features because of repeated pairings of the associated pair. For instance the prime, SLING, and the associated choice word, SHOT, were used in Experiment 1. In isolation, the semantic similarity between these two words appears to be weak, but their repeated use in the compound SLINGSHOT may nevertheless result in substantial proportion of shared features.

Experiment 2: Orthographic–Phonemic Priming

The results of Experiment 1 and the success of the ROUSE model suggest that some features from the prime words are confused with features from the flashed target. The difference between repetition and associative priming in Experiment 1 is explained by ROUSE in terms of the number of features in the primes that are shared with the features in the choice words (i.e., the number of confusable features). Presumably, primes and associated choice words only share semantic (and possibly contextual) features, whereas repetitions share these features as well as orthographic and phonemic features. In Experiment 1, preference effects were much smaller for associative priming implying that orthographic–phonemic features were primarily responsible for the preference effects with repetition priming. Therefore, in Experiment 2, we tested this idea through conditions in which orthographic–phonemic similarity between primes and choice words was largely retained but semantic similarity was removed. If, as predicted by ROUSE, shared features generally determine preference effects, and, as implied by Experiment 1, the shared features in repetition priming are primarily orthographic and phonemic, then highly similar orthographic–phonemic primes should lead to a pattern of results similar to those obtained in Experiment 1 with repetition priming.

Method

There were 52 participants in the passive priming condition and 56 in the active priming condition. Four categories of word pairs (one to appear as a prime word and a related word to appear as a choice word) were created, each consisting of 48 five-letter words and 48 four-letter words. The pairs in three of the categories were orthographically very similar because four out of five or three out of four letters were identical and in the same position within the letter string. The remaining category was used for *repetition* priming. The orthographically similar categories were further subdivided by degree of phonemic and morphologic (i.e., semantic) similarity. The first category, labeled *orthographic,* was not semantically similar, and less phonemically similar and included pairs such as ANGEL → ANGER. The second category, labeled *orthographic and phonemic,* was more phonemically similar but not semantically similar and included pairs such as, ALTAR → ALTER. The third category, labeled *orthographic and morphologic,* consisted of variants of words with the same morphology (mostly created through tense changes with verbs) and included pairs such as AWAKE → AWOKE. As with AWAKE and AWOKE, these pairs were semantically similar, although tense changes can induce different meanings when the words are considered in isolation such as with "a wedding RING" and "your mother RANG." The 4 basic priming conditions were run for each of the four categories of words resulting in a total of 16 conditions. These conditions are illustrated with examples using particular words in Table 4. Each participant received 12 trials on each of these conditions.

If one assumes that orthographic priming produces large preference effects, it is unclear whether visual or abstract orthographic features are responsible. Therefore, letter case between the primes and choice words was manipulated. For the passive priming condition, the orthographic (e.g., ANGEL → ANGER) and orthographic and phonemic (e.g., ALTAR → ALTER) pairs were combined into one category. This larger category was then split randomly in half for each participant. For half the pairs, the primes were lowercased and for the other half, the primes were uppercased. Target flash and choice words were always lowercased, hence, there was a case switch between primes and target for the latter group of words. As seen in Table 4, switching case did not lessen preference effects, and, therefore, the orthographic and orthographic and phonemic categories were kept separate and all words were presented in lowercase in the active priming condition (the active priming group was run following completion of the passive priming group, allowing this slight change in procedure).

Table 4
Experiment 2: Examples and Results

Priming condition	Primes	Target	Choice words	Passive p(c)	Active p(c)
Type of priming Passive: *orthographic* and *orthographic and phonemic*; case switch Active: *orthographic* (examples shown)					
Neither	DATA + FLAG	AWRY	AWRY or HALT	.757	.801
Both	AIRY + HALO			.740	.716
Target	AIRY + FLAG			.845	.762
Foil	DATA + HALO			.628	.763
Type of priming Passive: *orthographic* and *orthographic and phonemic*; same case Active: *orthographic and phonemic* (examples shown)					
Neither	PIER + COLT	HALE	HALE or DUEL	.750	.798
Both	HAIL + DUAL			.765	.673
Target	HAIL + COLT			.824	.750
Foil	PIER + DUAL			.668	.743
Type of priming: *orthographic and morphologic*					
Neither	LIFT + DIVE	BENT	BENT or KNOW	.791	.807
Both	BEND + KNEW			.731	.717
Target	BEND + DIVE			.867	.766
Foil	LIFT + KNEW			.648	.766
Type of priming: *repetition*					
Neither	BELL + KNEE	GRIP	GRIP or JURY	.833	.780
Both	GRIP + JURY			.699	.655
Target	GRIP + KNEE			.860	.722
Foil	BELL + JURY			.782	.765

Note. p(c) = forced-choice performance.

To ascertain the generality of the passive–active difference, a different active priming task was used. This task was to determine if the two prime words could serve as the same part of speech. If the two prime words could serve the same part of speech, participants were instructed to press a key labeled *match* and otherwise press a key labeled *mismatch*. No feedback was given on this task, and accuracy was not calculated.

In Experiment 1, we used a separate pool of words for unrelated prime words. To control against any confounds introduced by using different primes, in Experiment 2 and in all subsequent experiments we created unrelated prime words through a re-pairing technique. Conditions that primed only one or neither choice word were created by randomly re-pairing prime and choice words such that, across participants, the same prime words were used in each of the priming conditions. In other words, primes and their related choice words could be presented in an intact or rearranged form. For example, one participant might receive the intact pairs ANGEL priming ANGER as well as CHOIR priming CHAIR, whereas another participant might receive the rearranged pairs ANGEL priming CHAIR and CHOIR priming ANGER.[6] All other procedures were the same as Experiment 1.

Passive Priming Results

The passive priming results are shown in Table 4. There were differences across the four types of primes, $F(3, 49) = 8.44$, $p < .001$, and these interacted with the four priming conditions, $F(9, 43) = 5.39$, $p < .001$. The repetition priming results from Experiment 1 are replicated for passive priming: a large both-primed deficit, $t(51) = 4.94$, $p < .001$; an overall preference effect, $F(1, 51) = 38.12$, $p < .001$; and a preference to choose repeated words, $t(51) = 2.82$, $p < .005$.

Three types of orthographic priming were used: two types with or without a case change and a third type using the orthographic and morphologic word pairs. There were no differences among the three types of orthographic priming, $F(2, 50) = 0.94$, $p = .40$, and no interaction between these types and the four priming conditions, $F(6, 46) = 2.17$, $p = .06$. When we separately analyzed the two types of orthographic priming between which the case of the primes varied, we found no effect of switching the case of the prime, $F(1, 51) = 0.75$, $p = .39$, and no interaction between switching the case of the prime and the four priming conditions, $F(3, 49) = 1.59$, $p = .21$ (see Table 4). This result strongly suggests that abstract orthography rather than visual similarity is crucial to prime interference and, more generally, that abstract orthography is the appropriate level of analysis in this paradigm. Because the orthographic priming conditions did not differ, they are collapsed and the results depicted in Figure 7.

When we analyzed the collapsed orthographic conditions, we found that performance in the both-primed condition was no different than in the neither-primed condition, $t(51) = 1.47$, $p = .15$. There was, however, an overall preferential effect, $F(1, 55) = 77.35$, $p < .001$, that consisted of a strong preference to choose words that were orthographically similar to the primes, $t(51) = 7.97$, $p < .001$. In ROUSE terms, passive orthographic priming seems to have produced a positive preference, but little if any preferential variability.

Active Priming Results

Participants took an average of 2,962 ms performing the part-of-speech matching task. The active priming results are given in Table 4. There were differences across the four types of priming, $F(3, 53) = 3.74$, $p < .025$, although these differences did not interact with the four priming conditions, $F(9, 47) = 0.86$, $p = .57$. The pattern of results from active priming in Experiment 1 was replicated for repetition priming: a large both-primed deficit, $t(55) = 4.49$, $p < .001$; an overall preference effect, $F(1, 55) = 4.77$, $p < .05$; and a slight preference to choose the word that was not a repetition, $t(55) = 1.55$, $p = .06$.

When we excluded repetition priming, we found that there were no differences across the three types of orthographic priming, $F(2, 54) = 2.01$, $p = .14$, and no interaction between these priming types and the four priming conditions, $F(6, 50) = 0.40$, $p = .88$. Therefore, we collapsed the three types of orthographic priming and they are displayed in the lower panel of Figure 7. For the collapsed orthographic priming conditions, there was a large both-primed deficit, $t(55) = 6.89$, $p < .001$; an overall preferential effect, $F(1, 55) = 21.18$, $p < .001$; but no preference for or against words orthographically similar to the prime words (i.e., the target-primed condition was not different from the foil-primed condition), $t(55) = 0.09$, $p = .92$. In ROUSE terms, the both-primed deficit is due to preferential variability.

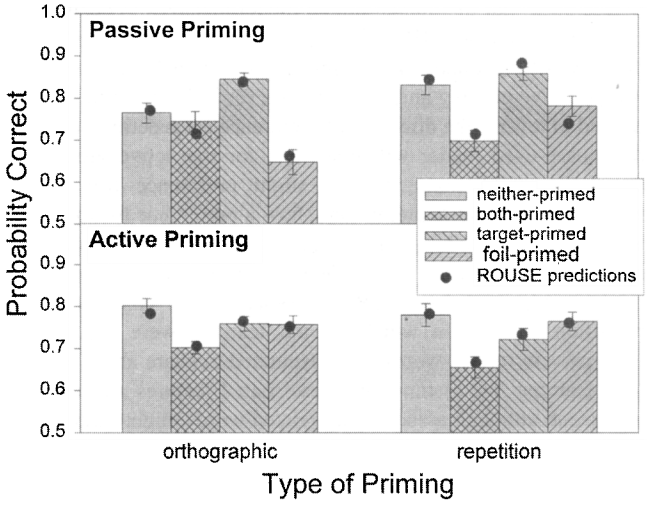

Figure 7. The accuracy results and predictions for Experiment 2, in which primes were identical to choice words (repetitions) or orthographically similar to choice words. Error bars are ± *SEM*. Passive versus active priming was a between-subjects manipulation. ROUSE = responding optimally with unknown sources of evidence.

Discussion

Experiment 2 results showed a general similarity between orthographic and repetition priming: both-primed deficits in both active and passive priming conditions, and a preference for a choice related to a passively processed prime and against a choice related to an actively processed prime. This general similarity

[6] In actuality, the re-pairing was more complicated than switching the primes between two sets of intact pairs. With a simple switching, participants could potentially use their memory of previous rearranged trials to discern the answer on later rearranged trials. Instead, a one-offset rearrangement was used (e.g., B′ → A, C′ → B, A′ → C).

suggests that both-primed deficits and preference effects are primarily due to orthographic prime interference. This conclusion is not entirely surprising given the visual nature of the identification task.

The failure to find significant differences for different degrees of morphologic and phonemic similarity suggests that phonemic and semantic features play a smaller role than orthographic features; these conditions produced similar results because of the overwhelming orthographic similarity. This conclusion appears to be somewhat in contradiction to Experiment 1 in which associative priming, although not nearly as strong as repetition priming, was nevertheless substantial. However, keep in mind that other, nonsemantic types of feature sharing might occur with associative priming. Furthermore, the morphologically related words used in Experiment 2 may have been much less semantically similar than repetitions, because they may have been much more likely to activate alternate meanings.

The invariance of the passive priming results when case was changed implies that comparisons are carried out at the level of abstract (not visual) orthography. The failure to obtain differences as a result of case change between primes and targets–foils is perhaps surprising conceptually (though similar results have been obtained in many other studies; see Evett & Humphreys, 1981). It should be noted that our procedure of placing primes in different screen locations from both targets and choice words might have reduced the importance of matching veridical visual features.

Several aspects of these results appear puzzling at first glance. In the passive priming results, the small both-primed deficit for orthographic priming compared with the large deficit with repetition priming suggests orthographic similarity plays less of a role in passive priming. At the same time in the passive priming results, the larger preference effect with orthographic priming suggests just the opposite. Part of the explanation may lie in the selection of words in the different conditions: In this experiment, the orthographic and repetition conditions used separate pools of target words. Differences in the overall perceptibility of each group of words could be inferred by comparing the neither-primed results. The participants in the active group apparently found these words equally perceptible, because neither-primed performance was the same in repetition and orthographic priming conditions, $t(55) = 1.08$, $p = .29$. However, the participants in the passive group found the perceptibility of the orthographic group of words lower than the repetition group of words, $t(51) = 3.83$, $p < .001$. Because different pools of words were used for the different types of priming, in the model fit we allowed a separate target-perception parameter, β, for each pool of words. We show in the next section that the different baseline levels of performance, corresponding to different β values, accounted for this pattern of results even though the same orthographic prime similarity, ρ, was required for both passive and active priming.

The ROUSE Model Applied to Experiment 2

As in Experiment 1, the use of different participants and a switch between passive versus active prime viewing was expected to induce different degrees of prime interference, α, and different degrees of prime discounting, α', so these parameters were separately fit to the passive and active groups. In Experiment 3 we used the same word pairs as Experiment 2 so the orthographic prime similarity probability, ρ, was required to be the same for both experiments. In the fitting of ROUSE throughout this study, we allowed the target perception parameter, β, to be estimated separately for different groups (and, hence, for different types of prime processing) and within group whenever a different pool of target words was used. (As can be seen in Table 3, the estimates of β across word pools within group were in some cases very similar, as in the active group in the present experiment. In these cases an assumption of equality would have sufficed. In other cases different β values across word pools within group were crucial for explaining the results, as in the passive group in the present experiment.) Table 3 gives the resulting parameter estimates, and the numerical predicted values can be found in Appendix B. The reasonably good fit of the model is illustrated by the dots in Figure 7.

For repetition priming, the both-primed deficit, the slight tendency to choose the unprimed word in the active priming group, and the tendency to choose the primed word in the passive priming group were similar to the results of Experiment 1. Thus, it is not surprising that the pattern of parameter estimates was generally similar in the two experiments. In particular, in both studies the estimate of α was higher than the actual value for active priming and lower than the actual value for passive priming. The estimated value of the prime similarity parameter, ρ, was reasonable: If orthography was the major (but not the sole) determinant of interference in repetition priming and the orthographically similar words shared three quarters or four fifths of their letters, one might expect a value of ρ to approximate 7.

For the passive priming group, a larger value of β for the repetition conditions allowed explication of a puzzle seen in the results. When the estimate of α was lower than the true value, as in the passive group, ROUSE predicts increasing separation between the target-primed and foil-primed conditions as prime similarity increases. Therefore, this factor should produce a greater disparity between these conditions in repetition priming (higher similarity) than in orthographic priming (lower similarity). The data showed the opposite. However, ROUSE predicts that the tendency to choose prime-related words will be larger for the orthographic condition because fewer target features are perceived in this condition (lower β); with fewer target-activated features, prime-activated features play a larger role, and preference is greater.

For the active priming group, the model provided a remarkably good fit. This seems surprising considering that there was considerable preferential variability for both orthographic and repetition priming (i.e., large both-primed deficits), but whereas repetition priming produced a slight preference against a repeated choice-word, orthographic priming was on average neutral with respect to the primed choice word. Naive expectations with ROUSE suggest that using the same α and α' should produce preference in the same direction regardless of prime similarity. This turned out to be false; for the γ, β, and α parameters used, there was an interaction between prime similarity and the direction of preference provided $\alpha' > \alpha$ (i.e., active priming). This interaction is explained below.

ROUSE and Prime Similarity

To take a closer look at the role of prime similarity, we produced predictions for the four critical priming conditions, for the case

when the estimate of α is lower than the true value (the upper panel of Figure 8, corresponding to passive priming), and for the case when the estimate of α is higher than the true value (the lower panel of Figure 8, corresponding to active priming) for values of prime similarity, ρ, ranging from zero to one. The parameter values are typical in that β was set to .05 and α was set to .1 (as in all simulations, γ was .02). To make the predictions clearly visible, two slightly exaggerated values of α' were used: .05 (corresponding to passive priming) and .3 (corresponding to active priming).

For passive priming, the predictions are shown in the upper panel of Figure 8; these predictions conform to expectations. As the value of ρ drops from 1, corresponding to the switch from repetition to orthographic priming in Experiment 2 or to associative priming in Experiment 1, one can see that there is no change in the neither-primed case (as must be the case because no features are shared), and the both-primed deficit decreases (as must be the case because fewer shared features produce smaller variation in the numbers activated by the primes). The strong preference for the primed alternative with high values of similarity gradually decreases to zero as the similarity decreases to zero.

The active priming predictions, shown in the lower panel of Figure 8 do not conform to naive expectations. Because α is the same in both panels, the neither-primed and both-primed conditions in the lower panel are the same as in the upper panel;

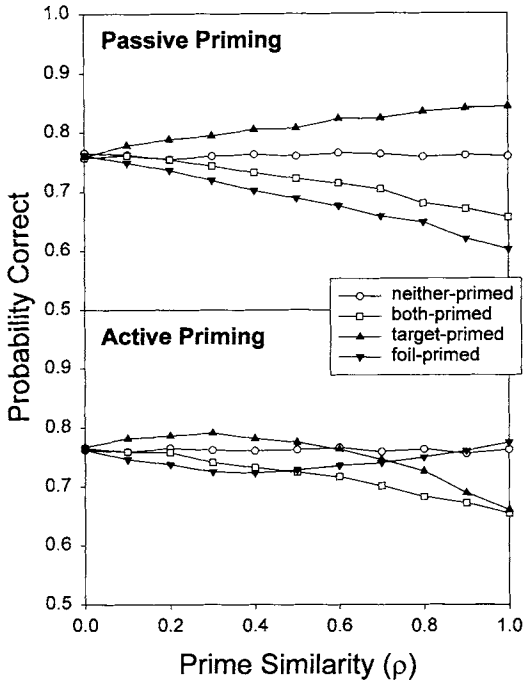

Figure 8. Predicted results as a function of prime similarity (ρ). The passive condition shows a preference for the related choice word and effects that simply increase in size with ρ. The active condition shows a crossover from a preference for a related choice word when similarity is low to a preference against a related choice word when similarity is high. The neither-primed and both-primed conditions are unaffected by discounting and identical in the two panels. The default parameters used in the simulations were as follows: $N = 20$, $\gamma = .02$, $\beta = .05$, $\alpha = .1$, α' (passive) = .05, and α' (active) = .3.

changing the estimate of α only affects the target-primed and foil-primed conditions. The preference predictions are the place where the failures of intuition appear: For high similarity (e.g., $\rho = 1$; starting at the right-hand side of the panel), there is a preference for the choice not related to the prime, as expected when $\alpha' > \alpha$. As similarity drops, however, the target-primed and foil-primed conditions change nonmonotonically, and the direction of preference changes from a preference against to a preference in favor of prime-related words.

To understand this crossover, we consider the two effects of priming: a gain in evidence that is due to prime-activated features versus a loss in evidence that is due to discounting target-activated and noise-activated features (refer to the right-hand four panels of Figure 5). As the similarity of the prime to a choice word decreases from 1.0, there is a reduction both in the gains due to prime-activated features (in the figure, decreasing the size of the $\rho\alpha$ gray circle), as well as a reduction in the loss that is due to discounting target-activated ($\rho\beta$) and noise-activated ($\rho\gamma$) features (in the figure, decreasing the size of the small embedded gray circles). With sufficient numbers of features available for priming and discounting, these two effects counterbalance each other, and the direction of preference remains constant as prime similarity is reduced. However, that the parameters are small (so that the size of the relevant circles are quite small) means that sometimes there are no features in the circles, and the probability of this event increases as ρ decreases. In particular, with α greater than β and γ, the absence of gains associated with prime-activated ($\rho\alpha$) features is less likely than the absence of losses associated with discounting target- ($\rho\beta$) and noise-activated ($\rho\gamma$) features. As prime similarity decreases and these absences become more likely, a disparity between the evidence gains and losses emerges, which causes the direction of preference to change.

That the preference crossover with active priming must occur for sufficiently low values of prime similarity is made clear by considering the situation in which only one feature of a primed alternative is expected to be shared with a prime. Only this shared feature is available for any kind of priming and discounting. Ignoring the possibility of more than one source activating the shared feature (which is an extremely low probability), the shared feature could be activated by a prime, target, or noise, or the feature might remain inactive. If the shared feature is activated by a prime (α), then it will result in a preference for the primed alternative, and if it is activated by the target or noise (β or γ), then it will result in a preference against the primed alternative. According to the equations found in Figure 4 this is true for all estimates of α because a discounted feature always provides more evidence than an OFF feature but less evidence than a nondiscounted ON feature. Because α is greater then β and γ, it follows that there will be a preference for prime-related words regardless of the level of discounting (i.e., for both passive and active priming).

As might be expected, the predicted crossover with active priming is contingent on the relative and absolute magnitudes of α, β, and γ. Simulations have shown that α must be greater than γ and β: Increasing β or γ or decreasing α can lessen or even eliminate the crossover (and even induce a crossover in passive priming from a positive to a negative preference as prime similarity decreases). In addition, a crossover at a reasonably high level of prime similarity requires that the absolute values of the parameters

must be scaled properly to the number of features, N: With appropriately small values of α, β, and γ compared with N, the nonlinearities created by the absence of evidence gains or losses are introduced even for sizable values of ρ.

Experiment 2 provides evidence in favor of this nonintuitive prediction of the ROUSE model. A similar effect was observed and predicted in Experiment 1: For the active priming results shown in Figure 2, it can be seen that although there was a preference reversal for repetition priming, there was essentially no preference or a slight preference in favor of related words with associative priming. ROUSE correctly predicted these results for the reasons explained here. Associative priming corresponds to a low level of prime similarity, and Figure 8 demonstrates that a preference for related words will result regardless of active versus passive priming. Further evidence concerning similar effects of changes in similarity between prime and choices is presented in Experiment 4. However, we first present a study in which we re-use the orthographically similar words of the present study to manipulate the similarity between the choice words. This is a critical study because ROUSE predicts similarly nonintuitive preference changes in active priming as a function of choice-word similarity.

Experiment 3: Repetition Priming and Choice-Word Similarity

Ratcliff and McKoon (1997) have used a paradigm similar to ours to explore long-term repetition priming. In their long-term priming paradigm, words were studied in a first phase; in the second phase of the experiment, target words were briefly flashed followed by a 2-AFC test. Ratcliff and McKoon varied the orthographic similarity of the two choice words and obtained results that tightly constrained possible models. We carried out such a manipulation in the present experiment. Before turning to our study, it is useful to review the findings of Ratcliff and McKoon.

Typical long-term priming results with other testing procedures (e.g., lexical decision or naming) show facilitation with repetition priming but not with associative or other sorts of priming (e.g., Ratcliff, Hockley, & McKoon, 1985; however for an alternate account see Becker, Moscovitch, Behrmann, & Joordens, 1997). For repetition priming, changes in the appearance of the choice word relative to the study word reduce the magnitude of the priming effect but do not eliminate it (e.g., Bowers & Michita, 1998).

In most long-term repetition priming studies, the facilitation observed could be due to a type of preference to choose a recently encountered word. This led Ratcliff and McKoon (1997) to use the 2-AFC design to determine if facilitation is due to a perceptual benefit or preference. They found no change in the both-primed condition, and the average of the target-primed and foil-primed conditions was about equal to performance in the neither-primed condition. Both results suggest that there was no perceptual facilitation for the target flash. Performance in the target-primed condition was higher than in the foil-primed condition, which demonstrated a preference for repetitions. That this advantage occurred when the choice words were orthographically similar, and was greatly reduced when the choice words had differing orthography, was a key factor in leading Ratcliff and McKoon to conclude that long-term repetition priming was a matter of "bias" rather than improved perception, and this assumption was a critical component of the model they developed to account for the results. Schooler et al. (2001) developed an alternative model to account for the same results (on the basis of Bayesian principles similar to those used in this article), but their model also assumed that perception was not improved by priming.

In the following experiments, it is important to note that both models make critical use of the distinction between diagnostic features (those that differ between the choices) and nondiagnostic features (those that are the same for the two choices). In fact Schooler et al. (2001) showed that the reduction in variance caused by the change in number of diagnostic features can itself account for the differences between similar and dissimilar choices.

Because the pattern of results observed by Ratcliff and McKoon (1997) differed markedly for similar and dissimilar choice words, and because these differences were critical to constraining possible models of the results, we decided to test the ROUSE model by varying the similarity of the choice words in our present short-term priming paradigm. In Experiment 3, we used only repetition priming. We used the word pairs of Experiment 2 but included conditions in which the orthographically similar word pairs of that study appeared as choice words. Instead of using pairs of dissimilar words selected to be as dissimilar as possible, as was done in Ratcliff and McKoon's study, we created our dissimilar word pairs by randomly pairing two words.

Method

There were 55 participants in the passive priming condition and 56 in the active priming condition. Repetition priming was used for all conditions. The four categories of word pairs found in Experiment 2 were re-used; in the 3 sets of orthographically similar choice word conditions, the members of a given pair were presented as choice words (one being the target, the other the foil). The dissimilar choice-word conditions were identical to the repetition conditions found in Experiment 2. All other procedures were the same as in Experiment 2. The active priming group used the same part-of-speech matching task. In Experiment 2, the letter-case manipulation was only administered to the passive priming group. To test the generality of case indifference, in Experiment 3 we used the letter-case manipulation for the active priming group; as in Experiment 2, this condition combined the orthographic and orthographic and phonemic word pairs and randomly assigned word pairs to conditions in which primes appeared in lowercase or conditions in which primes appeared in uppercase (resulting in same case and case switch between prime and target). For the passive priming group the word pairs were kept in their original categories and all words appeared in lowercase. The conditions are illustrated with examples using particular words in Table 5.

Passive Priming Results

The results are given in Table 5. There were three types of orthographic choice-word similarity; these did not differ statistically, $F(2, 53) = 3.11$, $p = .05$, and interacted only weakly with priming condition, $F(6, 49) = 2.84$, $p < .025$. Inclusion of the dissimilar choice-word similarity conditions led to differences, $F(3, 52) = 16.59$, $p < .001$, and these differences interacted with the four priming conditions, $F(9, 46) = 3.05$, $p < .01$. The main effect of choice-word similarity was due to better performance for the dissimilar than orthographically similar choice-word conditions (collapsing across priming conditions, $t[54] = 6.63$, $p < .001$). This is natural because the decision is more difficult when

Table 5
Experiment 3: Examples and Results

Priming condition	Primers	Target	Choice words	Passive p(c)	Active p(c)
Choice-word similarity Passive: *orthographic* (examples shown) Active: *orthographic* and *orthographic and phonemic*; case switch					
Neither	DATA + FLAG	AWRY	AWRY or AIRY	.724	.740
Both	AWRY + AIRY			.609	.613
Target	AWRY + FLAG			.709	.735
Foil	DATA + AIRY			.558	.625
Choice-word similarity Passive: *orthographic and phonemic* (examples shown) Active: *orthographic* and *orthographic and phonemic*; same case					
Neither	PIER + COLT	HALE	HALE or HAIL	.694	.740
Both	HALE + HAIL			.564	.607
Target	HALE + COLT			.692	.701
Foil	PIER + HAIL			.529	.537
Choice-word similarity: *orthographic* and *morphologic*					
Neither	LIFT + DIVE	BENT	BENT or BEND	.665	.723
Both	BENT + BEND			.633	.606
Target	BENT + DIVE			.750	.726
Foil	LIFT + BEND			.539	.542
Choice-word similarity: Dissimilar (same as *repetition* in Experiment 2)					
Neither	BELL + KNEE	GRIP	GRIP or JURY	.732	.793
Both	GRIP + JURY			.705	.719
Target	GRIP + KNEE			.797	.766
Foil	BELL + JURY			.624	.768

Note. p(c) = forced-choice performance.

the choice words are more similar. Because the differences among types of orthographic similarity were small, these conditions were combined, and the results are shown in Figure 9.

There was a highly significant both-primed deficit in the orthographically similar choice-word conditions, $t(54) = 6.04$, $p < .001$. However, in the dissimilar choice-word conditions, the small both-primed deficit did not reach significance, $t(54) = 0.90$, $p = .37$. There were preferential effects for both the orthographically similar, $F(1, 54) = 111.57$, $p < .001$, and dissimilar, $F(1, 54) = 40.44$, $p < .001$, conditions: Participants tended to choose the repeated word for both types of choice-word similarity: dissimilar, $t(54) = 6.75$ $p < .001$; orthographically similar, $t(54) = 8.24$, $p < .001$. These results largely replicate those from the passive groups in Experiments 1 and 2.

Active Priming Results

Participants in the active priming group took an average of 2,811 ms to perform the part-of-speech matching task.

Three types of orthographic priming were used: two types with or without a case change and a third type using the orthographic and morphologic word pairs. The results are shown in Table 5. There were differences across the three types of orthographic choice word similarity, $F(2, 54) = 5.18$, $p < .01$, but these differences did not interact with priming condition, $F(6, 50) = 2.11$, $p = .07$. In all cases, the qualitative trends across priming conditions were the same. Presenting primes in a different case than the target and choice words produced differences that interacted weakly with the four priming conditions, $F(3, 53) = 2.93$, $p < .05$, but as seen in Table 5, this was due to quantitative, not qualitative differences. These results again support the conclusion that feature comparisons occur largely at the level of abstract orthography. In light of the foregoing analyses, the results are collapsed across the three types of orthographic similarity and the case change manipulation, and graphed in Figure 9.

As with the passive group, there were differences with choice-word similarity when the dissimilar conditions were included, $F(3, 53) = 37.21$, $p < .001$, and these differences interacted with the four priming conditions, $F(9, 47) = 4.86$, $p < .001$. As with passive priming, the main effect of choice-word similarity was due to lower performance with orthographically similar choice words, $t(55) = 10.42$, $p < .001$.

For both dissimilar, $t(55) = 2.92$, $p < .005$, and orthographically similar choice words, $t(55) = 7.41$, $p < .001$, there were deficits in the both-primed conditions. Likewise, there were preferential effects for both dissimilar, $F(1, 55) = 4.08$, $p < .05$, and similar choice words, $F(1, 55) = 79.54$, $p < .001$. When comparing the target-primed and foil-primed conditions, participants preferred the repeated word when the choice words were orthographically similar $t(55) = 6.12$, $p < .001$ (it is important to note that this result is opposite to that found in the active group in Experiments 1 and 2), whereas there was no difference between these conditions when the choice words were dissimilar, $t(55) = 0.05$, $p = .96$ (in keeping with the preference removal or reversals seen in Experiments 1 and 2 with active priming).

Discussion: ROUSE and Choice-Word Similarity

For active priming, the observation of a preference for repeated words only for similar choice words and not for dissimilar words

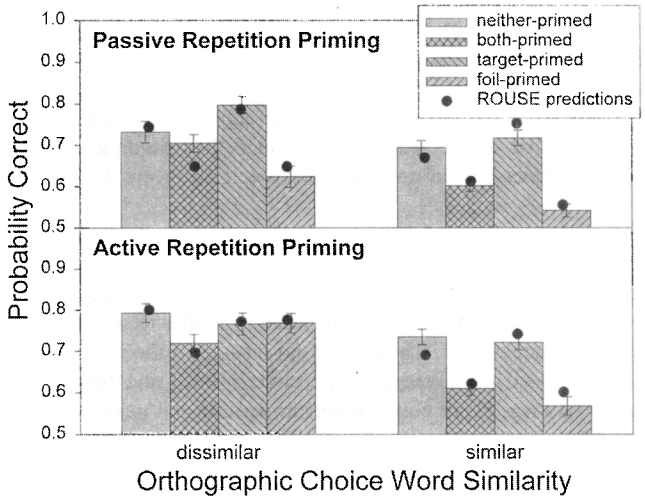

Figure 9. The accuracy results and predictions for Experiment 3, which used repetition priming and two choice words that were either orthographically dissimilar or similar. Error bars are two ± *SEM*. Passive versus active priming was a between-subjects manipulation. ROUSE = responding optimally with unknown sources of evidence.

is analogous to the Ratcliff and McKoon (1997) results but does not seem at first glance to be in accord with the predictions of ROUSE (or the results from the first 2 experiments). As we demonstrate next, however, the model provided a surprisingly good account of the findings.

How can ROUSE predict a preference for repeated words when the choices are similar, even though α' is larger than α (i.e., active priming)? The answer depends on the fact that similar choices contain many shared features. These features are nondiagnostic and play no role in decision making. Thus, the effect of making the choices similar is to reduce considerably the total number of features involved in the decision. In the discussion of similarity between prime and choice words we highlighted the nonlinear effects that come into play when the number of features becomes very low such that the evidence gains or losses associated with priming and discounting are likely to be absent. As was the case in that discussion, the effect is to shift preference in the direction of the related choice word.

This is demonstrated with simulations in Figure 10 using the same default parameters used in the construction of Figures 6 and 8 (prime similarity is set to 1.0 in order to match the repetition priming used in Experiment 3). The leftmost set of points in Figure 10 is the same as the rightmost set of points in Figure 8 (in both figures, these points are repetition priming with dissimilar choice words); as expected there is a preference for repeated words with passive priming and a preference against repeated word with active priming. As choice-word similarity increases, performance drops for both types of priming because of the decreasing number of diagnostic features (at choice-word similarity of 1.0, the choice words are identical and performance is at chance). For passive priming, the preference direction is unaffected by increasing choice-word similarity. For active priming, the preference against repeated words changes to a preference for repeated words as similarity increases and the number of diagnostic features drops.

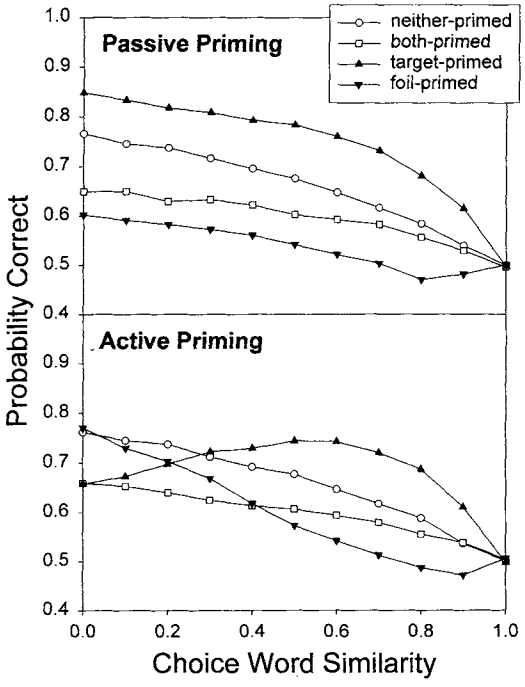

Figure 10. Predicted results as a function of choice-word similarity for repetition priming (ρ = 1). As choice-word similarity approaches 1.0, fewer diagnostic features remain and performance decreases. The passive priming effects are relatively unaffected by changes in choice-word similarity. The active priming effects are more complex: Similar to the effect of prime similarity as seen in Figure 8, increases in choice-word similarity cause a preference crossover: Low levels of similarity produce a preference for the choice word unrelated to the prime, and higher levels produce a preference for the choice word related to the prime. The default parameters used in the simulations were as follows: N = 20, γ = .02, β = .05, α = .1, α' (passive) = .05, and α' (active) = .3.

The predicted crossover follows from reasoning similar to that found in the *ROUSE and Prime Similarity* section. Prime similarity reduced the number of features available for priming and discounting such than on many trials, the evidence gains or losses associated with priming were missing. A similar effect occurs with increasing choice-word similarity. The number of diagnostic features decreases as choice-word similarity increases and, therefore, the number of features available for priming and discounting decreases. As choice-word similarity increases, the relative probabilities of the absence of evidence gains and losses become more important than the actual average balance of gains and losses. More specifically, the probability of evidence gain is related to α and the probability of evidence loss is related to β and γ, and therefore, the relative strengths of these parameters determines the direction of preference regardless of the estimate of α. As with the prime-similarity crossover, the same relative manipulations of the parameters serve to eliminate the crossover, and the same absolute manipulations of the parameters serve to influence the position of crossover as a function of choice-word similarity.

The ROUSE Model Applied to Experiment 3

The predictions of the ROUSE model, on the basis of the best-fitting parameters found in Table 3, are given in Figure 9. The model's predictions mimic the qualitative patterns of results and come close to many of the quantitative observations. The crucial prediction of the model occurs for the active priming group: The model predicts both (a) for dissimilar choice words, no differential preference for the repeated choice word and (b) for orthographically similar choice words, a preference for the repeated choice word. It is particularly striking that this second prediction differs from those for the active groups in Experiments 1 and 2, even though in all three cases α' was estimated to be higher than α. That is, the estimate of prime activation was higher than the actual value in all three studies, yet the direction of preference was predicted to reverse in this study, because of the similarity of the choice words. Even more compelling is that the reversal was observed within subject, and the model predicted the results using the same α and α' values for both directions of preference. That this unanticipated prediction matched the results lends credence to the model.

Because the same words were used in Experiments 2 and 3, the same orthographic similarity parameter, ρ, was used in the fit of both experiments. It is important to note that the orthographic similarity parameter, ρ, performed a different function in Experiment 3 than it did in Experiments 1 and 2. In Experiments 1 and 2, ρ determined the similarity of prime to choice word and was set at a level sufficient to produce the observed relation between repetition and associative priming or between repetition and ortho-

graphic priming. In Experiment 3, ρ determined similarity between the two choice words, and its value was most important for determining the proportion of features in the decision process (which determines performance levels and the switch in the active priming preference results). It is noteworthy that good fits are obtainable with a common similarity parameter in both situations.

As seen in Figure 9, in a few instances, the fit to data was quantitatively awry. An examination of Appendix B reveals that only in two instances across all the experiments was the fit more than two standard errors away from observed data. Both of these mispredictions occurred in the difference between the neither-primed and both-primed conditions in Experiment 3: The observed data show a larger both-primed deficit for orthographically similar than dissimilar choice words, whereas the model predicts the opposite (as seen in Figure 10). A subtle, but necessary, confound in the experiment might explain these mispredictions. The both-primed condition with orthographically similar choice words was qualitatively different than the other conditions because only in this condition were the primes orthographically similar to each other. The participants may have given more attention to the letters that differed between the two prime words, perhaps causing an increase in source confusion (α) for these features. This increase would occur only for the similar choice-word condition; because the both-primed deficit is directly related to the magnitude of α, this attentional account might explain the two quantitative discrepancies between predictions and data in this study. In our fit of the data, we did not allow for this possibility and the same α was used across all conditions.

Experiment 4: Associative and Orthographic Priming

Considering that Masson and Borowsky (1998) found facilitation in short-term semantic priming even when preference was controlled, our failure to find such an advantage in Experiment 1 was surprising. However, our three studies used two prime words, whereas previous priming studies used one. In Experiment 4, we used only one prime word to test whether the difference in results was due to a dilution of the force of associative priming when the number of prime words increased. Furthermore, in Experiment 4 we examined how associative priming varies with the association directionality between prime and target. Again, active and passive prime processing were used. In addition, in this experiment we compared associative priming to orthographic priming; in one set of conditions a single prime word was used that was orthographically similar to both choice words. The results of Experiment 2 suggested that orthography is a major determinant of preference effects, so we anticipated that primes sharing all but one letter would produce greater preference than associatively related primes.

Association directionality is broken into three categories. In a *symmetrical* association, the prime produces the target or foil as an associate, and the target or foil produces the prime as an associate. In an asymmetrical *forward* association, the prime produces the target or foil as an associate, but the target or foil does not produce the prime as an associate. In an asymmetrical *backward* association, the prime produces neither choice word as an associate, but the target or foil produces the prime as an associate. Most experiments select words on the basis of the forward association strengths without reference to the backward association strengths (however, see Thompson-Schill, Kurtz, & Gabrieli, 1998). Therefore, traditional priming studies have used some mixture of the symmetrical and forward categories. Different theories of representation will predict different results for these priming categories. For example, in the present version of ROUSE, all that matters is shared semantic features. Association directionality plays a role only in as much as it reflects differential degrees of semantic similarity. Thus, this version of ROUSE predicts that these three categories of priming should produce qualitatively similar results.

Method

There were 62 participants in the active priming group and 87 in the passive priming group. The procedures in Experiment 4 were similar to those of Experiment 1, but only one word appeared as the prime instead of two. The one prime word was repeated in two locations on the screen corresponding to the two locations where separate prime words were displayed in other experiments. Unlike Experiments 1–3, chin rests were not used.

Four categories of 40 triples were created: three categories with a prime and two associates and a fourth category with a prime and two orthographically similar words. For the associative categories, all associates were four or five letters in length and primes were three to five letters in length. In the forward category, primes associated to two associates with an average strength of .168 (Nelson et al., 1994). Searching through the norms with these associates as potential cue words revealed that these associates did not associate back to the primes with any known strength. This does not necessarily imply that the primes and associates were semantically dissimilar. For example, the prime ASHES associates to DUST, but DUST does not associate back to ASHES. In selecting these asymmetrical associates, we required that all associates exist in the norms as primes. In the backward category, two choice words associated to the single prime, but the prime did not associate to either choice word. The average association strength was .260. In the symmetrical category, association strengths were found in both directions with an average forward strength of .217 and a backward correlation of at least .07. The orthographic category was created from five-letter words. The prime was required to share four out of five letters (in the same positions) with each of the choice words (but not necessarily the same four letters). An additional pool of 336 five-letter words was created for use in the practice sessions and threshold duration block.

A necessity created by the use of a single-prime design was differential choice-word similarity for the neither-primed and both-primed conditions compared with the target-primed and foil-primed conditions. For example, the prime CURVE would be followed by a choice between CARVE and CURSE in the orthographic, both-primed condition. In this example, it is impossible to test the target-primed condition with this same high level of choice-word similarity; the foil would necessarily be similar to the prime. To create the target-primed and foil-primed conditions, dissimilar choice words were randomly selected. Despite this difference, the neither-primed and both-primed conditions could be compared with each other, and likewise the target-primed and foil-primed conditions could be compared with each other. To a lesser degree, the analogous situation exists for associative priming (i.e., the choice words are necessarily somewhat semantically similar to each other only in the both-primed and neither-primed conditions).

With only one prime, a new task was necessary for active priming. Participants performed a simple affect task by rating each prime word as positive or negative. In an experiment to be published elsewhere (Experiment 4 in Huber, 2000), it was found that a single prime yielded a small but reliable both-primed benefit with associative passive priming. That experiment examined the relationship between the both-primed benefit and prime duration, and the largest both-primed benefit was observed for prime durations of 50–500 ms. This finding led us to use a prime duration of 250 ms in the passive priming condition of the present experiment (as in

Experiments 1–3, active priming consisted of presenting primes until response followed by a re-presentation of the primes for 500 ms). Each of the four categories of words was tested in each of the 4 basic priming conditions. These 16 conditions each appeared 10 times distributed throughout the experiment. Table 6 contains specific examples of these 16 conditions. All other procedures were as explained in Experiment 1.

Results

For the active priming group, participants took an average of 1,179 ms to perform the affect rating task.

There were no significant differences across the three types of associative priming: For the passive group, $F(2, 85) = 1.01$, $p = .37$, and for the active group, $F(2, 60) = 0.55$, $p = .58$. There were also no interactions between these priming types and the four priming conditions: For the passive group, $F(6, 81) = 1.21$, $p = .31$, and for the active group, $F(6, 56) = 0.71$, $p = .65$. Therefore, we collapsed the three types of associative priming.

Figure 11 shows the priming results separately for orthographic priming and the average of the three types of associative priming, and Table 6 shows the separate results. Including orthographic priming, there were differences across the four types of priming: For the passive group, $F(3, 84) = 27.78$, $p < .001$, and for the active group, $F(3, 59) = 11.86$, $p < .001$. These differences interacted with the four priming conditions: For the passive group, $F(9, 78) = 6.40$, $p < .001$, and for the active group, $F(9, 53) = 2.99$, $p < .01$. Because there were no significant differences among the three types of associative priming, these differences with the inclusion of orthographic priming largely reflect differences between associative and orthographic priming.

The results for the orthographic priming conditions appear similar to the repetition priming results of Experiments 1–3, for both the passive and active groups. A both-primed deficit was observed for the passive group, $t(86) = 6.74$, $p < .001$, as well as for the active group, $t(61) = 2.27$, $p < .025$. The usual check for the existence of preferential effects cannot be performed because the neither-primed–foil-primed and target-primed–both-primed comparisons were confounded with choice-word similarity. However, the tendency to respond with a word related, or unrelated, to the prime is also indicative of preferential effects, and these were clear: For the passive group, there was a tendency to choose the word orthographically related to the prime, $t(86) = 3.63$, $p < .001$, whereas for the active group there was a tendency to choose the word not orthographically related to the prime, $t(61) = 2.56$, $p < .01$.

For associative priming in the passive group, there was a both-primed benefit (although small in magnitude), $t(86) = 2.78$, $p < .005$; for the active group, there was no both-primed deficit or benefit, $t(61) = 0.48$, $p = .64$. There was a tendency to choose the word associatively related to the prime for both passive and active prime viewing: for the passive group, $t(86) = 7.97$, $p < .001$, and for the active group, $t(61) = 2.86$, $p < .005$.

Discussion

The data show that the tendency to choose prime-related choice words occurs in both passive and active conditions, for each

Table 6
Experiment 4: Examples and Results

Priming condition	Prime	Target	Choice words	Passive p(c)	Active p(c)
		Type of priming: *Forward* associative			
Neither	WOUND	CABIN	CABIN or HOTEL	.749	.797
Both	LODGE			.756	.802
Target	LODGE		CABIN or BEACH	.797	.829
Foil	SHELL			.705	.781
		Type of priming: *Backward* associative			
Neither	BOARD	PSALM	PSALM or VERSE	.735	.811
Both	BIBLE			.755	.810
Target	BIBLE		PSALM or BRAWL	.794	.798
Foil	FIGHT			.692	.771
		Type of priming: *Symmetric* associative			
Neither	WHOLE	HAPPY	HAPPY or FROWN	.730	.824
Both	SMILE			.787	.802
Target	SMILE		HAPPY or ABOVE	.831	.836
Foil	BELOW			.691	.771
		Type of priming: *Orthographic*			
Neither	HEDGE	DUSTY	DUSTY or MISTY	.717	.757
Both	MUSTY			.591	.700
Target	MUSTY		DUSTY or SHAPE	.739	.710
Foil	SHAKE			.618	.779

Note. p(c) = forced-choice performance.

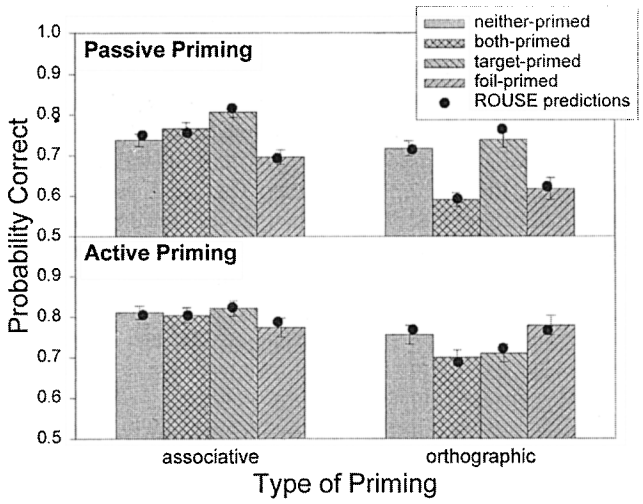

Figure 11. The accuracy results and predictions for Experiment 4, which used one prime word that was either associatively or orthographically related to the choice words (the choice words were incidentally somewhat similar to each other for the neither- and both-primed conditions, and dissimilar for the target- and foil-primed conditions). Error bars represent 2 *SEM*s. Passive versus active priming was a between-subjects manipulation. ROUSE = responding optimally with unknown sources of evidence.

direction of association. This result is in agreement with a recent study that found associative priming in lexical decision regardless of association direction (Thompson-Schill et al., 1998). In our experiment, the perceptual benefit occurred across different visual locations, so the finding of such a benefit was not dependent on the prime and target occupying the same screen location. One needs to be cautious in drawing any strong conclusions, however, because in both passive and active conditions, the effect was strongest for the symmetric association.

This study replicates Masson and Borowsky's (1998) observation of semantic facilitation when preference is controlled, even though our prime(s) were displayed in a different location than the target. In particular, associative–semantic priming with a single, passively viewed prime presented in two locations for 250 ms produced a small both-primed benefit. This benefit with passive priming was observed regardless of the direction of association but was largest for the symmetric associative condition (see Table 6: forward = .007; backward = .020; symmetric = .057). The fact that these effects are quite small makes it difficult to reach any firm conclusions on the basis of the failure of the differences among these numbers to reach statistical significance. If priming were indeed uniform across directions of association, the result would be consistent with a theory, such as ROUSE, in which feature overlap, not ease of associative production, is the crucial determinant of preferential and perceptual aspects of priming; but the present data do not allow any strong conclusion to be drawn.

The ROUSE theory in its present form cannot predict both-primed benefits. The theory would have to be augmented by additional mechanisms to deal with such effects. We did not extend ROUSE to do so in this study, but present possible mechanisms for such perceptual benefits in the General Discussion section. It is important to note that the both-primed benefit only occurred with passive associative priming and was small compared with the preference effects discussed next.

Similar to the repetition priming results and most of the orthographic priming results in Experiments 1–3, orthographic priming in this study produced both-primed deficits for both the passive and active groups. Also similar to the results of earlier studies, orthographic priming produced a switch in the direction of preference between passive and active priming. In contrast, associative priming in the present study produced no such switch; for both the passive and active groups, there was a tendency to choose related words.

Figure 8 (and the related discussion following Experiment 2, titled *ROUSE and Prime Similarity*) explains ROUSE's ability to predict this complex pattern of interactions. It also helps to keep in mind that target-primed versus foil-primed conditions in this study involved dissimilar choices, whereas neither-primed versus both-primed conditions involved moderately similar choices. Thus, the target-primed and foil-primed predictions (dissimilar choice words) can be gathered through reference to Figure 8: High values of ρ correspond to orthographic priming and produce predictions as shown in the right-hand area of the upper and lower panels of Figure 8 for passive and active priming. Low values of ρ correspond to associative priming and produce predictions for passive and active priming as shown in the left-hand area of the upper and lower panels of Figure 8. Predictions for the neither-primed and both-primed cases (moderate choice-word similarity) would have to be inferred through reference both to Figures 8 and 10.

The ROUSE Model Applied to Experiment 4

The parameter estimates for Experiment 4 are given in Table 3. Because the same words were used for both the active and passive conditions, the same prime-similarity values, ρs, were required to handle both sets of results. One prime-similarity value was used for the associative conditions and a second value was used for the orthographic conditions. In addition, this experiment necessarily introduced a degree of choice-word similarity for the neither-primed and both-primed conditions because of the requirement that both choice words be similar to a single prime word (see the *Method* section). Because it was not clear what assumption was appropriate for deriving this similarity (e.g., assuming independent feature sampling would result in a choice word similarity of ρ^2), an additional similarity parameter was allowed for orthographic and associative priming to capture this incidental choice word similarity (see Table 3). As with Experiments 1–3, separate prime activation, α, and estimated prime activation, α', values were used for the separate active and passive groups of participants and separate target-activation parameters, βs, were used for the associative and orthographic conditions because different target words were used in each case.

The predictions given in Figure 11 show that ROUSE captures the results found in Experiment 4 with the exception of the small both-primed benefit with passive associative priming (although the fitting routine managed to find a compromise such that predicted performance lies within the error bars for both the neither-primed and both-primed conditions). Across the experiments reported in this article, this is the only observation of a both-primed benefit. Nevertheless, we have performed an experiment, to be reported elsewhere (Huber, 2000), which replicated this both-primed ben-

efit with passive associative priming and a single prime word. In that study, we observed the both-primed benefit for prime durations ranging from 29 ms to 457 ms. Any both-primed benefit is outside the scope of ROUSE model in its present form, and mechanisms capable of producing such perceptual benefits are discussed in the *Short-Term Priming Theories and Perceptual Effects* section in the General Discussion.

The most remarkable aspect of the observed and predicted results was found in the active priming group. As predicted by Figure 8 and explained in the *Rouse and Prime Similarity* section, the model predicts that for overestimates of α (i.e., active priming), preference removal breaks down for low prime similarity. Orthographically similar words sharing four out of five letters in position were expected to have a high degree of similarity, given the results of Experiment 2. The results of Experiment 1 suggested associative similarity should be much less. The parameters seen in Table 3 confirm these expectations, and the low ρ value with associative priming resulted in a preference for associative choice words, whereas the high ρ value with orthographic priming resulted in a preference against orthographically similar choice words. That the associative similarity parameter (.078) was much less than the associative similarity parameter in Experiment 1 (.296) was expected given the weaker association strengths found in this experiment (.378 association strength in Experiment 1 vs. .215 in Experiment 4).

Finally, we note that the estimate of orthographic choice-word similarity (.027) for the neither- and both-primed cases were much smaller than one would expect on the basis of letter overlap. If a more conceptually sensible estimate (say in the 0.50–0.70 range) were used instead, this would cause a selective performance reduction in the neither-primed and both-primed conditions (because of the reduction in number of diagnostic features) that would reduce the quantitative goodness of fit (although the qualitative pattern of predictions would remain fine). In simulations, we have explored the possibility that parameters vary across participants, which reduces the force of effects and allows us to attain a good fit with a sensibly high estimate for choice-word similarity. However, such an approach seemed to us to go too far in adding complexity simply for the sake of attaining a good fit, and we leave this issue as a target for future research.

General Discussion

In this article, we have presented a model, ROUSE, to predict data collected from a task involving short-term priming of words presented at threshold, followed by a forced-choice decision. The general Bayesian approach we used to generate the model can of course be used in a variety of other settings, and some of these have already been developed. In long-term priming, Schooler et al. (2001) used a very similar model for perceptual identification followed by forced choice, but also for perceptual identification followed by yes–no decisions, and by naming. Schooler et al.'s model did not use discounting of repeated features–words, although the present results suggest it might prove useful to do so (as discussed later). In lexical decision, Wagenmakers et al. (2000) used a related model to predict response time and accuracy data for different kinds of words and nonwords, and for repetitions of these. In a different setting, Shiffrin and Steyvers (1997, 1998) used a Bayesian model with many similarities to ROUSE to predict recognition memory data. In none of these examples were participants assumed to be less than optimal (i.e., using inaccurate estimates of parameters). The use of over- and underestimates of prime interference is a unique contribution of the ROUSE model.

For many researchers and theorists, the present findings may suggest a new way of interpreting short-term priming. With a few notable exceptions (Hochhaus & Johnston, 1996; Johnston & Hale, 1984; Masson & Borowsky, 1998), previous short-term priming studies have been ambiguous as to whether the effects of priming are due to perception or preference. The best method we have found to distinguish these effects involves the use of prime-related foil items within a forced-choice task. Using such methods, we obtained strong conclusions and believe the theoretical implications extend to other short-term priming paradigms, such as lexical decision, naming, and the naming version of perceptual identification, and possibly to other tasks as well.

We observed preference effects to be ubiquitous, with large preference effects found in all experiments and with all types of primes. One way in which preference effects were manifested was as a both-primed deficit. This occurred in almost all cases (except associative priming, discussed below). In ROUSE, this both-primed deficit was interpreted as variability caused by prime activation favoring one or another choice word by chance.[7] Only in the case of a single associative prime viewed passively did we observe the both-primed benefit that is expected with perceptual enhancement. In this one case, the both-primed benefit was small, suggesting that perceptual effects, when they exist, are small compared with preference effects (perceptual enhancements might exist with other types of primes but be masked by the much larger preference effects). Therefore, it seems likely that the advantage for related targets found in many prior studies is largely due to what we have termed preference effects. These effects are not to be confused with bias in a signal-detection sense and are not to be cast aside as uninteresting; they are important and highly constraining for theory, and form the basis for the ROUSE model.

It is particularly important that we were able to reverse the direction of preference through the use of passive versus active processing of primes. Interpretation of prior studies is even more difficult because in many cases it is unclear whether the procedures induced participants to process primes actively or passively (in ROUSE terms, over- or under estimate α). Under such circumstances, it is not surprising that overviews of this literature (e.g., Neely, 1991) reveal a complex and often contradictory array of results.

It seems fairly clear how a ROUSE-based approach might be used to deal with traditional short-term priming studies, and we simply mention what may be the most straightforward approach. First, consider the naming version of perceptual identification. In this paradigm, participants must name the flashed target word.

[7] In theory, the existence of the both-primed deficit would require that the prime information remain so that it could be confused with the choice words. Huber (2000) reported studies in which compound words appeared as primes and choice words were identical to both or neither of the constituent words. The first studies, carried out with Kirstin Ruys, showed both-primed deficits. Later studies placed either masks or other words between the primes and the target flash, and both manipulations eliminated the both-primed deficit.

Schooler et al. (2001) calculated likelihood ratios for all words (in a similarity neighborhood) in the lexicon; the largest of these was named, if the likelihood ratio was above a threshold. If the primes provide features that match a particular word in the lexicon (in addition to noise and features activated by the presentation), then in the absence of discounting, that word has an advantage in accuracy or speed of responding. If the situation fosters sufficient discounting of the prime information (and other criteria are met), then the normal priming benefit can be removed or even reversed (Humphreys et al., 1988).

Next, consider the task of lexical decision. Wagenmakers et al. (2000) assumed that the test string is compared with words within an orthographic similarity neighborhood of the test string. As in the Bayesian models for memory (e.g., Shiffrin & Steyvers, 1997), the likelihood ratios for each word in the neighborhood are summed and divided by the total; a "word" decision is made if the result is greater than 1.0 (i.e., if the odds in favor of word are above 1.0). If primes contribute features that match a given lexical entry, and there is no discounting, this tends to increase the likelihood (or speed) of a word decision. As before, discounting could remove or even reverse this tendency (Dominguez & de Vega, 1997; Lukatela & Turvey, 1996; Lupker & Colombo, 1994; O'Seaghdha & Marin, 1997).

Effects of Passive Versus Active Prime Processing

At present, we do not have the data to explicate fully the nature of the passive–active manipulation. For example, all our experiments confounded the passive–active manipulation with prime duration. A recent study in our laboratory to be reported elsewhere, varied prime duration under passive instructions; the results showed that longer prime durations produced effects similar to those of active priming but somewhat smaller in magnitude. Interpretation remains ambiguous because longer prime exposure might lead participants to engage in more elaborate (i.e., more active) processing of primes. In a follow-up experiment, we had participants make judgments (i.e., active priming) about briefly presented (< 200 ms) primes, and the active priming pattern of results was nevertheless observed. This demonstrates that the passive–active distinction is more than just low-level visual presentation duration. This result also demonstrates that large preference effects emerge even with substantially shorter prime–target Stimulus Onset Asynchronies (SOAs).

We surveyed extant models and did not locate any candidates that appeared to have mechanisms that were capable of predicting the present passive–active differences. In ROUSE, we assume that there are multiple possible sources for an activated feature, and the possible sources must be taken into account when calculating the evidence values that enter into a Bayesian decision process. Under- and overestimation of the probability that a feature can be activated by the primes is the key factor that allows ROUSE to predict the differences between active and passive prime processing.

There are explanations for the active–passive differences that can be stated within a ROUSE framework but that differ somewhat from the theory we have proposed. In ROUSE, we assume that one system (perhaps a version of working memory) has full access to the primes and their features, presumably because of the primes being presented well above perceptual threshold (and perhaps because of their being attended to sufficiently well). When the prime knowledge in this system is compared with the knowledge in the word-identification system (which we have assumed contains some of the prime features), the features common to the primes and choices can be identified and preference can be removed or reversed within the decision process. This reasoning suggests an alternate interpretation in which prime knowledge in the working memory system is not fully available in all conditions. In particular, in the passive condition, primes only appear for a half second and participants are not instructed to pay close attention to the primes. As a result, in the passive condition, all the features of the primes might not be available at the time of the 2-AFC decision for comparison with the output of the word-identification system. This implies that some features activated by primes are not discounted and that a preference for prime-related words results. In this view, participants are always excessive in their estimation of α, in both active and passive conditions, but features that are shared between the primes and choices are not discounted when they are not noticed to be present in primes. This idea is further discussed in the subliminal priming section below. In any event, we lack the data at present to separate these two, nearly equivalent explanations.

Orthography Versus Associations

One puzzle in the short-term priming literature has been an almost universal observation of facilitation for semantic and associative priming, whereas orthographic, phonemic, and repetition priming often produce facilitation but sometimes produce deficits (e.g., Dominguez & de Vega, 1997; Humphreys et al., 1988; Lukatela & Turvey, 1996; Meyer et al., 1974). Our implementation of the ROUSE model suggests an answer to this asymmetry of results. Two factors must both be present for the unrelated choice to be favored: First, evidence consistent with the prime must be discounted to a greater degree than optimal (which occurred in our studies in the active priming conditions); second, the similarity of the prime to the choice must be very high. These conditions can both be satisfied for orthographic and repetition priming, but associative and semantic priming necessarily involve lower levels of similarity, at least for the types of tasks we used.

To reiterate the role played in ROUSE by similarity of prime to target, consider again Figure 8 and the attendant discussion. Almost all of the studies in the literature correspond to the target-primed condition of Figure 8. For low similarity primes (i.e., low ρ), corresponding to associative-semantic priming, the target-primed condition is always above the neither-primed condition (i.e., a positive effect of priming) regardless of the passive–active manipulation. Thus, associative-semantic priming will always produce facilitation. For high similarity primes (i.e., high ρ), corresponding to repetition or orthographic-phonemic priming, the target-primed condition is above the neither-primed condition with passive priming and below with active priming (i.e., priming can be positive or negative). Thus, orthographic-phonemic and repetition priming can lead to results in either direction (or a lack of priming) depending on the exact procedures used.

Preference, Prime Duration, and "Subliminal Priming"

All of the present experiments used long-duration, above-threshold primes. The 2-AFC paradigm allowed us to test repeti-

tion priming while controlling strategic responding through our use of repeated foil words. Traditional short-term priming paradigms typically do not test repetition priming with above-threshold primes due to concern for such strategic responding (although see Humphreys, et al., 1988; Ratcliff et al., 1985). However, with associative-semantic, orthographic-phonemic, and other forms of similarity priming, the concern for strategic responding is less and these have been extensively tested with long-duration primes in traditional priming paradigms (e.g., McKoon & Ratcliff, 1992; McNamara, 1994; Meyer et al., 1974; O'Seaghdha & Marin, 1997). Our results and theory can be seen as providing an interpretation of these results in terms of preferential and perceptual effects. It is less obvious that our results are applicable to short duration, subthreshold (i.e., subliminal) primes, but as discussed next, extending the ROUSE model to address this situation is straightforward.

To minimize the use of explicit strategies in short-term priming (i.e., guessing an answer related to a prime), many priming studies use short-duration, masked, subthreshold primes (e.g., Evett & Humphreys, 1981; Marcel, 1983; Perea & Gotor, 1997). In these experiments, participants were not able to identify the prime on most trials, and may not even have been able to determine at an above-chance level whether a prime was presented. The ROUSE model is readily applicable to this situation. If the system cannot identify the presence of a prime, then it is unlikely that any discounting of ON features will occur (and it would certainly be difficult to know which features to discount). Applying ROUSE to this situation is easily accomplished through use of the equations in Figure 4 with the estimate of prime activation, α', set to zero. Without discounting, any extra evidence from the prime, even if very weak, is certain to tip the scales in favor of prime-related words. This does not imply that perceptual facilitation occurs with subliminal priming but that preference can play a role in subliminal priming (indeed it might even play a stronger role because discounting cannot take place). Furthermore, although preferential variability (the factor producing the both-primed deficit) might still occur, the amount of such variability would be very small because, at most, a very small number of features could be activated by the prime. Thus, ROUSE applied to the subliminal priming paradigm would almost certainly predict a "beneficial" result from priming (i.e., a preference for prime-related items). Because discounting has been eliminated, this benefit should hold for all types of related primes (i.e., high and low similarity).

A puzzle in the literature is that the magnitude of facilitation with subliminal priming is often similar to that observed with above threshold priming. ROUSE provides a possible account of such results. For long-duration explicit primes, there is much activation from the primes, but the decision process substantially discounts the evidence from this activation, reducing the size of the beneficial effect (i.e., the preference for prime-related items has been mostly removed). For short-duration subliminal primes, there is little activation from the primes, but there is no discounting, so the net effect might be similar.

Relevant research in evaluative judgments lends some support to ROUSE's interpretation of subliminal priming. Murphy and Zajonc (1993) used near-threshold and above-threshold prime pictures with positive or negative affect; such primes were followed by Chinese letters that had to be evaluated on a positive–negative dimension. Near-threshold primes led to congruent evaluation, but above-threshold primes tended toward an incongruent evaluation. The authors interpreted the results in terms of two different neurological–emotional routes. An interpretation with ROUSE is more parsimonious, holding that the change in the direction of priming is a function of prime availability: With above-threshold primes, participants can use knowledge of prime features to discount interfering prime activation, which results in a neutral preference or even a reversal; with near-threshold primes, the prime features are often unknown and no discounting takes place, which results in a preference for targets of similar affect. This same pattern of results has been found with orthographic–phonemic priming as measured with naming (O'Seaghdha & Marin, 2000). Using high-frequency, beginning-related prime–target pairs (e.g., prime: STORAGE; target: STORY), 400-ms prime exposures resulted in increased latencies, whereas brief (57 ms) forward-masked primes resulted in decreased latencies. The authors interpreted the results in terms of facilitation that was due to shared orthographic–phonemic features for near-threshold primes and phonological competition for above-threshold primes. Again, ROUSE might provide a more parsimonious account.

Short-Term Priming Theories and Preferential Effects

For many years, spreading-activation theories (Anderson, 1983; Collins & Loftus, 1975) and compound-cue theories (Dosher & Rosedale, 1989; Ratcliff & McKoon, 1988) have been the predominant theoretical accounts of short-term priming in lexical decision. These theories map a measure such as summed strength or activation to reaction time; they have not been applied to 2-AFC data. To predict the positive preference in our passive priming 2-AFC data, the mechanisms that apply to target processing could be applied to foil processing.

In spreading-activation theory, a primed foil could also receive preactivation and the choice made, according to which choice-word acquired the greater activation. If only one choice was preactivated, this choice word would tend to be chosen. If both choices were preactivated, the result would depend on details of the theory; some sort of stochastic activation would be necessary to produce preferential variability (and a corresponding both-primed deficit). The primary difficulty in accounting for our results with spreading activation is the lack of a decision mechanism for reversing preference with active prime processing. In addition, we have evidence (Experiments 2 and 4) that the extent of representational overlap between prime and target plays a crucial role in determining the magnitude and even the direction of preference. It is not clear how a one-dimensional construct like strength of activation could handle such effects.

Compound-cue theory shares these difficulties. The theory posits that the prime plus target are used to probe memory. This could be extended to two-prime studies by using both primes plus the target flash to probe memory. It is not clear how to use the theory to deal with 2-AFC testing. One method that is seemingly in keeping with the spirit of the model could use the compound cues to prompt a recall process. The recalled word could then be matched to both choice words. If the primes themselves can be recalled, a preference for a primed word would naturally result. Nonetheless, this theory has no decision mechanism capable of producing a preference reversal with active priming, and it is also

hard to see how the theory could predict the different preference results for varying degrees of prime–target similarity.

A distributed model of short-term priming developed by Masson (1991, 1995) has been extended to the same–different testing procedure (i.e., is the response word the same or different from the target flash?). This model assumes a Hopfield (1982) network of word identification; every word has a multidimensional energy-well learned through Hebbian connections. The effect of a prime is to place the network closer to the location of the target word. Once the target is presented, the system has less distance to travel to settle into the target word attractor. In this manner, facilitation is predicted in reaction-time data. Masson and Borowsky (1998) extended the model to same–different data by assuming that the pattern of activation just prior to the presentation of the response word is stored and then compared with the pattern of activation elicited by the response word. It would be easy to apply this model to 2-AFC data by comparing the stored pattern with each choice word and choosing the word with the best match.

Masson's (1991, 1995) theory assumes a lack of knowledge for the source of activation, which is similar to the ROUSE theory. Thus, the stored pattern reflects the prime as well as target activation, and this produces a preference for prime-related words. Unlike spreading activation and compound-cue theory, this distributed account contains the distributed representation necessary to produce preferential differences with the extent of prime–target overlap. However, to produce preference removal or reversal, Masson's theory requires a discounting mechanism similar to that found in ROUSE. If such a mechanism were implemented, the resultant theory would be some combination of the two theories: the decisional aspects of ROUSE (i.e., discounting) and the dynamics and learning associated with Masson's model. However, the proper Bayesian calculation for discounting in Masson's model would be very complex; ROUSE was formulated with its ON–OFF activation rule to simplify the situation.

Short-Term Priming Theories and Perceptual Effects

We have noted that ROUSE cannot predict the small both-primed benefit with associative–semantic passive priming that was observed in Experiment 4. Although it is tempting to dismiss this result, we have replicated this both-primed benefit in an unpublished experiment (Huber, 2000), and Masson and Borowsky (1998) found a similar semantic benefit when preference was controlled.[8] We assume that a both-primed benefit implies some sort of interaction between prime and target, resulting in improved target perception (however, see Wagenmakers et al., 2000, for a theory that produces both-primed benefits through prime–target interactions that most people would not label as perceptual in nature). Although current theories do not seem to have mechanisms that are capable of explaining our preference findings, the theories tend to include mechanisms that produce prime–target interactions, which raises the possibility that these theories could be adapted to provide a mechanism for explaining the benefits of associative priming on target perception.

Most accounts of spreading-activation theory are not sufficiently specified to determine whether prime-induced target activation is purely additive or interactive. In contrast, compound-cue theory is necessarily interactive; the match between target and each memory trace multiplies the match between prime and each memory trace. If the weighted match values are greater than 1, compound-cue theory predicts facilitation above and beyond an additive effect of the prime.

Masson and Borowsky (1998) performed an experiment similar to our Experiment 4 and found, as we did, a both-primed benefit with semantic priming (they used a same–different judgment and calculated the A′ measure of sensitivity). Masson's distributed model of short-term priming (Masson 1991, 1995) predicted this benefit; the authors defined perception in such a way that the gain was not, in their terms, perceptual. In their theory, the lower level features that are fed into the net before settling occurs are not affected by the prime. However, we define a perceptual benefit as an interaction between prime and target such that more or less evidence is obtained from the target presentation. In Masson's attractor model, it is true that the input is not affected by priming, however the orthographic–phonemic-semantic feature activations interact with target presentation such that the gain in activation that is due to the target input is greater when a related prime has been presented first. We would term such an interaction perceptual.

Regardless of the terminology, Masson's (1991, 1995) distributed model predicts an advantage with priming above and beyond the additive effect of the prime; the prime moves the system to a position near the target, but this advantage does more than provide a fixed advantage to all similarly related words. Once the target is presented, the model settles into the target's attractor more quickly (i.e., more is gained from the target flash). One possible extension of ROUSE that would be capable of predicting a both-primed benefit would be to take a dynamic approach to prime–target processing such as in Masson's distributed model. This would have the advantage of making explicit predictions about priming as a function of the temporal sequence of events. Nevertheless, as mentioned previously, the proper discounting equations become overly complicated if more sophisticated dynamics are introduced.

Whether in Masson's framework or otherwise, it is important to ask why the perceptual benefit should be specific to associative priming. Alternatively, if a perceptual benefit exists for repetition priming but is outweighed by preferential variability, then it is important to ask why the perceptual benefit should be proportionately larger for associative priming. To explore this question, we performed simulations of ROUSE with a simple modification. The target-flash activation parameter, β, was multiplied by some ratio greater than 1 for all features shared by prime and target, in all conditions in which the target was primed. With this addition and with new estimates of parameter values, ROUSE comes close to predicting the same patterns of results that were predicted by the original version. Unfortunately, the new version is not able to predict simultaneously the both-primed deficit with orthographic priming and the both-primed benefit with associative priming (both of these results were found with the passive group in Experiment 4). If this version of ROUSE is required to predict a

[8] Note that both-primed benefits for associative priming (or equivalent evidence) are often difficult to obtain. In our work, the use of multiple primes, or single primes with weaker associations (Ruys, 1998), produced data in which a benefit could not be seen. Note that the failure by Ruys to find a both-primed benefit occurred in a passive priming study that found reliable preferences for choosing the related choice word, for choice words that were similar (orthographically or semantically) or dissimilar.

both-primed deficit for one level of prime similarity, then it necessarily predicts a both-primed deficit for all levels of prime similarity. The key distinction is whether preferential interference or perceptual gain is a stronger influence factor; if a both-primed deficit is observed, preferential interference is stronger and ROUSE predicts this will be true for all types of priming (i.e., all values of ρ). One could assume that the perceptual multiplier is higher for associative features, but this assumption does little more than repeat the data and does not provide a very satisfying answer to the question.

A more informative explanation involves a dissociation between word-level and feature-level effects. Suppose perceptual enhancement arises largely because of preactivation at the word level (perhaps because of top-down support) and not because of activation of lower level features. In fact, the lower level activations harm performance because of preferential variability. For repetition priming, there is a perceptual gain that is due to word-level effects but this is strongly outweighed by the harm caused by variable activation of all of the orthographic and other low-level features. For orthographic priming, there is no perceptual gain, but the amount of harm is somewhat lower than in the repetition priming case because there is less orthographic feature overlap. The net effect could make the amount of both-primed deficit similar in these two cases. For associative priming, there is no harm caused by variability of orthographic–phonemic feature activation, and there might be enough word-level activation to produce the perceptual gains that are seen (i.e., the both-primed benefit). These and related issues would have to be investigated in future research.

Repetition Blindness

The deficits we observed with repetition priming are similar in some respects to the phenomenon of repetition blindness (Kanwisher, 1987). In the repetition blindness paradigm, participants view a RSVP of words and attempt to determine which of the words is repeated (because participants are attempting to identify and remember each word, this corresponds to our active priming). Repetition blindness is the relative failure to detect the repeat of a word for some period of time following its first presentation (the failure is relative because participants are good at detecting the first presentation of the words). We provide a sketch of how the ROUSE model could be extended to capture this result.

To perform this task, the presentation of each word must be identified and labeled as a separate occurrence. The question asked of the word-identification system is whether enough additional activation has accrued to suggest that a new (or repeated) word has been presented. Because source confusion is inherent in the system, as specified in ROUSE, the optimal way to determine whether additional activation has accrued is to discount activation from known previous presentations. For example, suppose a presentation sequence consists of SEAM followed by PIT. Further suppose that one of the features under consideration is *S* at the beginning of a word. Because of lingering activation of this feature, the word-identification system might wrongly conclude that SPIT was presented unless the *S* feature was appropriately discounted.

Because it is initially unknown which, if any, word has been presented during the initial moments of a new presentation, all lexical entries are candidates, and likelihoods for each entry can be calculated in a manner similar to the naming version of perceptual identification (e.g., Schooler et al., 2001). In addition, all features known to have recently occurred are discounted to an appropriate extent (i.e., the features of the most recently presented word are discounted the most). As explained above, discounting is crucial for accurate identification of most words in the RSVP stream, but a side effect is reduced detection for repeated words (i.e., all features of a repeated word are discounted). As the lag between repetitions increases, the estimate of interference (as well as the actual interference) from the first presentation lessens, serving to reduce discounting and allowing greater identification for repetitions.

This extension of ROUSE to repetition blindness might shed some light on competing extant theories of repetition blindness (e.g., Downing & Kanwisher, 1995; Whittlesea, Dorken, & Podrouzek, 1995). In Kanwisher's (1987) type–token, model repeated words access the same word-type but there is a failure to individuate repetitions as separate word-tokens. In terms of tokens, there is a perceptual deficit for the second occurrence. ROUSE uses a similar idea, as applied at the feature level; similar to the idea that both presentations contact the same type without being expressed as separate tokens, there is a failure to individuate the source(s) of activation for a given feature. The same feature (corresponding to types) might be activated by the first or second presentation of a word, and the effort to remove this source confusion through discounting gives rise to the repetition deficit. Inherent in the discounting process is the assumption that some other system (corresponding to tokens) has already noted that the first presentation occurred. In future work, applying ROUSE to repetition blindness data might provide additional constraints on the theory. In addition, the extension of ROUSE to the RSVP procedure demonstrates that discounting is more than just a useful process for priming phenomena and might more generally be an automatic process used by identification systems to overcome source confusion. In this respect, discounting can be thought of as a mechanism that clears out the activation of identified objects such that subsequent objects can be identified with minimal interference.

Long-Term Repetition Priming

In long-term repetition priming, a preference for repeated words was observed with similar but not dissimilar choice words (Ratcliff & McKoon, 1997). Contrary to this result, Bowers (1999) found a preference for repeated words regardless of the similarity between choice words. Our Experiment 3 finds each of these patterns, one for active priming and the other for passive priming. If extended to the long-term priming domain, the decision mechanism contained within ROUSE could provide an explanation for the conflicting results. One could argue that, for some reason (perhaps the nature of the instructions), participants in Ratcliff and McKoon's studies properly estimated the effect of previously studied words (similar to our active priming participants), whereas the participants in Bowers's studies underestimated the effect of previously studied words (similar to our passive priming participants). Ratcliff and McKoon (in press) reported on an experiment that tested the effect of instructions on long-term repetition priming. Participants who were told to passively read a study list displayed a preference for repetitions regardless of choice-word similarity, whereas partici-

pants who were told to actively study the list of words for a later memory test displayed a preference for repetitions only when the choice words were orthographically similar. Ratcliff and McKoon (in press) interpreted this difference in terms of a strategy to choose repeated words given the passive study instructions versus the normal workings of the word-identification system with the active study instructions. Our Experiment 3 provided a short-term version of this instructional difference and demonstrated results that shift. If discounting also occurs in long-term repetition priming, then the ROUSE theory could explain both sets of long-term repetition priming results without implicating strategic responding in either case.

The effect of prime presentation in ROUSE is activational and presumably decays with delays between prime presentation and perceptual identification. As such, the model does not apply to long-term priming. To produce long-term priming, an additional mechanism is needed to reinstate activation for previously seen features. The mechanism used by Schooler et al. (2001) is context matching. The lexical–semantic code for a word is updated with current context features when it is first studied. At test, the choice words have current context features added to their representation, thereby improving the overall match for previously studied words. If the Schooler et al. theory could be modified to include differential discounting of evidence from context features, depending on the instructions, it might be possible to explain the conflicting results of Ratcliff and McKoon (1997) versus Bowers (1999).

Discounting and Recognition Memory

In the study of recognition memory, it has been proposed that when different sources of information contributing to performance are placed in opposition, one source of information may discount the other (e.g., Jacoby, McElree, & Trainham, 1999; Jacoby & Whitehouse, 1989; Jacoby, Woloshyn, & Kelley, 1989). For example, Jacoby and Whitehouse had participants study lists of words for a later memory test. During the memory test, the to-be-recognized test words were immediately preceded either by a subthreshold (16 ms) or well-above threshold (600 ms) presentation of a context word. For the above-threshold presentation, participants were asked to read out loud the context word (similar to active priming). On half the trials, the context word was identical to the subsequent recognition test word, and on half the trials the context word was different. For above-threshold context-word presentations, both hits and false-alarm rates decreased for matching compared with mismatching context words. In contrast, for below-threshold context-word presentations, both hits and false-alarm rates increased for matching context words. Essentially, participants were more willing to respond *old* to the test word when they were unaware that the context word matched, whereas if they were aware, they were less willing to respond old. The authors interpreted this finding in terms of an illusion of familiarity induced by the recent presentation of the matching context word. This illusion of familiarity was allowed through when participants were unaware of the match, whereas it was discounted when participants were aware.

Although a recognition test is clearly different than the perceptual identification procedure we used, the reversal in the effect of the context word is very similar to the passive–active reversal that we observed with repetition priming. Furthermore, the verbal description of these effects provided by Jacoby and Whitehouse (1989) is similar in some respects to our computational theory, although the ROUSE model supposes discounting at the level of features instead of whole words. The predicted similarity effects seen in Figures 8 and 10 and their realization in all our experiments lend support to the idea that feature-level, not word-level, discounting applies in the domain of perceptual identification. Given the complexity of the predicted and observed data as function of prime and choice-word similarity, we would be surprised if participants had any pronounced awareness of this discounting process. In future work, experiments manipulating the similarity of items in an opposition-memory paradigm may help determine the level at which discounting takes place in recognition memory.

Is the ROUSE Theory Too Powerful (i.e., Too Complex)?

We intentionally applied the simplest possible version of the ROUSE theory to our results and did not try to augment it with more sophisticated mechanisms and additional processes that would probably make the model more cognitively plausible. Therefore, the theory can be thought of as a demonstration proof of the power of the core assumptions to predict the findings and as a stand-in for a class of more complicated models that would incorporate the same processes. Nonetheless, the success of the ROUSE model, particularly in accounting for data that at first glance appears incomprehensible, might lead one to question whether the model is too complex (i.e., capable of predicting almost any result). This issue has recently (Myung, Forster, & Browne, 2000) been addressed in the ongoing pursuit of an error measure to reflect model complexity as well as quantitative error. In assessing model complexity, the number of free parameters is one important factor (ROUSE was fine on this dimension because it was implemented with few free parameters). A second and critical factor is the proportion of data space that can be predicted by the free range of the parameters. In other words, is the model limited to a specific subset of predictions? If a small number of free parameters is nevertheless capable of predicting almost any data pattern, the model is too complex and therefore untestable. The technical methods to assess this possibility are still under development, so next we provide more informal arguments that ROUSE is indeed testable.

In this article, we have done our best to explain why ROUSE makes the specific predictions that it does. The curious changes in preference as a function of prime and choice-word similarity are the natural result of the model under an appropriately sized vector of features and with minimal noise. Nevertheless, the skeptical reader might still worry about model complexity. We address a small part of this concern in the following manner. The parameter estimates reported in Table 3 are the best-fitting parameters, tailored appropriately to each study on the basis of different participants, different stimuli, and different procedures. However, we found that it was possible to use a set of default parameter values for all experiments and still capture the correct qualitative pattern

of results.[9] This observation suggests that the ROUSE model fit the data for reasons inherent in the basic structure of its assumptions.

Summary

Through the use of four priming conditions and 2-AFC testing in a perceptual identification task, preferential and perceptual priming effects were distinguished. We verified the existence of perceptual enhancement in the case of associative primes viewed passively, but the effect was quite small and difficult to obtain reliably. On the other hand, preferential effects were large and ubiquitous; the size of such effects increased in magnitude across associative, orthographic, and repetition priming. Preferential effects included at least two components: an increase in variability that reduced performance for all primed conditions and a preference on average for or against choosing prime-related words. When primes were processed passively, the average preference was always in favor of prime-related words. When primes were processed actively, for a considerable period of time, the average preference was much smaller and even reversed for orthographic and repetition priming. Regardless of how the primes were viewed, preferential variability resulted in sizable deficits for orthographic and repetition priming. These patterns suggest that the facilitation typically reported in the literature for short-term priming after passive prime viewing is largely due to a preferential effect rather than a perceptual one. Because our results showed that preference can exist in either direction depending upon the manner in which primes are viewed, traditional short-term priming results are difficult to assess without the use of conditions designed to assess the direction and magnitude of preference. In the absence of these conditions, subtle differences between paradigms could lead to differences in the magnitude and even the direction of priming.

We interpret our results in terms of the ROUSE theory. This theory does not yet attempt to incorporate mechanisms for perceptual enhancements and, therefore, cannot explain the (small) both-primed benefit we found with passive associative priming. In future research, such mechanisms will be appended to the theory. Instead, ROUSE is used in this article to explain the (large) preference effects found in all the studies. The theory supposes features of the choice words are activated by the primes, by the target, and by visual noise, but the participant is unsure which source(s) activated a given feature. Given this to be the case, the participant discounts the evidence provided by an active feature known to have existed in a prime word. Depending on the level of discounting, a level that is assumed to vary with passive and active priming, this simple theory explains a wide range of preferential effects and interactions, both positive and negative.

---

[9] These default parameters were the same as those listed with and used in the creation of Figures 6, 8, and 10. The only parameter that cannot be given a default value is prime similarity, ρ. It makes sense that this parameter is highly dependent on the type of primes used.

References

Anderson, J. R. (1983). *The architecture of cognition.* Cambridge, MA: Harvard University Press.

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,* 1059–1082.

Bowers, J. S. (1999). Priming is not all bias: Commentary on Ratcliff and McKoon (1997). *Psychological Review, 106,* 582–596.

Bowers, J. S., & Michita, Y. (1998). An investigation into the structure and acquisition of orthographic knowledge: Evidence from cross-script Kanjii-Hiragana priming. *Psychonomic Bulletin & Review, 5,* 259–264.

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

Curran, P. J., West, S. J., & Finch, J. H. (1996). The robustness of test statistics to nonnormality and specification error in confirmatory factor analysis. *Psychological Methods, 1,* 16–29.

Dominguez, A., & de Vega, M. (1997). Lexical inhibition from syllabic units in Spanish visual word recognition. *Language and Cognitive Processes, 12,* 401–422.

Dosher, B. A., & Rosedale, G. (1989). Integrated retrieval cues as a mechanism for priming in retrieval from memory. *Journal of Experimental Psychology: General, 118,* 191–211.

Downing, P., & Kanwisher, N. (1995). A reply to Whittlesea, Dorken, and Podrouzek (1995) and Whittlesea and Podrouzek (1995). *Journal of Experimental Psychology: General, 121,* 1698–1702.

Evett, L. J., & Humphreys, G. W. (1981). The use of abstract graphemic information in lexical access. *Quarterly Journal of Experimental Psychology, 33,* 325–350.

Hochhaus, L., & Johnston, J. C. (1996). Perceptual repetition blindness effects. *Journal of Experimental Psychology: Human Perception and Performance, 22,* 355–360.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences, U.S.A., 79,* 2554–2558.

Huber, D. E. (2000). *Perception and preference in short-term word priming.* Unpublished doctoral dissertation, Indiana University, Bloomington.

Humphreys, G. W., Besner, D., & Quinlan, P. T. (1988). Event perception and the word repetition effect. *Journal of Experimental Psychology: General, 117,* 51–67.

Jacoby, L. L., McElree, B., & Trainham, T. N. (1999). Automatic influences as accessibility bias in memory and Stroop tasks: Toward a formal model. In D. Gopher & A. Koriat (Eds.), *Attention and performance XVII: Cognitive regulation of performance: Interaction of theory and application* (pp. 461–486). Cambridge, MA: MIT Press.

Jacoby, L. L., & Whitehouse, K. (1989). An illusion of memory: False recognition influenced by unconscious perception. *Journal of Experimental Psychology: General, 118,* 126–135.

Jacoby, L. L., Woloshyn, V., & Kelley, C. (1989). Becoming famous without being recognized: Unconscious influences of memory produced by dividing attention. *Journal of Experimental Psychology: General, 118,* 115–125.

Johnson, M. K., Hashtroudi, S., & Lindsay, D. S. (1993). Source monitoring. *Psychological Bulletin, 114,* 3–28.

Johnston, J. C., & Hale, B. L. (1984). The influence of prior context on word identification: Bias and sensitivity effects. In H. Bouma & D. G. Bouwhuis (Eds.), *Attention and performance X: Control of language processes* (pp. 243–255). Hillsdale, NJ: Erlbaum.

Kanwisher, N. G. (1987). Repetition blindness: Type recognition without token individuation. *Cognition, 27,* 117–143.

Lombardi, W. J., Higgins, E. T., & Bargh, J. A. (1987). The role of consciousness in priming effects on categorization: Assimilation versus contrast as a function of awareness of the priming task. *Personality and Social Psychology Bulletin, 13,* 411–429.

Lukatela, G., & Turvey, M. T. (1996). Inhibition of naming by rhyming primes. *Perception & Psychophysics, 58,* 823–835.

Lupker, S. J., & Colombo, L. (1994). Inhibitory effects in form priming: Evaluating a phonological competition explanation. *Journal of Experimental Psychology: Human Perception and Performance, 20,* 437–451.

MacMillan, N. A., & Creelman, C. D. (1991). *Detection theory: A user's guide.* New York: Cambridge University Press.

Marcel, A. J. (1983). Conscious and unconscious perception: Experiments on visual masking and word recognition. *Cognitive Psychology, 15,* 197–237.

Martin, L. L. (1986). Set/reset: Use and disuse of concepts in impression formation. *Journal of Personality and Social Psychology, 51,* 493–504.

Masson, M. E. J. (1991). A distributed memory model of context effects in word identification. In D. Besner & G. W. Humphreys (Eds.), *Basic processes in reading: Visual word recognition.* (pp. 233–263). Hillsdale, NJ: Erlbaum.

Masson, M. E. J. (1995). A distributed memory model of semantic priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21,* 3–23.

Masson, M. E. J., & Borowsky, R. (1998). More than meets the eye: Context effects in word identification. *Memory & Cognition, 26,* 1245–1269.

McKoon, G., & Ratcliff, R. (1992). Spreading activation versus compound cue accounts of priming: Mediated priming revisited. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 1155–1172.

McNamara, T. P. (1992). Theories of priming: I. Associative distance and lag. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 1173–1190.

McNamara, T. P. (1994). Theories of priming: II. Types of primes. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,* 507–520.

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

Meyer, D., & Schvaneveldt, R. W. (1971). Facilitation in recognizing parts of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology, 90,* 227–234.

Meyer, D., Schvaneveldt, R. W., & Ruddy, M. G. (1974). Functions of graphemic and phonemic codes in visual word-recognition. *Memory & Cognition, 2,* 309–321.

Murphy, S. T., & Zajonc, R. B. (1993). Affect, cognition, and awareness: Affective priming with optimal and suboptimal stimulus exposures. *Journal of Personality and Social Psychology, 64,* 723–739.

Mussweiler, T., & Neumann, R. (2000). Sources of mental contamination: Comparing the effects of self-generated versus externally-provided primes. *Journal of Experimental Social Psychology, 36,* 194–206.

Myung, I. J., Forster, R. M., & Browne, W. M. (Eds.). (2000). Model selection [Special issue]. *Journal of Mathematical Psychology, 44.*

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* (pp. 264–336). Hillsdale, NJ: Erlbaum.

Nelson, D. L., McEvoy, C. L., & Schreiber, T. A. (1994). *The University of South Florida word association, rhyme and word fragment norms.* Unpublished manuscript.

Norris, D. (1995). Signal detection theory and modularity: On being sensitive to the power of bias models of semantic priming. *Journal of Experimental Psychology: Human Perception and Performance, 21,* 935–939.

O'Seaghdha, P. G., & Marin, J. W. (1997). Mediated semantic-phonological priming: Calling distant relatives. *Journal of Memory and Language, 36,* 226–252.

O'Seaghdha, P. G., & Marin, J. W. (2000). Phonological competition and cooperation in form-related priming: Sequential and nonsequential processes in word production. *Journal of Experimental Psychology: Human Perception and Performance, 26,* 57–73.

Perea, M., & Gotor, A. (1997). Associative and semantic priming effects occur at very short stimulus-onset asynchronies in lexical decision and naming. *Cognition, 62,* 223–240.

Ratcliff, R., Hockley, W., & McKoon, G. (1985). Components of activation: Repetition and priming effects in lexical decision and recognition. *Journal of Experimental Psychology: General, 114,* 435–450.

Ratcliff, R., & McKoon, G. (1988). A retrieval theory of priming memory. *Psychological Review, 95,* 385–408.

Ratcliff, R., & McKoon, G. (1997). A counter model for implicit priming in perceptual word identification. *Psychological Review, 104,* 319–343.

Ratcliff, R., & McKoon, G. (in press). Bias in the counter model. *Psychological Review.*

Ratcliff, R., McKoon, G., & Verwoerd, M. (1989). A bias interpretation of facilitation in perceptual identification. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 378–387.

Ruys, K. (1998). *Orthographic and semantic components of associative priming and bias in forced choice perceptual identification.* Unpublished masters dissertation, University of Amsterdam, the Netherlands.

Schooler, L. J., Shiffrin, R. M., & Raaijmakers, J. G. W. (2001). A Bayesian model for implicit effects in perceptual identification. *Psychological Review, 108,* 1, 257–272.

Schwarz, N., & Bless, H. (1992). Constructing reality and its alternatives: An inclusion/exclusion model of assimilation and contrast effects in social judgment. In L. L. Martin & A. Tesser (Eds.), *The construction of social judgments* (pp. 217–245). Hillsdale, NJ: Erlbaum.

Shiffrin, R. M., & Steyvers, M. (1997). A model for recognition memory: REM—retrieving effectively from memory. *Psychonomic Bulletin & Review, 4,* 145–166.

Shiffrin, R. M., & Steyvers, M. (1998). The effectiveness of retrieval from memory. In M. Oaksford & N. Chater (Eds.), *Rational models of cognition* (pp. 73–95), Oxford, England: Oxford University Press.

Thompson-Schill, S. L., Kurtz, K. J., & Gabrieli, J. D. E. (1998). Effects of semantic and associative relatedness on automatic priming. *Journal of Memory and Language, 8,* 440–458.

Wagenmakers, E. M., Zeelenberg, W., Schooler, L. J., & Raaijmakers, J. G. W. (2000). *A criterion-shift model for enhanced discriminability in perceptual identification: A note on the counter model.* Manuscript submitted for publication.

Wegener, D. T., & Petty, R. E. (1995). Flexible correction processes in social judgment: The role of naive theories in corrections for perceived bias. *Journal of Personality and Social Psychology, 68,* 36–51.

Whittlesea, B. W. A., Dorken, M. D., & Podrouzek, K. W. (1995). Repeated events in rapid lists: Part I. Encoding and representation. *Journal of Experimental Psychology: General, 121,* 1670–1688.

Appendix A

ROUSE Simulation Procedures

The first step in producing a simulation for a particular trial is to determine the number of unique (diagnostic) features in each of the choice words. In Experiments 1 and 2, this was simply N, because we assumed that the choice words shared no features ($\rho = 0$). In Experiments 3 and 4, the similarity between the choice words was manipulated, resulting in a stochastically determined number of features (less than or equal to N) for conditions with similar choice words. In general, if the similarity is ρ, then the probability is $1 - \rho$ that each of the N features will be diagnostic.

The next step is to determine, for each diagnostic feature in both the target and foil, whether the feature is shared with the primes. This will depend on the prime-similarity parameter, ρ (in theory this would be a different similarity parameter than that for choice-word similarity, although across Experiments 2 and 3 the same words were used to produce prime similarity in one case and choice-word similarity in the other case, resulting in our use of the same parameter for both manipulations). If the condition mandates that the target or foil is primed, then ρ will stochastically determine which features in a primed choice word are shared with a prime. Next, knowing which features are shared with a prime, an ON–OFF state is stochastically determined for all the diagnostic features. Target features not shared with a prime are OFF, with probability $(1 - \beta)(1 - \gamma)$ and ON with 1 minus this probability. Target features shared with a prime are OFF with probability $(1 - \alpha)(1 - \beta)(1 - \gamma)$ and ON with 1 minus this probability. For foil features, these same probabilities apply except that $(1 - \beta)$ is removed in each case (i.e., the target presentation is not a potential source of activation).

Last, the odds are calculated. It may be instructive to keep separate the likelihoods for the target and foil. According to Equation 2, the likelihoods are the result of product terms for each feature. However, depending on the parameter values used and the number of product terms, numerical problems can result during the calculation of the products. It is computationally useful to convert Equation 2 into a sum by taking the log (being careful to watch out for boundary conditions when parameters are set to 0 or 1 so that no attempt is made to calculate a log of 0). The individual feature likelihood ratios appearing in Figure 4 are converted by taking the log of each ratio (this can be done at the beginning of the simulated trials and the logs stored for repeated use). Separate sums of the log feature-likelihood ratios are tallied for the target and foil features contingent on whether each features appeared in the primes and whether each feature is ON or OFF. These sums provide the separate target and foil log-likelihood ratios. The target and foil log-likelihood ratios are then compared with each other, and the choice word with the higher value is selected (i.e., with the log transformation, the normative criterion is 0). By subtracting the foil log-likelihood ratio from the target log-likelihood ratio, the log-odds appearing in Figure 6 were produced.

Performance is determined by tallying the proportion of trials on which the log-odds is above 0 (i.e., trials in which the target log likelihood was the greater). If there is a tie (i.e., log-odds is 0 or, alternately, Equation 2 is 1.0), then the evidence in favor of each alternative is equal, resulting in random responding (i.e., the target is chosen with a probability of .5). In practice, participants might adopt a strategy such as "choose the word on the left" when they were unsure or when the evidence in favor of each word is equal. However, because factors such as left–right position of the target were randomly counterbalanced, this would be the equivalent of random responding.

(Appendixes continue)

Appendix B

ROUSE Predictions

Experiment	Type of priming and priming condition	Passive priming				Active priming			
		Obs.	Pred.	*SEM*	*N*	Obs.	Pred.	*SEM*	*N*
	Neither	.692	.700	.023	1,320	.760	.770	.022	1,248
1	Associative								
	both	.671	.679	.025	660	.793	.763	.023	624
	target	.733	.739	.024	660	.790	.790	.025	624
	foil	.670	.659	.024	660	.776	.755	.029	624
	Repetition								
	both	.626	.621	.022	660	.647	.666	.024	624
	target	.770	.771	.019	660	.716	.703	.025	624
	foil	.567	.576	.031	660	.781	.774	.023	624
	Repetition[a]								
	target	.712	.725	.024	660	.686	.711	.028	624
	foil	.594	.590	.029	660	.720	.727	.028	624
2	Orthographic								
	neither	.765	.771	.024	1,872	.802	.783	.017	2,016
	both	.745	.715	.023	1,872	.702	.706	.015	2,016
	target	.845	.839	.015	1,872	.759	.765	.017	2,016
	foil	.648	.663	.030	1,872	.757	.752	.021	2,016
	Repetition								
	neither	.832	.845	.023	624	.780	.782	.027	672
	both	.699	.715	.025	624	.655	.667	.025	672
	target	.859	.884	.016	624	.722	.734	.026	672
	foil	.782	.740	.024	624	.765	.761	.022	672
3	Repetition								
	neither	.732	.744	.026	660	.793	.801	.023	672
	both	.705	.649	.021	660	.719	.696	.022	672
	target	.797	.786	.021	660	.766	.772	.027	672
	foil	.624	.649	.026	660	.768	.776	.024	672
	Repetition[b]								
	neither	.694	.670	.017	1,980	.734	.690	.019	2,016
	both	.602	.613	.014	1,980	.609	.621	.016	2,016
	target	.717	.752	.019	1,980	.721	.741	.018	2,016
	foil	.542	.556	.016	1,980	.568	.601	.022	2,016
4	Associative[c]								
	neither	.738	.753	.015	2,610	.811	.804	.016	1,860
	both	.766	.753	.015	2,610	.804	.801	.019	1,860
	target	.807	.809	.014	2,610	.821	.822	.019	1,860
	foil	.696	.700	.018	2,610	.774	.782	.023	1,860
	Orthographic[d]								
	neither	.717	.709	.019	870	.756	.779	.023	620
	both	.591	.593	.017	870	.700	.671	.018	620
	target	.739	.740	.019	870	.710	.717	.022	620
	foil	.618	.629	.027	870	.779	.773	.024	620

Note. Obs. = observed; Pred. = predicted.
[a] The alternative choice word was associatively primed. [b] Orthographically similar choice words. [c] Associatively similar choice words for neither and both. [d] Orthographically similar choice words for neither and both.

Received November 23, 1999
Revision received June 13, 2000
Accepted June 14, 2000 ■