

Mechanisms of Source Confusion and Discounting in Short-Term Priming 2: Effects of Prime Similarity and Target Duration

David E. Huber
University of Colorado at Boulder

Richard M. Shiffrin and Keith B. Lyle
Indiana University Bloomington

Raushanna Quach
University of Colorado at Boulder

D. E. Huber, R. M. Shiffrin, K. B. Lyle, and K. I. Ruys (2001) tested two-alternative, forced-choice (2-AFC) perceptual identification in a short-term priming task. For repetition priming, passive viewing of primes resulted in a preference to choose repeated words, but actively responding to primes resulted in a preference against choosing repeated words. These results were explained with a computational model, responding optimally with unknown sources of evidence (ROUSE), using the offsetting mechanisms of source confusion and discounting. An analysis of ROUSE revealed conditions under which discounting efficacy should diminish, causing a preference for primed words even with active prime processing. Two new studies confirm 2 such conditions: very short target flash durations and very low similarity between primes and primed choice words. These a priori predictions contrast with the a posteriori data fits of a multinomial model developed by R. Ratcliff and G. McKoon (2001).

Recent use of forced-choice testing in perceptual-identification priming tasks suggests that priming (both short- and long-term) is due more to decisional biases than to enhanced perceptual processing (e.g., Huber, Shiffrin, Lyle, & Ruys, 2001; Ratcliff & McKoon, 1997; Ratcliff, McKoon, & Verwoerd, 1989). A typical test trial consists of a briefly flashed target word (or picture) followed by a pattern mask and then a choice between the target and a foil. The target, the foil, both, or neither choice can be primed by an earlier presentation of the choice word itself or by some word similar to the choice word. If the earlier presentation occurs about 1 s before the test trial, the paradigm is termed *short-term priming*, or if it occurs minutes or more before the test trial the paradigm is termed *long-term priming*. In long-term priming the relation between prime and choice word is usually one of identity; for short-term priming, the relation can be identity, association, or surface-feature similarity.

For long-term priming, and for the traditional method of short-term priming, performance increases following target priming and

decreases following foil priming. Ratcliff and McKoon (1997) observed equal costs and benefits with long-term repetition priming and concluded that priming acts to bias the choice. For short-term priming, Huber et al. (2001) obtained a similar result when participants passively viewed primes briefly, but the bias switched direction when participants actively responded to the primes. This result made it clear that the effect was a form of bias, but the many existing and conflicting uses of the term *bias* led the authors to adopt the more impartial term *preference*.

Huber et al. (2001) explored short-term word priming using the display sequence seen in Figure 1. In this paradigm, two prime words are presented on every trial. The use of two primes allows priming of only the target (target primed), only the foil (foil primed), both target and foil (both primed), or neither target nor foil (neither primed). In the neither-primed and both-primed conditions, both primes are unrelated to the choices, or both primes are equally related to the choices, therefore these conditions are said to be *unbiased*. In the target-primed and foil-primed conditions, only one of the choice words is primed, and these conditions are used to assess the magnitude and direction of preference. Across groups of participants, Huber et al. (2001) manipulated the manner in which primes were viewed. One group viewed the prime words for 500 ms and was told that the primes were merely a warning to prepare for the target flash. This was termed *passive priming*. Another group was shown the primes and asked to give a non-speeded response in relation to the primes; following this response the prime words reversed position on the screen, and the trial proceeded exactly as in the passive condition. This was termed *active priming*. Different studies used different active-priming tasks (e.g., *Do the prime words match in animacy?* and *Could the prime words serve the same part of speech?*) with similar results.

David E. Huber and Raushanna Quach, Department of Psychology, University of Colorado at Boulder; Richard M. Shiffrin and Keith B. Lyle, Department of Psychology, Indiana University Bloomington.

Keith B. Lyle is now at the Department of Psychology, Yale University.

This research was supported by Public Health Service, National Institute of Mental Health (NIMH) Grant 12717, National Science Foundation (NSF) Grant SBR 9512089, NIMH Training Grant MH19879, NSF Grant KDI/LIS IBN-9873492, and Office of Naval Research Grant N00014-00-1-0246.

Correspondence concerning this article should be addressed to David E. Huber, Department of Psychology, University of Colorado at Boulder, 345 UCB, Boulder, Colorado 80309-0345. E-mail: dhuber@psych.colorado.edu

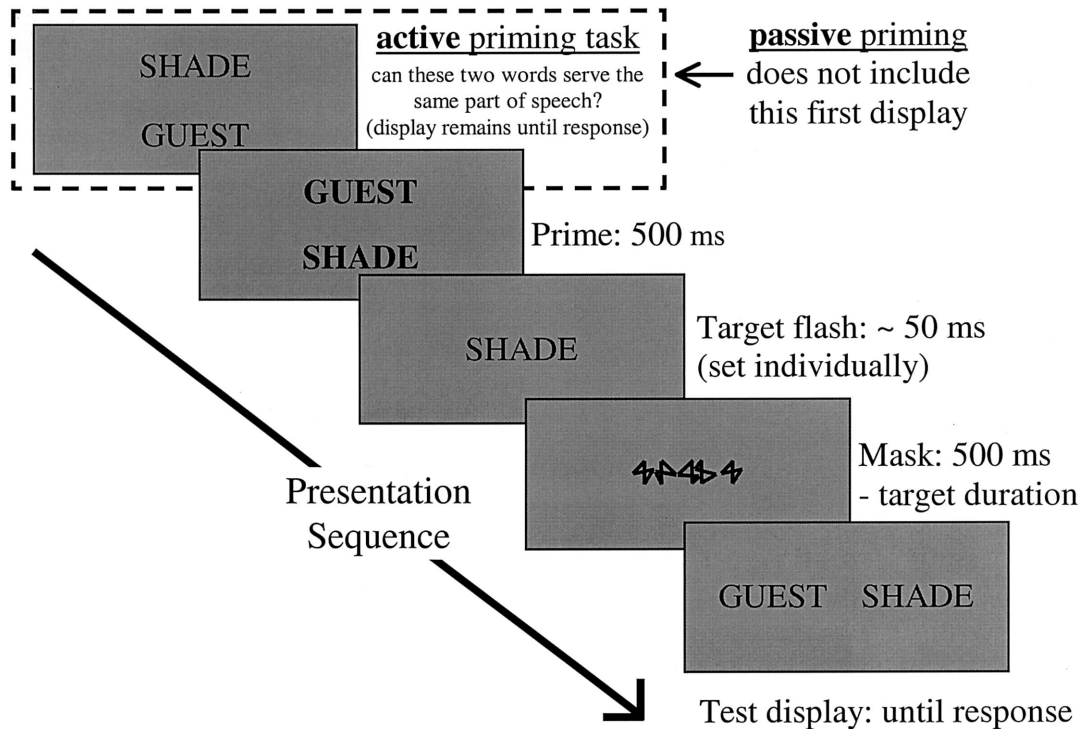


Figure 1. The sequence of visual displays used in the present studies and in Huber et al. (2001). The particular words in Figure 1 provide an example of repetition priming with the both-primed condition. Varying the relationship between the prime words and the choice words, such that only the target, only the foil, or neither choice word is primed produces the target-primed, foil-primed, and neither-primed conditions.

The pattern of results obtained by Huber et al. (2001) for short-term word priming deviated from the patterns for long-term priming in two important ways. First, as mentioned previously, the direction of costs and benefits switched depending on how the primes were processed. With passive priming there was a preference to choose repeated (i.e., primed) words, whereas with active priming there was a preference against choosing repeated words. This change can be seen in the preference conditions in the right-hand panels of Figure 2, which shows the repetition-priming conditions of Huber et al. (2001, Experiment 1). Second, priming of both choices resulted in performance deficits for both passive and active prime processing. In Figure 2, this can be seen in the unbiased conditions in the left-hand panels, for which both-primed deficits (compared with the neither-primed condition) are observed for both passive and active priming. In addition to both-primed deficits, a deficit was often found when comparing the average of the target-primed and foil-primed conditions with the neither-primed condition, although such an effect was small in the data shown in Figure 2.

Huber et al. (2001) explained such findings with the responding optimally with unknown sources of evidence model (ROUSE). In ROUSE, preference for primed choices is the result of source confusion between features perceived because of (a) presentation of the primes and (b) the flash of the target. General priming deficits result from variability in this source confusion. The system attempts to correct for such source confusion by discounting evidence associated with primed words. More discounting is assumed to occur with active priming than with passive priming.

In this article we present new studies designed to test several a priori predictions of the ROUSE model. In addition, we analyze the predictions of a model proposed by Ratcliff and McKoon (2001), a model that uses mechanisms analogous but not identical to ROUSE's source confusion and discounting. Specifically, our studies tested predictions derived from ROUSE concerning circumstances in which discounting ought to become ineffective, resulting in a preference for primed words even with active priming. This outcome ought to occur in each of the following situations: (a) when the similarity of the choice words with each other is high, (b) when the target-flash duration is very short or zero, and (c) when the similarity of the primes to the choice words is low. Huber et al. (2001) observed the predicted discounting failures for choice-word similarity and similarity of the primes to the choice words. We report two new studies: One study tests these predictions by manipulating target-flash duration, and a second more precisely tests the prediction concerning similarity of the primes to the choice words.

Both ROUSE and Ratcliff and McKoon's (2001) multinomial model contain source confusion and discounting, suggesting these complementary mechanisms are necessary to account for the complex pattern of results and are possibly a requirement for accounts of short-term priming studies generally. There are many reasons to think that the visual processing system might confuse inputs that occur in very close temporal (and spatial) proximity. If this does occur, then it might be assumed that the decision part of the visual processing system has acquired mechanisms to cope with these confusions. One potential coping mechanism is what we imple-

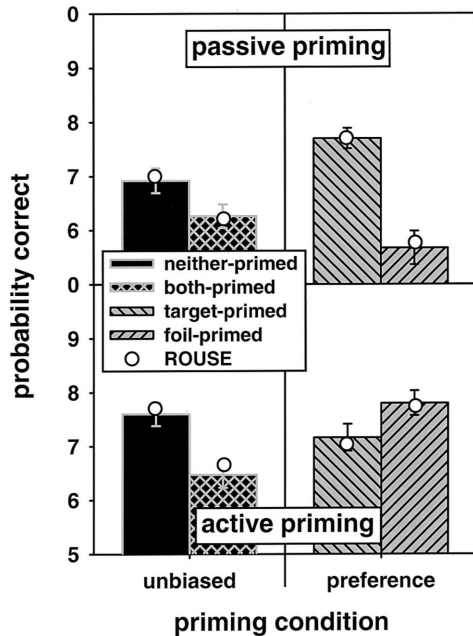


Figure 2. The repetition-priming conditions from Huber et al. (2001, Experiment 1). Passive versus active priming was a between-subjects manipulation. The white circles show the results of the best-fitting responding optimally with unknown sources of evidence model (ROUSE) parameters. The error bars represent two standard errors of the mean. The best-fitting parameters were α (passive) = .073, α' (passive) = .054, α (active) = .085, α' (active) = .152, β (passive) = .034, and β (active) = .054. The estimates of target activation (β') were set to their actual values. Both the estimate and the actual probabilities of noise activation (γ' and γ) were set to the default value of .02.

ment as the discounting of evidence associated with recently presented features. For example, the identification process might implement discounting to clear the system following identification, thus allowing the future identification of highly similar or identical objects. Huber et al. (2001) discussed the implications of such a discounting mechanism in terms of a variety of paradigms, including subliminal priming, repetition blindness, long-term priming, and recognition memory. Using short-term priming as a model task, Huber and O'Reilly (in press) proposed a physiological mechanism that produces appropriately discounted neural activation levels, and they discussed applications to the above phenomena, as well as visual masking, image aftereffects, perception of ambiguous stimuli, semantic satiation, affective priming, and inhibition of return.

From this broad perspective, it could be asked what paradigmatic circumstances require the use of discounting. Huber et al. (2001) demonstrated that actively responding to prime words results in a high degree of discounting at a subsequent identification stage. Huber, Shiffrin, Quach, and Lyle (in press) went on to demonstrate that increasing prime duration, in the absence of overt prime responding, elicits results similar to, but smaller than, active priming. Furthermore, in those studies, the accurate identification of prime words was a strong predictor of discounting (i.e., choice words that were correctly identified as having also been presented as prime words were more strongly discounted). The two new

studies reported here, combined with the orthographic choice-word-similarity experiment of Huber et al. (2001), identify three additional requisites for strong discounting: (a) low similarity between the choice words, (b) sufficiently long target durations, and (c) sufficiently high similarity between primes and choices.

ROUSE

The general framework of the ROUSE approach assumes words are represented as vectors of features. Presentation of the primes activates the features contained in the prime words. Later, some of this activation remains, becoming mixed with the activation induced by the target flash and other, randomly induced, sources of noise activation (e.g., the pattern mask). This source confusion (i.e., unknown sources of evidence) results in a preference for primed words because of extra, prime-induced feature activation. The unwanted effects of source confusion are counteracted, on average, through the discounting of features known to have been presented in prime words (i.e., responding optimally). Because the primes, rather than the target flash, may have been the source of activation, a lowered level of evidence is assigned to primed features. With excessive discounting, source confusion is excessively counteracted, resulting in a preference against primed words.

First, we consider source confusion through feature activation in greater detail, as seen in Figure 3. The probability that noise will activate any feature in either choice word is γ , and the probability that the target flash will activate a feature in the target choice word is β . If a prime feature matches a choice-word feature, then that choice-word feature will become active with probability α . A mediating parameter, ρ , representing the proportion of shared features, probabilistically determines which features are shared between a particular prime and a particular related choice word ($\rho = 1$, for repetition priming). This situation results in an activation of primed features that is indistinguishable from an activation due to perception of the target flash. Thus, source confusion produces a preference for primed choice words. Furthermore, because activation of prime features is probabilistic, the number of choice-word features that become active due to presentation of a related prime is variable. This variability injects noise into the decision process, lowering primed performance (i.e., both-primed deficits).

Next, we consider feature discounting in greater detail, as seen in Figure 4. Because the prime words are always presented well above threshold, for a minimum of 500 ms, we assume that the system knows which features appeared in prime words (but not which features are active as a result of priming). If one of the choice-word features is the same as a feature in a prime word, activation could have arisen from the prime rather than the target flash, therefore such a feature should be trusted to a lesser extent. In this way, primed features are discounted, and therefore primed words are discounted.

The mathematical formalization of discounting is depicted in Figure 4. Each panel gives the likelihood ratio that a feature belongs to the target, instead of the foil, interpretable as evidence in favor of the choice word containing that feature. The calculation of the feature-likelihood ratio depends on the state of activation of the feature and whether the feature appeared in a prime. The actual activation-confusion parameters are α , β , and γ , and the estimates

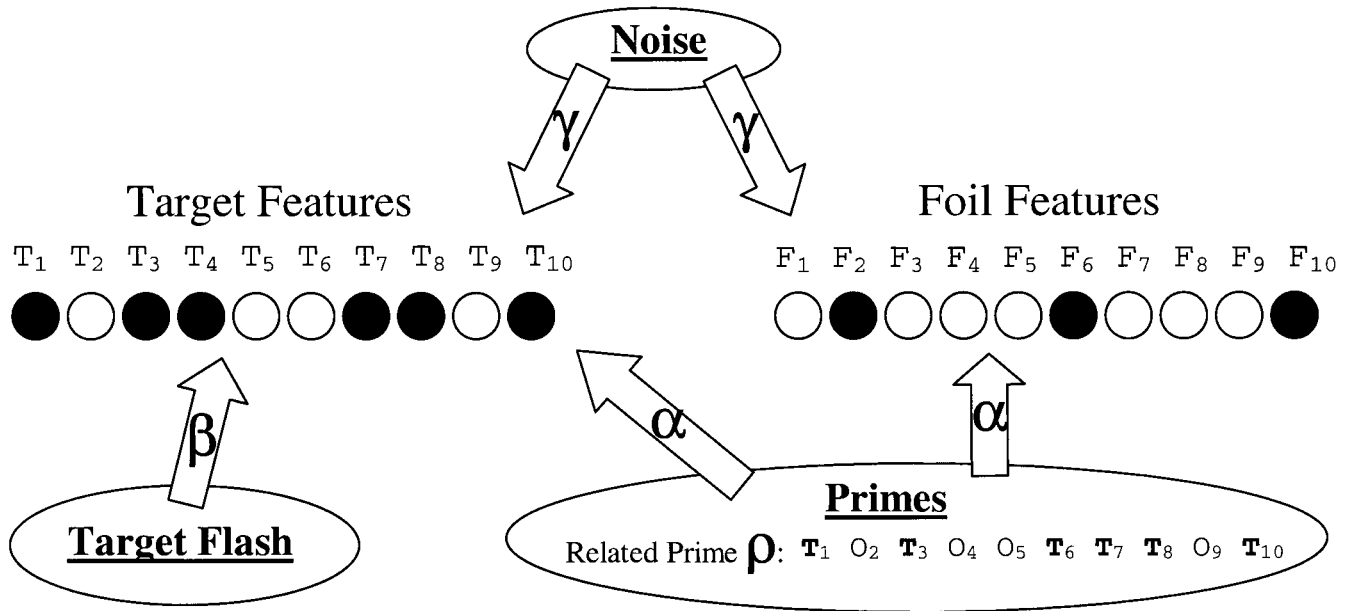


Figure 3. The three sources of choice-word feature activation in the responding optimally with unknown sources of evidence model (ROUSE). With probability β , every feature in the target word is activated. Features shared between a prime and a choice word are activated with probability α , and a mediating parameter, ρ , probabilistically determines the proportion of shared features. Noise activation is applied to all features with probability γ . Ten features per word are shown, although 20 features per word were used in the simulations. The figure shows an example of priming the target with an intermediate level of priming ($0 < \rho < 1$). For this particular stochastic simulation, six features were shared between prime and target, as indicated by the bold Ts in the prime word representation. T = target feature; F = foil feature; and O = other feature, appearing in neither choice word.

		State of Feature Activation	
		inactive	active
Feature Appeared in Prime	no	$\frac{(1-\gamma')(1-\beta')}{(1-\gamma')} = (1-\beta')$ <p>A. less than 1.0</p>	$\frac{1-(1-\gamma')(1-\beta')}{1-(1-\gamma')}$ <p>C. greater than 1.0</p>
	yes	$\frac{(1-\gamma')(1-\alpha')(1-\beta')}{(1-\gamma')(1-\alpha')} = (1-\beta')$ <p>B. less than 1.0</p>	$\frac{1-(1-\gamma')(1-\alpha')(1-\beta')}{1-(1-\gamma')(1-\alpha')}$ <p>D. closer to 1.0</p>

Figure 4. The four possible feature likelihood ratios in the responding optimally with unknown sources of evidence model (ROUSE) contingent on the state of feature activation and whether the feature appeared in a prime. The variables are labeled with primes because they are estimates of the true activation probabilities. Inactive features yield feature likelihood ratios less than 1.0 (Cells A and B) and provide evidence against the choice word to which they belong. Active features yield feature likelihood ratios greater than 1.0 (Cells C and D) and provide evidence in favor of the choice word to which they belong. In ROUSE, word likelihood ratios are determined by multiplying the constituent feature likelihood ratios, and the choice word with the higher word likelihood ratio is chosen. The dashed cell, labeled D, shows the discounted feature likelihood ratio, which is closer to 1.0 than the cell labeled C. This discounting occurs because active features that also appeared in a prime should be mistrusted to an extent dictated by the estimate of prime activation (α').

of these used by the system to calculate likelihood ratios are α' , β' , and γ' . Discounting is seen in the lower right-hand panel: The evidence here incorporates the estimate of source confusion, α' , in such a way that a higher estimate produces less evidence (i.e., the higher the probability that a perceived feature was perceived because of the prime, the less evidence that feature provides). The overall evidence in favor of a particular choice word being the target is calculated by multiplying together all the component feature likelihood ratios to provide a choice-word likelihood ratio (i.e., the likelihood that a particular choice word is the target instead of the foil). The choice word with the higher word-likelihood ratio is chosen and, in the event of a tie, a word is randomly selected.

It is important to emphasize that the equations in Figure 4 use estimates for the probability of activation by the primes, target, and noise (α' , β' , and γ' , respectively) rather than the true probabilities of activation. The true values determine the distribution of activated features of different types. However, we assume that the word-identification system does not have access to the true probabilities of activation and must use estimates to assign evidence values to the decision process.

In this article we focus on the effects of misestimating α . We assume that β and γ are estimated correctly, and we have demonstrated through simulations that misestimates of β and γ have a small, barely noticeable effect. In contrast, misestimates of α have a large effect on performance in the target-primed and foil-primed conditions. When α' is set to its true value (i.e., $\alpha = \alpha'$), and

enough features are activated by the various sources, discounting is set to its optimal level, and performance in the target-primed and foil-primed conditions is equal. When α' is set less than α , such as with passive priming, discounting is insufficient, resulting in target-primed performance better than foil-primed performance. Conversely, when α' is set greater than α , such as with active priming, discounting is too great, resulting in target-primed performance worse than foil-primed performance.

Thus, it is through the mechanisms of source confusion and discounting that ROUSE accounts for the pattern of observed data: (a) Source confusion (as determined by prime activation, α , and prime similarity, ρ) produces a tendency to mistake a prime for the flash; (b) variability in source confusion degrades both-primed relative to neither-primed performance; and (c) discounting tends to undo the tendency to mistake a prime for the flash. (Too little discounting, assumed to occur in passive conditions, causes a tendency to choose a choice word related to a prime; too much discounting causes a tendency to choose a choice word not related to a prime.)

The Multinomial Model

Ratcliff and McKoon (2001) developed an alternative model to account for the results obtained by Huber et al. (2001). This model uses mechanisms similar to ROUSE's source confusion and discounting but places them within the framework of a multinomial decision tree. Source confusion, in Ratcliff and McKoon's (2001) view, is the result of confusion within short-term memory such that a prime is sometimes remembered as the target. Discounting in the multinomial model is an all-or-none process that is probabilistically applied to prime-related choice words. There is a decision branch for discounting, another for source confusion, and a third for target perception (see Figure 5). The simplest form of the multinomial model orders the decision tree such that discounting occurs before source confusion, which in turn occurs before per-

ception. However this does not necessarily imply the order of occurrence, because there are other tree representations that yield identical equations.

Discounting is probabilistically applied to the target and the foil according to the parameters D_T and D_F (which are set to the value D for primed words and zero for unprimed words). If only one choice word is discounted the other choice word is selected, and if both are discounted a choice word is randomly selected (i.e., a guess). Therefore discounting results in a preference against primed words. Different types of priming (e.g., associative priming vs. repetition priming) can result in different discounting probabilities (i.e., different D values). For instance, Ratcliff and McKoon (2001) assumed that there is no discounting of associatively primed words but that there is discounting of orthographically primed words. Furthermore, they assumed discounting applies to active priming only and is not used in passive priming. To account for the experiments reported by Huber et al. (in press), in which discounting was observed without an active-priming task, one could relax this last assumption, allowing a more general form of the multinomial model.

If discounting does not apply to either choice word, then the source-confusion decision branch is reached. Similar to discounting, source confusion is differentially applied to each choice word depending on whether that choice word is primed ($S = 0$, for an unprimed choice word). In addition, it is assumed that source confusion is some function of prime similarity, and therefore, the type of priming affects the level of source confusion (e.g., S was estimated to be higher for repetition priming than for associative priming). Contrary to discounting, source confusion results in the selection, rather than the rejection, of prime-related choice words. Therefore source confusion results in a preference in favor of primed words, and correspondingly, target-primed performance is greater than foil-primed performance. This explains the preference switch between passive and active priming, because only with

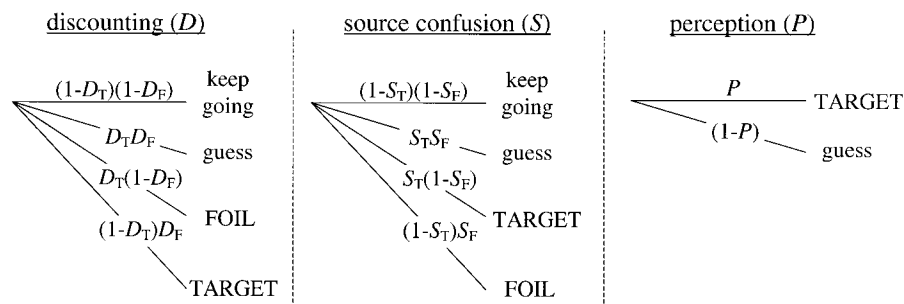


Figure 5. A decision tree giving rise to the equations found in the multinomial model. Probability-correct performance is obtained by multiplying through all the paths that result in selection of the target word. In the case of a guess, the target is selected with probability .5. D_T is the probability of discounting the target word, and D_F is the probability of discounting the foil word. A discounted word results in the selection of the other choice word. The D_T and D_F values are set to the probability of discounting a primed word, D , or set to zero depending on the priming condition under consideration. The same rules apply to the source confusion branch, but the confused choice word is selected instead of rejected. S_T is the probability of confusing the target word, and S_F is the probability of confusing the foil word. Because discounting takes precedence over source confusion, which in turn takes precedence over perception, source confusion and discounting block the use of target perception, resulting in both-primed deficits and a deficit in the average of the target- and foil-primed conditions. The discounting branch is used in active priming, but not passive priming, explaining the observed preference changes.

active priming is the discounting branch invoked to counteract source confusion.

If neither choice word is discounted or source confused, the perceptual-decision branch is reached, and with probability P the target flash is perceived, and the target is correctly chosen. If the target is not perceived, then a choice word is randomly selected.

For given values of D , S , and P , performance can be determined from Figure 5 by adding up the probabilities of all decision paths that ultimately result in selection of the target (guessing results in selecting the target with probability .5). The path probabilities are determined by multiplying the appropriate branch probabilities at each point along a given path. Because discounting and source confusion take precedence over target perception, the source confusion and discounting decision branches block the use of target perception and result in both-primed deficits as well as a deficit in the average of the target-primed and foil-primed conditions. This is analogous to the general priming deficits in ROUSE that are due to source-confusion variability.

ROUSE and the Multinomial Model: A Priori Versus A Posteriori Predictions

The above descriptions of ROUSE and the multinomial model are meant to give an intuitive grasp of the essentials of the two models. The brevity was intended to reduce redundancy with the original articles, and the reader is referred to Huber et al. (2001) and Ratcliff and McKoon (2001) for more detailed expositions. Nonetheless, we hope the following arguments comparing the two models can be understood on the basis of the present descriptions. The two models fit the data equally well, on the basis of roughly equal number of parameters, therefore a choice between the two models must depend on other factors and, in a sense, becomes a matter of taste. In our view, the multinomial model provides an easier transition from parameters to predictions, involving a multiplication of probabilities along branches, whereas ROUSE has a clearer conceptual linkage of processes to parameters (in many

situations it is not clear how the decision-tree probabilities in the multinomial should vary as a function of various empirical manipulations). Different investigators might or might not agree with this characterization and might weight these two factors differently. In our view, the best reason for preferring ROUSE lies in another direction: ROUSE makes advance predictions that if not borne out by the data would contradict the model, whereas the multinomial model is largely descriptive, handling new findings by adopting new assumptions. That is, ROUSE makes a priori predictions, whereas the multinomial model makes a posteriori predictions. Next, we explain the basis of certain of these ROUSE predictions and report studies designed to test them.

ROUSE and the Effectiveness of Discounting

Huber et al. (2001) found that the best fits resulted from setting noise activation to a small fixed value ($\gamma = .02$). Given this assumption, ROUSE predicts there are experimental manipulations that will eliminate the effectiveness of discounting, resulting in a preference for primed words even in the active-priming condition (i.e., even with $\alpha' > \alpha$).

For simplicity, one should consider only target features and ignore noise activation (which can be shown to produce no important distortions when γ is small). The Venn diagram of feature activation and evidence evaluation appearing in Figure 6 illustrates this situation. In the Venn diagram, lighter shaded regions represent evidence in favor of the target word. The left-hand panel shows target evidence without priming, in which case there are only target-activated features, because there are no features shared between primes and the target. The middle panel shows target evidence when primes have a similarity, ρ , to the target but when discounting is not used. Here there are extra features activated by primes and, therefore, more target evidence resulting in an increased tendency to choose the target, relative to the first panel. In principle the two circles in the second panel should overlap, but we do not depict the small region of overlap, considering that reason-

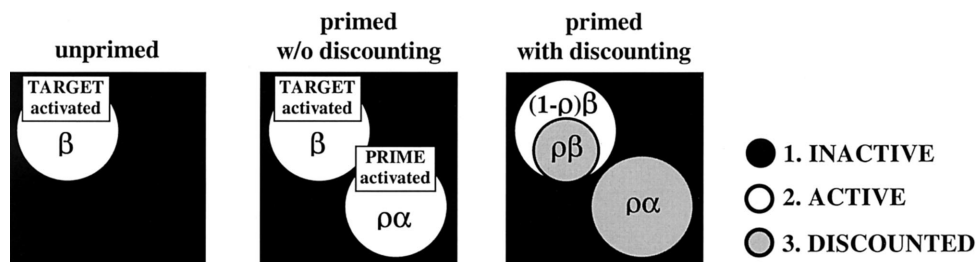


Figure 6. The space of feature activation and evidence evaluation for a target word under the assumption of no noise activation ($\gamma \approx 0$). The purpose of this figure is to demonstrate why the erroneously discounted target-activated features (area $\rho\beta$) are critical for effective discounting. The equations appearing in each area represent the proportion of features that fall into that classification. Because activation probabilities giving rise to reasonable performance levels are very small, the sizes of the displayed circles are larger than they should be, and features activated by multiple sources are not portrayed. The coloring corresponds to the feature likelihood ratios, as seen in Figure 4, with lighter colors representing greater values (i.e., more evidence). Comparing the first and third panels demonstrates that the discounting of prime-activated features (area $\rho\alpha$) cannot by itself eliminate the extra evidence provided by prime activation. It is only with the discounting of the target-activated features (area $\rho\beta$) that discounting results in the same or less evidence than that seen in the unprimed panel. Effects of variability are not portrayed, although such effects become important, particularly as an area becomes sufficiently small such that no features are likely to fall into that classification. w/o = without.

able parameter values make activation of a feature by both the flash and the prime very unlikely. The right-hand panel also shows a primed target, but evidence is calculated with discounting. There are two effects of discounting: As seen in the gray region labeled $\rho\alpha$, discounting lessens the extra evidence resulting from prime activation (white turns to gray, representing the “mistrust” of these features). Nonetheless, if this were the only effect of discounting, priming would still produce a target preference, because even discounted evidence (the gray $\rho\alpha$ circle) is better than no evidence (compared with no prime-activated features in the left-hand panel). However, when discounting is used, all features that are shared with the primes must be discounted, even those that were actually activated by the flash of the target word. Such features are represented by the gray $\rho\beta$ circle. Discounting these valid, target-activated features amounts to “throwing the baby out with the bathwater,” which is unavoidable because the source of activation is unknown. This effect lowers evidence for targets relative to the left-hand panel.

This description demonstrates that preference reversals from discounting depend on discounting target-activated features (the gray $\rho\beta$ circle). A manipulation making this type of discounting less likely (i.e., one shrinking the $\rho\beta$ circle) can reduce, or even eliminate, the effectiveness of discounting. This can be accomplished by reducing the number of features that could fall in this region. One way to identify the critical variables involves calculating the probability that the $\rho\beta$ circle completely disappears for a given set of words on a particular trial. Each target feature has a probability $\rho\beta$ of existing in a prime and also receiving activation from the target flash. Thus if N is the number of target features, Equation 1 expresses the probability of observing no such features on a given trial:

$$(1 - \rho\beta)^N. \quad (1)$$

According to this equation, reducing ρ , β , or N will increase the probability that the $\rho\beta$ circle is absent, thereby reducing the efficacy of discounting. Such an effect can produce a preference for primed words even when α' is greater than α , such as might occur with active priming.

The three studies we discuss use manipulations that reduce N , β , and ρ . Making the choice words similar to each other in effect reduces N because this manipulation reduces the number of diagnostic features (features shared between the choice words play no differential role in the decision process). This study was reported in Huber et al. (2001). Reducing the flash time lowers β (Experiment 1). Making the primes less similar to the choice words reduces ρ (Experiment 2). For each experiment we show ROUSE predictions using default parameter settings, where *default* means parameters that captured the qualitative trends in the previous experiments. In the case of the two new experiments, the predicted curves were produced in advance of data collection. Following each experiment we quantitatively model the results, using both ROUSE and the multinomial model, to contrast the behavior of the models.

Reducing N by Increasing Orthographic Choice-Word Similarity

We begin by revisiting Experiment 3 of Huber et al. (2001), which manipulated the level of similarity between the choice

words, but we recharacterize the ROUSE predictions in terms of reducing the efficacy of discounting. Increasing the similarity between the choice words produces an increase in the number of shared features (e.g., *lied* and *died* share three letters, but *lied* and *sofa* share no letters). In ROUSE, shared features contribute the same likelihood to each choice and, therefore, play no role in the decision (i.e., including a shared feature only results in multiplying the odds in favor of the target by 1.0). In contrast, the nonshared features are potentially diagnostic (e.g., *l* vs. *d*, in comparing *lied* with *died*) and often differentially affect the odds in favor of the target. The number of diagnostic features drops as similarity is increased, producing one of the factors that should reduce the effectiveness of discounting. Thus, for active priming, ROUSE predicts that sufficiently similar choice words should cause a switch from a preference against repeated words to a preference for repeated words.

Figure 7 shows simulated predictions of the ROUSE model, for repetition priming, using default parameter settings (these parameter values are listed in the figure caption).¹ As seen in the figure, increasing the similarity of the choice words lowers performance, as might be expected. The key prediction is seen in the lower right-hand panel of Figure 7, which shows the target-primed and foil-primed conditions when primes are processed actively. For dissimilar choice words (i.e., choice-word similarity = 0; e.g., choosing between *lied* and *sofa*), there is the usual preference against choosing a repeated word. However, as choice-word similarity increases (e.g., *lied* and *died*), the direction of preference reverses, and there is a preference for choosing repeated words. This reversal is predicted despite the overestimation of prime activation (i.e., even with excessive discounting). Reducing the number of diagnostic features produces a significant number of cases in which there are no (or very few) discounted target features, and therefore, discounting is ineffective.

Experiment 3 of Huber et al. (2001) tested this prediction. For the most part, the methods used in this experiment were identical to those used in the two new studies reported in the present article. Huber et al.'s (2001) experiment included both passive and active priming as a between-subjects manipulation, and active priming was induced by a part-of-speech matching task. High choice-word similarity was achieved with word pairs that shared all but one of the same letters in the same letter positions. Half the words were four letters long, and the other half were five letters long.

The results of this study are replotted in Figure 8. The critical finding is the large difference between similar and dissimilar choice words that is seen in the lower right-hand panel (target- and foil-primed active-prime processing). This finding contrasts with the passive-priming results in the upper right-hand panel, in which similarity of the choice words makes little difference. The pattern of findings qualitatively matches those illustrated in Figure 7. The ROUSE predictions shown in Figure 8 are a best fit with the parameter values given in the figure caption (e.g., choice-word similarity was set to $\rho = .7$). In sum, the results validate the first of three ROUSE predictions concerning discounting efficacy.

¹ These graphs, and all other newly reported ROUSE simulations, were generated using the analytic method for simulating ROUSE, outlined by Huber (2002).

Ratcliff and McKoon (2001) addressed these results with their multinomial model. If the same level of discounting (D) was assumed for both dissimilar and similar choice words, the model failed to capture the observed data. The situation was remedied by assuming that with highly similar choice words participants occasionally discount the wrong choice word. With an extra parameter for this erroneous discounting (and separate levels of source confusion, S , for dissimilar and similar choice words), the multinomial model was able to fit the preference changes about as well as ROUSE. However, ROUSE obtained its fit using a single level of source confusion (α) and a single level of discounting (α') for both similar and dissimilar choice words, with the predictions naturally arising from the reduced efficacy of discounting when choice-word similarity is high.

Experiment 1: Reducing β by Reducing Flash Duration

In ROUSE, another way to reduce the efficacy of feature discounting involves lowering the probability of feature activation by the flashed target (i.e., lowering β). In Experiment 1 we tested this prediction by varying target-flash duration; flash duration was set at its usual duration, at half that duration, or at zero.²

The qualitative ROUSE predictions, for repetition priming, as a function of decreasing target-flash duration (β), are shown in Figure 9. The usual default parameter settings were used, as reported in the figure caption. The estimate of target activation, β' ,

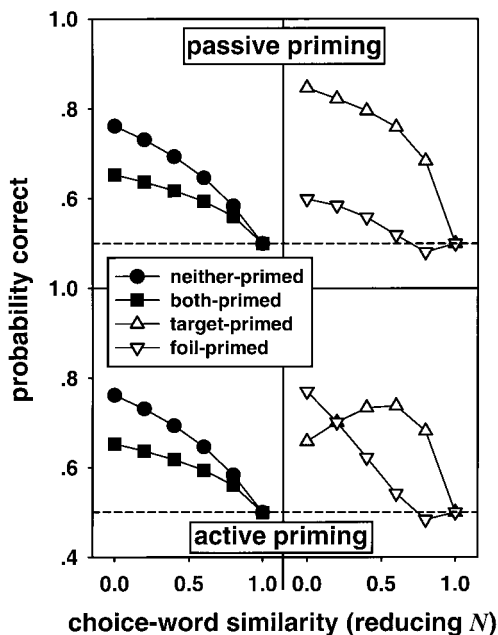


Figure 7. The responding optimally with unknown sources of evidence model (ROUSE) predictions for increasing the similarity of the choice words to each other with repetition priming ($p = 1$). The horizontal dashed line indicates chance performance. These predictions were generated using the analytic method for simulating ROUSE (Huber, 2002) with the default parameters that capture the qualitative trends in all previous experiments. These default parameters are $\beta = \beta' = .05$, $\gamma = \gamma' = .05$, $\alpha = .1$ (both passive and active), $\alpha' = .05$ (passive), and $\alpha' = .3$ (active). As explained in the text, choice-word similarity probabilistically determined the number of diagnostic features. N = number of target features.

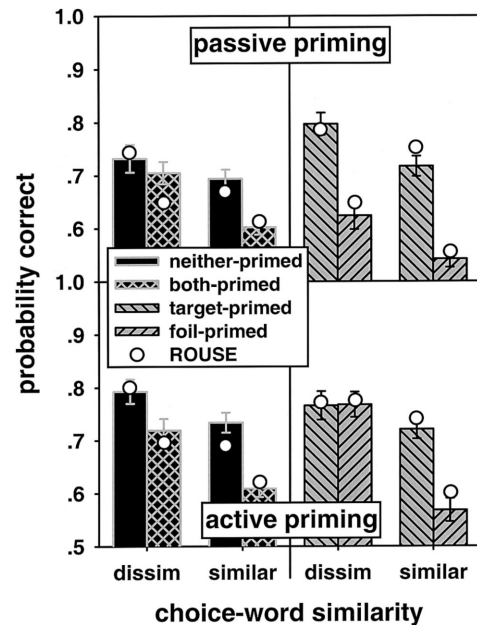


Figure 8. Huber et al.'s (2001) Experiment 3 with repetition priming. Similar choice words shared all but one letter, in position, and dissimilar (dissim) choice words were randomly paired. Passive versus active priming was a between-subjects manipulation. The white circles show the results of the best-fitting responding optimally with unknown sources of evidence model (ROUSE) parameters. The error bars represent two standard errors of the mean. Different pools of words were used in the similar and dissimilar choice words, and therefore, different probabilities of target activation were allowed for these conditions. The best-fitting parameters were α (passive) = .112, α' (passive) = .097, α (active) = .090, α' (active) = .125, β (passive-dissimilar) = .055, β (passive-similar) = .074, β (active-dissimilar) = .062, and β (active-similar) = .083, and the proportion of shared features between similar choice words was .7. The estimates of target activation (β') were set to their actual values. Both the estimate and the actual probabilities of noise activation (γ' and γ) were set to the default value of .02.

was kept at the default value of .05, whereas the true value of target activation, β , decreased. However, setting β' equal to β at each target duration produced nearly identical predictions. Again the critical predictions are found in the lower right-hand panel (target-primed and foil-primed performance for active priming). The left side of that graph corresponds to target-flash duration (β) set to its usual default setting of .05, which produces approximately 75% correct performance in the neither-primed condition. Because of excessive discounting, target-primed performance is lower than foil-primed performance. However, moving to the right on the graph corresponds to decreasing β , and the preference changes direction because of the reduced efficacy of discounting.

In addition to these preference changes with active priming, on which we have been focusing, the breakdown in discounting efficacy also produces preference changes with passive priming

² In the zero condition, nothing was flashed prior to the mask and participants were told that sometimes the flash would be so short that they might not realize anything was flashed but that they should do their best to respond correctly anyway.

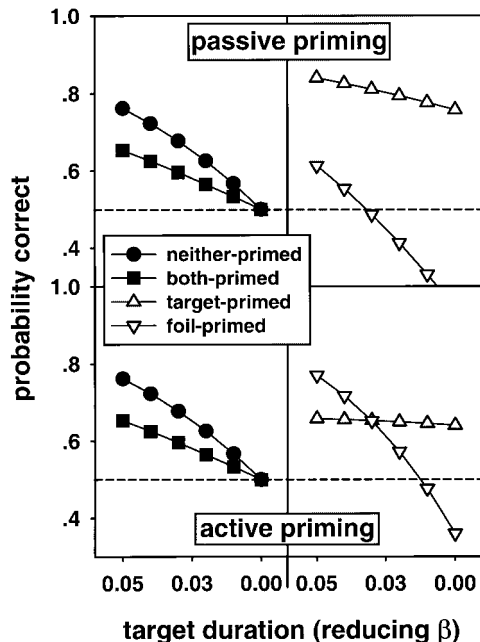


Figure 9. The responding optimally with unknown sources of evidence model predictions for decreasing target-flash duration (reducing β) with repetition priming ($\rho = 1$). The horizontal dashed line indicates chance performance. The default parameters listed for Figure 7 were used. The estimate of β was set to .05 and kept at that value as β decreased.

that are shown in the upper right-hand panel of Figure 9. In the case of passive priming, the effect of lowering β is to increase the difference between the target-primed and foil-primed conditions. In sum, reducing target duration was predicted to result in a preference reversal with active priming and a magnification of preference with passive priming.

Method

Participants. One hundred Indiana University Bloomington undergraduates participated in the experiment, receiving introductory psychology course credit for their participation. Fifty-one of the participants received the passive-priming version of the experiment, and 49 received the active-priming version. All participants were native English speakers with normal or corrected-to-normal vision.

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. The stimuli were displayed as black against a gray background to avoid phosphor persistence. Subject booths were enclosed, and the lighting was dim to avoid eyestrain. The resulting visual contrast was close to 100%. Monitor distance and font size were chosen such that prime words, choice words, and the center target word encompassed slightly less than 3° of visual angle horizontally (and even less vertically). All responses were collected through response boxes with four keys.

Materials. The stimuli for all presentations were drawn from a pool of 1,000 five-letter words. These words had a minimum written-language frequency of four, as defined and measured by Kučera and Francis (1967). Randomly generated letterlike pattern masks were used to avoid pattern-mask habituation (see Figure 1 for examples of the pattern masks). All words were displayed in capitalized Times Roman 22-point font.

Procedure. All variables, except passive versus active priming, were within subject. Repetition priming was the only type of priming used in this experiment. The basic design used the following two variables: priming condition, with four levels (neither primed, both primed, target primed, and foil primed) and flash duration, with three levels (full, half, and zero). Neither-primed trials were created by randomly selecting two prime words, a target word and a foil word, from the pool of 1,000 five-letter words. Word selection occurred without replacement such that a given word appeared only once within the experiment, thus avoiding contamination from long-term repetition priming. In the both-primed conditions only two words were selected because the primes were repeated as the choice words. For the target-primed and foil-primed conditions, the target or foil was accordingly repeated and the other choice word randomly selected (these conditions required three words).

As seen in Figure 1, two prime words were presented on every trial: one above and one below the fixation point. During the two-alternative, forced-choice (2-AFC) procedure, the target and foil were presented on the left and right of fixation. The top-down position of the primes and the left-right position of the choice words were both fully counterbalanced. On every trial, a sequence of events occurred as shown in Figure 1. Prior to the first display of the prime words (i.e., the first screen in Figure 1), a fixation point was displayed for 250 ms followed by a blank screen for 250 ms.

The full condition was analogous to standard repetition priming as used in the experiments of Huber et al. (2001) and Huber et al. (in press). The duration of the target flash was set individually for each participant as determined by a block of trials during which flash time was progressively reduced until 75% performance was obtained. The half condition was identical except that the flash duration was set to half the adjusted full duration. The flash duration in the half condition was determined by dividing the full duration by two and rounding down to the next lowest number of screen refresh cycles. In the zero condition, no target was flashed, although all other presentations and masks were the same as in the full and half conditions.

For passive priming, the experiment began with 16 trials of perceptual-identification practice using 100-ms target flash presentations, and priming was always neither primed. Participants in the passive-priming version were instructed that prime words were a warning to prepare for the flash of the target word. Participants in the active-priming version received a block of 16 active-priming-task practice trials, in isolation, prior to the 16 practice trials that included perceptual identification. The active-priming task was a self-paced judgment of whether the two prime words could possibly serve the same common part of speech (i.e., noun, verb, adjective, other). On average, participants took 2,412 ms giving this response. Following their response, the sequence of events unfolded just as with passive priming, as shown in Figure 1. Participants were given feedback for the part-of-speech matching task during the initial block of practice trials but not during the rest of the experiment (the instructions throughout the experiment stressed the importance of continued performance in the part-of-speech matching task). On average, correct performance in the part-of-speech matching task was 63% (with a single word potentially serving many parts of speech, this is a difficult task).

Following practice, participants received a block of 64 neither-primed trials. The purpose of this block was to find the target-flash duration at which performance was 75% (i.e., the target-flash duration for the full condition). Appropriate durations averaged 61.9 ms for the passive-priming participants and 58.8 ms for the active-priming participants, although there were large individual differences, with times ranging from 41.7 ms to 141.7 ms. On average, target-flash duration for the half condition was 28.1 ms for the passive-priming participants and 26.0 ms for the active-priming participants. A staircase method was used to find the appropriate target-flash duration for the full condition. Participants were fully informed about the procedure. Following these 64 trials, two blocks of 96 experimental trials were presented. The entire experiment took approximately 40 min.

Feedback for the perceptual-identification 2-AFC procedure was given on every trial throughout the entire experiment. In the case of the zero condition there could be no objectively correct answer because no word was flashed; instead, a nominally correct answer was assigned and used to provide feedback. Although the target- and foil-primed conditions could have been combined for the zero-flash duration (i.e., both scored as the probability of choosing the repeated choice word), we kept these two nominal conditions separate in the scoring such that the figure would be symmetric, and the power associated with every condition would be identical.

Results

Average probability correct values for the various conditions and passive versus active priming appear graphically in Figure 10 and numerically in the Appendix. The error bars show two standard errors of the mean. For both passive and active priming there was a main effect of prime condition, passive, $F(3, 150) = 23.20$, $MSE = 0.77$, $p < .001$; active, $F(3, 144) = 2.70$, $MSE = 0.089$, $p < .05$; and a main effect of target-flash duration, passive, $F(2, 100) = 42.31$, $MSE = 1.26$, $p < .001$; active, $F(2, 96) = 63.97$, $MSE = 2.48$, $p < .001$; and a significant Priming Condition \times Target-Flash Duration interaction, passive, $F(6, 300) = 5.01$, $p < .001$; active, $F(6, 288) = 8.44$, $MSE = 0.12$, $p < .001$.

An ANOVA was run breaking down the four priming conditions into the separate effects of priming the target (target- and both-primed vs. neither- and foil-primed) and priming the foil (foil- and both-primed vs. neither- and target-primed). This analysis examines the effect of priming the foil in isolation, which can be taken

as an indication of whether priming has an effect on the decision process independent of the target flash (i.e., a preferential effect). One should see Huber et al. (2001) for more discussion of this and related analyses. Should this analysis suggest there was a preferential effect, the direction of preference was determined by comparing the target-primed and foil-primed conditions (assuming that perceptual benefits do not contribute greatly to target-primed performance). Finally, unbiased performance deficits with priming were indexed by comparing the both-primed condition with the neither-primed condition. Such a result is interpreted by the ROUSE model in terms of increased variability caused by priming that harms performance in general.

With passive priming there was a significant effect of priming the foil (i.e., a preference) for the full condition, $F(1, 50) = 12.86$, $MSE = 0.30$, $p < .0025$; the half condition, $F(1, 50) = 21.21$, $MSE = 0.51$, $p < .001$; and the zero condition, $F(1, 50) = 16.17$, $MSE = 0.40$, $p < .001$. In each case, comparing the target-primed and foil-primed conditions, this preference was to choose the repeated choice word, full, $t(50) = 2.55$, $SE = .036$, $p < .001$; half, $t(50) = 5.80$, $SE = .038$, $p < .001$; and zero, $t(50) = 6.71$, $SE = .031$, $p < .001$. In addition to these direction of preference effects, there was a general performance deficit with priming for the full condition, as revealed by comparing the both-primed condition with the neither-primed condition, $t(50) = 2.34$, $SE = .026$, $p < .025$. There was no difference between the neither- and both-primed conditions for the half condition, $t(50) = .88$, $SE = .022$, $p = .39$, and for the zero condition there was no correct target, and therefore, performance in the neither- and both-primed conditions was necessarily at chance.

With active priming in the full condition, preference played a role, as demonstrated by the main effect of priming the foil, $F(1, 48) = 5.05$, $MSE = 0.082$, $p < .05$. A comparison of the target- and foil-primed conditions revealed this as a preference against choosing repeated words, $t(48) = 2.55$, $SE = .035$, $p < .01$. In addition there was a both-primed deficit for the full condition, $t(48) = 7.32$, $SE = .024$, $p < .001$. For the half and zero conditions, with active priming, there was no effect of priming the foil, half, $F(1, 48) = 0.14$, $MSE = 0.0039$, $p = .71$; zero, $F(1, 48) = 0.69$, $MSE = 0.018$, $p = .41$; no difference between target-primed and foil-primed conditions, half, $t(48) = .44$, $SE = .038$, $p = .67$; and zero, $t(48) = .87$, $SE = .041$, $p = .39$; and no both-primed deficits, half, $t(48) = .05$, $SE = .026$, $p = .96$.

Discussion

The full target-duration condition replicates the results of Huber et al. (2001) and Huber et al. (in press), demonstrating a preference for repeated words with passive priming, a slight preference against repeated words with active priming, and both-primed deficits for both passive and active priming. In addition, this experiment confirms the ROUSE prediction that shorter target flash duration will result in diminished discounting efficacy. For passive priming, this is revealed by the increase in the target-primed greater than foil-primed difference for the half and zero target-duration conditions compared with the full target-duration condition, $t(50) = 4.66$, $SE = .026$, $p < .001$. For active priming this is revealed by the decrease in the foil-primed greater than target-primed difference for the half and zero target-duration conditions compared with the full target-duration conditions, $t(48) = 3.53$,

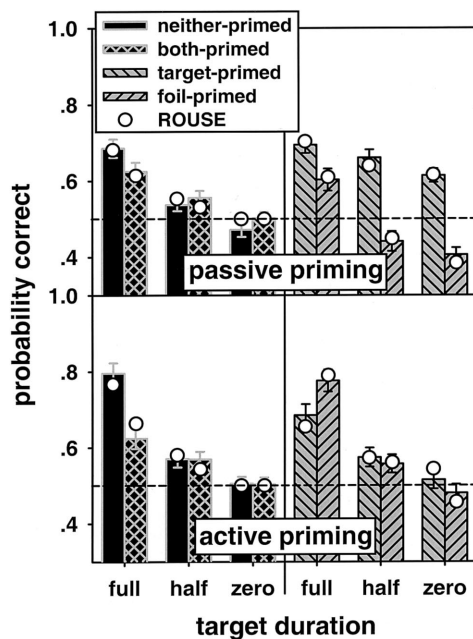


Figure 10. The results and best-fitting responding optimally with unknown sources of evidence model (ROUSE) predictions for Experiment 1, which tested repetition priming at three different target-flash durations. The horizontal dashed line indicates chance performance. Passive versus active priming was a between-subjects manipulation. The white circles show the results of the best-fitting ROUSE parameters found in Table 1. The error bars represent two standard errors of the mean.

Table 1
Best Fitting ROUSE Parameters

Parameter	Experiment 1		Experiment 2	
	Passive	Active	Passive	Active
α	.127	.162	.114	.462
α'	.094	.171	.086	.656
β				
Full duration	.045	.074	.024	.027
Half duration	.012	.018		
γ		.062		.030
ρ				
4 of 5			.745	
3 of 5			.654	
2 of 5			.012	
$\Sigma\chi^2$ (error)		30.314	18.424	

Note. α represents actual prime; α' represents estimated prime; β represents target; γ represents noise; ρ represents prime similarity. ROUSE = responding optimally with unknown sources of evidence.

$SE = .033$, $p < .0025$. As seen in Figure 9, the default-parameter ROUSE prediction was for the preference against repeated words with active priming to fully crossover to a preference in favor of repeated words with shorter target durations. Although this full preference crossover was not established statistically, the data clearly moved in the direction of a preference for repeated words, which conforms to the qualitative ROUSE prediction of diminished discounting efficacy. In the next section we demonstrate that ROUSE can quantitatively capture these results.

Besides confirming specific ROUSE predictions, these results provide strong evidence against a strategic interpretation of preference effects in the 2-AFC paradigm. For instance, suppose that participants in the passive-priming version adopted an explicit strategy to choose whichever word repeated a prime, given insufficient information from the target flash. For shorter target durations there is less perceptual information, and participants would rely more heavily on this strategy. This conforms to the observed increase in the target-primed greater than foil-primed difference for the half and zero conditions and supports such an explicit strategy account. However, the active-priming results are opposite of that expected by an explicit-strategy account. If, with the full target duration and active priming, participants adopted a strategy not to choose repeated words, then this strategy would result in even larger foil-primed greater than target-primed differences for the half and zero durations. Instead, the opposite occurred, as predicted by ROUSE.

Following disclosure of these data, Ratcliff and McKoon (2001) developed a mechanism for handling these results by assuming source confusion linearly increases with decreasing target perception (i.e., decreasing target duration). The multinomial model does not specify the source-confusion process, but the authors state "confusion should be less likely as the amount of perceptual information increases" (Ratcliff & McKoon, 2001, p. 841). In contrast, ROUSE provides the mechanism of source confusion, locating it in prime activation of shared features. As such, source confusion is unaffected by target-flash duration manipulations in ROUSE. Instead, discounting efficacy decreases with decreasing target-flash duration. This occurs not because of additional assumptions, but rather as the byproduct of probabilistic feature

activation, which is the essence of the ROUSE model. Details of applying the two models to the data are taken up in the next section.

The results of this experiment have potential implications that extend beyond the local context of this paradigm. ROUSE predicted that the extent of discounting (i.e., the relation of α' to α) plays a minor role for very short target durations (examine the $\beta = 0$ portion of the right-hand graphs of Figure 9) but is crucial to the direction and size of priming for longer target durations (examine the $\beta = .05$ portion of the right-hand graphs of Figure 9). By extension, these results (and the theory) suggest priming paradigms using even longer target durations may be more sensitive to discounting effects. For instance, in lexical decision and naming tasks, targets remain until a response is given and therefore the extent of discounting may strongly determine the magnitude and direction of priming. However, these extrapolations require additional theoretical work and testing in future studies.

Data Fitting With ROUSE and the Multinomial Model

ROUSE was fit to the data from Experiment 1 using an analytic simulation method (Huber, 2002). As explained in Huber et al. (2001), the log-likelihood method for calculating chi-square was used as a measure of fitting error.³ The simplex algorithm was used to adjust the parameters. The best-fitting parameters appear in Table 1, and the results with these parameters appear graphically in Figure 10, as the white circles, and numerically in the Appendix.

Because passive versus active priming was a between-subjects manipulation, different levels of α , α' , and β were allowed for the two groups. Given only three target durations, one of which must correspond to $\beta = 0$, we could not sensibly estimate the functional relation between target-flash duration and β . We therefore estimated separate β s for the full and half target durations. The

³ There are two errors in the Huber et al. (2001) article in specifying the log-likelihood method of calculating chi-square. Specifically, Equations 3 and 4 were written with the error-rate term subtracted from the correct rate term. Instead, the error-rate term should be added to the correct-rate term.

estimate of β (i.e., β') was set equal to the average of the three β s (i.e., $\beta' = [\beta_{\text{full}} + \beta_{\text{half}}]/3$). This was not a critical choice, however, because inaccurate estimates of β do not greatly affect ROUSE's predictions. The same value of γ was used for both the passive and active group of participants (γ is more likely related to the visual characteristics of the display and masks than to individual differences), and it was assumed that the estimate of γ was accurate ($\gamma' = \gamma$).

Unlike the data fitting appearing in Huber et al. (2001) and Huber et al. (in press), noise activation, γ , was a free parameter. In fitting this experiment (and Experiment 2) γ was allowed to vary, because the experiments reported in this article are particularly concerned with discounting efficacy. More noise confusion causes the crossover point appearing in the lower right-hand panel of Figure 9 to shift or even disappear. This occurs because the erroneous discounting of noise-activated features will begin to have a noticeable effect if there are a substantial number of noise-activated features. When this noise factor is largely missing (because of a small γ), erroneous discounting of target-activated features is the only method for maintaining discounting efficacy, and therefore, the effects caused by reducing the values of the terms in Equation 1 are obtained; these same effects are relatively diminished when the noise factor plays a large role. We use this argument to explain the failure in Experiment 1 to observe a large preference crossover with active priming as target duration decreased. That is, the best-fitting value of γ , .062, although small, was nevertheless sufficiently larger than the default value of .02 to bring the predictions in line with the data. If the default value was used, the predicted interaction would be too large for a good quantitative fit.

Next we fit the results of Experiment 1 with the multinomial model. We replicated unpublished fits of this data performed by Ratcliff and McKoon (personal communication, March 24, 2000). To handle the increased preference to choose repeated words as target-flash duration decreased, Ratcliff and McKoon (2001) assumed that the source confusion parameter (S) linearly increases as perception (P) decreases, up to some maximum value (S_{max}) when P equals zero (i.e., $S = S_{\text{max}}[1 - P]$). They adopted this same

assumption in fitting their replication of the passive-priming version of Experiment 1 (Ratcliff and McKoon, 2001). In fitting these data with the multinomial model, we allowed separate perception parameters for the passive and active groups at each target-flash duration (in the zero condition, P is set to zero). As assumed by Ratcliff and McKoon (2001), we allowed a free parameter for discounting (D) for the active group and set this parameter to zero for the passive group. Different S_{max} values were allowed for the active and passive groups.

The multinomial-model results, with best-fitting parameters, are seen in the Appendix, and the best-fitting parameter values are found in Table 2. A comparison of Tables 1 and 2 reveals that the multinomial model fit the data better than ROUSE despite using two fewer parameters, a fact that should not be overlooked and speaks in favor of the multinomial model. However, the fits were obtained by adding specific assumptions to the model that were designed to produce the correct pattern. In this sense, the model is descriptive rather than prescriptive. In contrast, ROUSE contains no parameters that specifically mandate decreased discounting as target duration decreases. Instead, the decreased discounting efficacy in ROUSE for shorter target durations is a property of the model in its original form.

Experiment 2: Reducing ρ by Reducing Orthographic Prime Similarity

The third factor in Equation 1 that affects discounting efficacy, and hence determines the presence or absence of a preference reversal in the active condition, is the similarity of prime(s) to choice(s), designated by ρ . Huber et al. (2001) varied ρ in several ways, by comparing repetition with orthographic priming in Experiment 2, by comparing repetition with associative priming in Experiment 1, and by comparing orthographic with associative priming in Experiment 4. However, each of those manipulations involved prime- to choice-word relations of different feature types: word identity, orthographic overlap, and associative relations. If repetition priming reflects the highest degree of similarity, orthographic priming (sharing all but one letter) reflects an intermediate

Table 2
Best Fitting Multinomial Model Parameters

Parameter	Experiment 1		Experiment 2	
	Passive	Active	Passive	Active
P				
Full duration	.354	.595	.270	.298
Half duration	.110	.182		
D (repetition priming)		.180		
D				
4 of 5				.100
3 of 5				.068
2 of 5				.015
S_{max}	.209	.271		
S				
4 of 5			.094	
3 of 5			.124	
2 of 5			.029	
$\Sigma\chi^2$ (error)	23.154		15.050	

Note. P = perception; D = discounting; S_{max} = repetition prime with $P = 0$; S = source confusion.

level of similarity, and associative priming reflects the lowest degree of similarity, then these experiments lend support to the prediction that discounting becomes less effective as prime similarity is reduced. However, a test within one feature type is needed. In Experiment 2 we manipulated similarity within the orthographic-phonemic feature type. This was accomplished by using primes that share two, three, or four of five letters in their respective letter positions.

Using the default ROUSE parameters, Figure 11 shows the predicted results as a function of decreasing prime similarity (decreasing p). As with the other predicted graphs, the most important result appears for the target- and foil-primed conditions with active priming, shown in the lower right-hand panel. With repetition priming ($p = 1$) there is a preference against choosing repeated words because of excessive discounting ($\alpha' > \alpha$). As prime similarity decreases, the target- and foil-primed conditions cross over, because of decreased discounting efficacy, and there is a preference to choose primed words. Finally, as prime similarity approaches zero, all prime effects necessarily disappear and all the conditions become equal, both for passive and active priming. Given that our experimental manipulation of orthographic-phonemic similarity does not include repetition priming, we did not necessarily expect to observe the full crossover depicted in Figure 11. Regardless, in the active-priming condition we expected to observe an increase in preference for primed words as prime similarity decreases.

Method

Participants. One hundred six University of Colorado at Boulder undergraduates participated in the experiment, receiving introductory psychology course credit for their participation. Fifty-three participants re-

ceived the passive-priming version of the experiment, and 53 received the active-priming version. All participants were native English speakers with normal or corrected-to-normal vision.

Equipment. Unlike Experiment 1, Experiment 2 used PC monitors running at 70 Hz yielding 14.3 ms per screen refresh. Responses were collected through the computer keyboard, and participants were tested four at a time in a well-lit room with open-air dividers between each computer.

Materials. The stimuli for all presentations were drawn from a pool of 240 five-letter word quadruples. Each quadruple consisted of a choice word (e.g., *drown*) and three potential prime words that shared four (e.g., *drawn*), three (e.g., *known*), or two (e.g., *trout*) of the same letters in their respective letter positions. An additional 352 five-letter words were used during the practice trials and the target-duration threshold-determination block of trials.

Procedure. Except as noted, all procedures were the same as those used in Experiment 1. All variables, except passive versus active priming, were within subject. The basic design used the following two variables: priming condition, with four levels (neither primed, both primed, target primed, and foil primed) and orthographic-phonemic prime similarity, with three levels (four, three, or two letters kept the same, and in the same letter position, out of five). The top-down position of the primes and the left-right position of the choice words were randomly counterbalanced across trials.

Participants in the active-priming version received a block of 16 active-priming-task practice trials, in isolation, prior to the 16 practice trials that included perceptual identification. The active-priming task was to determine the number of potential verbs among the two prime words (i.e., zero, one, or two were potential responses). On average, participants took 3,534 ms giving this response. At the end of each trial, throughout the entire experiment, participants were given feedback on the actual number of potential verbs, their guessed number of potential verbs, and whether they selected the correct choice word in the perceptual-identification task. On average, correct performance in the number of verbs task was 67%.

As set by a threshold-determination block of trials, appropriate target-flash durations averaged 23.5 ms for the passive-priming participants and 20.5 ms for the active-priming participants, although there were large individual differences, with times ranging from 14.3 to 185.8 ms. Following the threshold determination block of trials, the experiment consisted of two blocks of 60 trials in which the 12 experimental conditions were repeated five times in each block in random order.

Results

Average probability correct values for the various conditions and passive versus active priming appear graphically in Figure 12 and numerically in the Appendix. For passive priming, there was a main effect of priming condition, $F(3, 156) = 4.30$, $MSE = 0.12$, $p < .01$; but the main effect of orthographic-phonemic prime similarity, $F(2, 104) = 2.40$, $MSE = 0.038$, $p = .096$, fell just shy of significance. There was a significant Priming Condition \times Orthographic-Phonemic Prime Similarity interaction, $F(6, 312) = 2.51$, $MSE = 0.043$, $p < .025$. For active priming there was a main effect of priming condition, $F(3, 156) = 5.15$, $MSE = 0.13$, $p < .0025$; but not orthographic-phonemic prime similarity, $F(2, 104) = 1.77$, $MSE = 0.029$, $p = .18$; and no Priming Condition \times Orthographic-Phonemic Prime Similarity interaction, $F(6, 312) = 0.79$, $MSE = 0.014$, $p = .58$. However, these analyses are relatively uninformative because we predicted a nonlinear interaction. Specific post hoc contrasts were accordingly carried out.

With passive priming there was a significant effect of priming the foil (i.e., a preference for four-letter prime similarity, $F(1, 52) = 5.84$, $MSE = 0.18$, $p < .025$; and three-letter prime similarity, $F(1, 52) = 7.37$, $MSE = 0.15$, $p < .01$; but not with

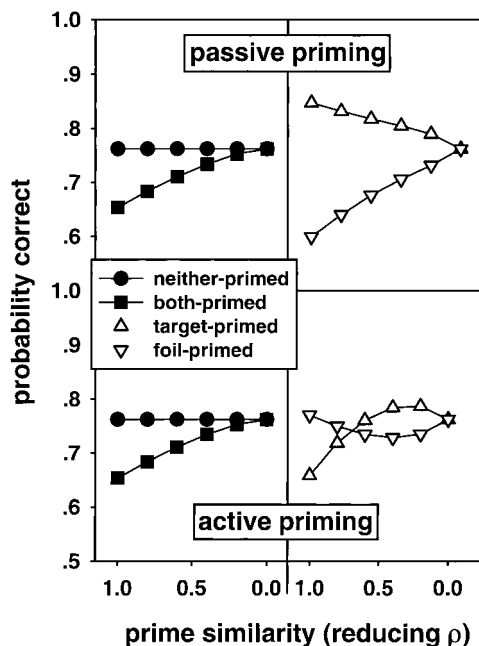


Figure 11. The responding optimally with unknown sources of evidence model predictions for decreasing prime similarity (reducing p). The default parameters listed for Figure 7 were used.

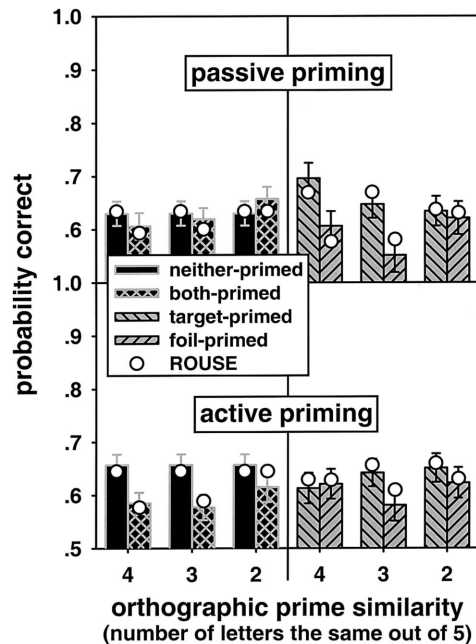


Figure 12. The results and best-fitting responding optimally with unknown sources of evidence model (ROUSE) predictions for Experiment 2, which tested three different levels of orthographic–phonemic prime similarity by maintaining four, three, or two out of five letters the same, and in the same letter position, between primes and related choice words. Passive versus active priming was a between-subjects manipulation. The white circles show the results of the best-fitting ROUSE parameters found in Table 1. The error bars represent two standard errors of the mean.

two-letter prime similarity, $F(1, 52) = 0.14$, $MSE = 0.0023$, $p = .712$. In fact, for two-letter prime similarity, there was no effect of priming condition, $F(3, 156) = 0.70$, $p = .556$, and no further tests were warranted for that level of prime similarity. For four- and three-letter prime similarity, comparing the target-primed and foil-primed conditions, the direction of the preference was to choose the primed choice word, four letter, $t(52) = 2.24$, $SE = .040$, $p > .025$; three letter, $t(52) = 3.01$, $SE = .032$, $p < .0025$. Despite these preference effects, there were no both-primed deficits in either case, four letter, $t(52) = 1.06$, $SE = .023$, $p = .296$; three letter, $t(52) = .45$, $SE = .025$, $p = .656$.

With active priming and four-letter prime similarity, there was no effect of priming the foil, $F(1, 52) = 2.02$, $MSE = 0.055$, $p = .161$, and correspondingly, no difference between the target-primed and foil-primed conditions, $t(52) = .21$, $SE = .036$, $p = .834$. However, there was a both-primed deficit, $t(52) = 3.03$, $SE = .024$, $p < .0025$. We interpret this pattern of results to be consistent with the predictions of ROUSE (as verified by the quantitative predictions shown in Figure 12): Preferential variability (i.e., source-confusion variability) caused a both-primed deficit, whereas on average the preference was neutral in its direction. With three-letter prime similarity, there was an effect of priming the foil, $F(1, 52) = 11.63$, $MSE = 0.26$, $p < .0025$, which was revealed as a preference to choose the primed choice word by a comparison of target-primed to foil-primed, $t(52) = 2.10$, $SE = .029$, $p < .025$. In addition, there was a both-primed deficit, $t(52) = 3.19$, $SE = .025$, $p < .0025$. With two-letter prime

similarity, there was no effect of priming the foil, $F(1, 52) = 2.49$, $MSE = 0.065$, $p = .120$, and more generally, no effect of priming condition, $F(3, 156) = 1.21$, $MSE = 0.022$, $p = .309$.

Discussion

In Experiment 2, for four out of five letters shared between prime and choice word, target-primed performance was better than foil-primed performance with passive priming, but there was no difference between these conditions with active priming. Furthermore, with passive priming, there was no both-primed deficit, but with active priming there was a significant both-primed deficit. These four out of five letter-priming results fully replicate Experiment 2 of Huber et al. (2001), which used primes that shared four out of five, or three out of four letters in the same letter positions.

The less interesting ROUSE prediction concerned passive priming. ROUSE predicted that as prime similarity decreased, the target-primed greater than foil-primed preference would gradually decrease until these conditions become equal in the absence of priming. This pattern was found in Experiment 2 (statistically, a decrease in the target-primed minus foil-primed difference for the four-letter and three-letter levels of prime similarity compared with the two-letter level of prime similarity, $t[52] = 2.33$, $SE = .034$, $p < .025$). However, this result would be the prediction of any theory monotonically relating the preference for primed words to the level of prime similarity.

The critical prediction occurred in the active condition. Comparing the active-priming results with the predicted graphs, it appears that the equal target-primed and foil-primed performance observed for the four-letter level of prime similarity corresponds to the crossover point in the lower right-hand panel of Figure 11. Moving from that crossover point to even lower levels of prime similarity, ROUSE predicted that the preference should switch to a preference in favor of primed words and then converge to no preference. The observed active-priming data in Figure 12 show just this pattern. Because we expected this nonlinear pattern, and because there appears to be a very small amount of priming remaining with two-letter prime similarity, we tested the prediction statistically with contrast weights of -3 , 2 , and 1 on the target-primed minus foil-primed difference for the four-letter, three-letter, and two-letter levels of prime similarity (i.e., the prediction for the difference between target primed and foil-primed is no difference, a large positive difference, and a small positive difference). This contrast fell just shy of significance, $t(52) = 1.48$, $SE = .12$, $p = .073$. Nevertheless, as shown in the results section, and as might be expected with this nonlinear pattern, there was no significant difference between the target-primed and foil-primed conditions for the four-letter and two-letter levels of prime similarity, but for the intermediate three-letter level of prime similarity, target-primed performance was significantly better than foil-primed performance.

Data Fitting With ROUSE and the Multinomial Model

ROUSE was fit to the data from Experiment 2 using the same method as in Experiment 1. The best fitting parameters appear in Table 1, and the results with these parameters appear graphically in Figure 12, as the white circles, and numerically in the Appendix.

As in Experiment 1, different levels of α , α' , and β were allowed for the passive and active groups of participants. It was assumed that the estimate of β was accurate ($\beta' = \beta$). The same value of γ was used for both the passive and active group of participants, and it was assumed that the estimate of γ was accurate ($\gamma' = \gamma$). As explained in the data-fitting section following Experiment 1, γ was a free parameter, although the same value was applied across both passive and active priming. The values of α and α' followed the pattern of all previous fits with ROUSE: $\alpha' < \alpha$ with passive priming (i.e., too little discounting) and $\alpha' > \alpha$ with active priming (i.e., too much discounting).

Three different similarity parameters (ρ) were used, one for each of the levels of orthographic-prime similarity, and these same parameters were used for both passive and active priming (i.e., we assumed similarity is largely a function of the stimuli and not subject to large differences across groups of participants). These parameters were allowed to range freely and produced the sensible ordering appearing in Table 1 with four-, three-, and two-letter similarity set to .745, .654, and .012, respectively. These values also make sense in terms of the predicted graph appearing in Figure 11. In that graph, the crossover point occurs around .7 (i.e., target primed equals foil primed with active priming), and with four-letter similarity no difference was observed between the target-primed and foil-primed conditions.

To compare the two models, we fit Ratcliff and McKoon's (2001) multinomial model to these data. As with Experiment 1, we adopted their assumption of no discounting with passive priming. Because the multinomial model does not contain a mechanism with which to manipulate similarity, it cannot be specified in advance how source confusion and discounting should separately vary as a function of prime similarity. Thus we allowed both the source-confusion and discounting parameters to range freely as a function of prime similarity. Allowing separate source-confusion parameters for passive and active priming produced an overfitting of the active-priming results because both source confusion and discounting contribute to deficits with priming, and these parameters play off each other. Our solution was to equate the source-confusion parameters across passive and active priming (similar to fixing the prime similarity parameters with ROUSE across passive and active priming). The best fitting parameters for the multinomial model appear in Table 2, and the results with these parameters appear in the Appendix. As with Experiment 1, the multinomial model fit the data with less fitting error than ROUSE despite using fewer parameters, again an impressive demonstration of the ability of the multinomial model to describe the patterns of data.

Considering that phonemic similarity may play a key role in orthographic priming, the ROUSE similarity parameters seem sensible: With four (*drawn-drown*) and three (*known-drown*) letters the same, out of five, there tend to be syllables in common between primes and primed choice words, but with only two (*trout-drown*) letters the same, the orthographic-phonemic relationship is less clear and perhaps not significantly different than the random pairings found in unprimed conditions (unprimed pairs were allowed to contain letters in common). ROUSE captured this intuition by estimating high levels of similarity for the four- and three-letter conditions (.745 and .654) but negligible similarity for the two-letter condition (.012). In contrast, the multinomial model handled the results by estimating that source confusion should be highest for the three-letter condition (.124) and lower in the four-

and two-letter conditions (.094 and .029). It is difficult to imagine why prime similarity should result in such a nonlinear pattern of source confusion. It is for reasons such as this that we feel that the multinomial model provides a description of the data, rather than capturing the underlying processes.

General Discussion

Huber et al. (2001) used 2-AFC testing of short-term word priming and observed a preference to choose repeated words with passive priming and a preference against choosing repeated words with active priming. These new experiments provide further replications of those results and identify additional variables that are important for such preference changes. In addition, the new experiments provide a strong case against a strategic bias interpretation of preference changes. For instance, if preference results from a strategy against choosing repeated words, it is hard to understand (for active priming) why this strategy would reverse because of increasing orthographic similarity of the choice words, decreasing target-flash duration, or decreasing similarity of primes to choices. It could be argued that changes in similarity (of either type) make it difficult to determine which of the choice words was primed, but this would only result in a lack of preference and not the observed switch from a preference against to a preference in favor of choosing primed words. Such an argument would not apply to changes of target-flash duration, and it is hard to see how any strategic argument could explain the observed result. Therefore, preference changes appear to be an interesting and complex behavior that may further elucidate the identification process.

In this article we have attempted to highlight the subtle differences between Huber et al.'s (2001) ROUSE model and Ratcliff and McKoon's (2001) multinomial model. However, one should not lose sight of the considerable success of both models, a success we attribute to the use in both models of the mechanisms of source confusion and discounting. The ROUSE model proposed source confusion between primes and the target flash to explain the passive-priming pattern of results and the discounting of primed words to explain the switch with active priming. The multinomial model followed up on this work, borrowing the same two mechanisms but implemented as all-or-none decision branches rather than as part of feature evaluation. That different model implementations of these mechanisms are equally successful provides converging evidence that source confusion and discounting are important aspects of word identification in particular and, perhaps, important aspects of perceptual processing in general.

Nevertheless, comparison of the models is important, and such comparisons are the key to motivating informative experiments. In this case, it was our attempt to falsify the ROUSE model by testing several nonintuitive predictions that led to the reported experiments. Specifically, ROUSE predicted that increasing the similarity of the choice words to each other, decreasing the target-flash duration, and decreasing the similarity of the primes to the choice words would each result in a reduction of discounting efficacy that would be observed as a change in preference direction. Each of these predictions was empirically confirmed. In each case, ROUSE quantitatively captured the data by varying parameters related to the specific experimental manipulations (i.e., a similarity parameter or the target flash activation parameter), without changing either the source confusion (α) or discounting (α') parameters.

Comparisons of ROUSE to the multinomial model can be made on several dimensions. Both models provide good quantitative accounts of the data, although the multinomial actually fit slightly better than ROUSE. However, we prefer the ROUSE model because, with variables that are more closely tied to experimental manipulations, it is more tightly constrained. This results in a priori predictions that if violated empirically would require significant changes in the model. Because the multinomial model does not precisely specify how source confusion (S) and discounting (D) relate to similarity manipulations and other physical characteristics of the display situation, the multinomial model makes relatively few a priori predictions. The multinomial model is capable of extrapolation once it has been adjusted to predict a set of data; it can be used to “predict” the results of those same manipulations in other experiments. For new manipulations, however, the model must usually be adapted once the pattern of data is known.

Although predictions are easy to derive mathematically with the multinomial model, the parameters are not clearly mapped to the kinds of experimental manipulations that we often use to test aspects of perception and priming. Therefore, new relationships between the model variables and experimental manipulations were generated following each new manipulation. For example, to handle high choice-word similarity, Ratcliff and McKoon (2001) assumed that the wrong choice word might be discounted with some probability. To handle decreasing target duration, they assumed that source confusion increased, in a linear fashion, as target perception decreased. To handle decreasing prime similarity our fit of the multinomial model specified a nonlinear function relating prime similarity to source confusion.

It is for these reasons that we prefer ROUSE to the multinomial model. In recent years the search for appropriate criteria for selecting one model over another has received renewed attention (e.g., Myung, Forster, & Browne, 2000; Roberts & Pashler, 2000). We believe that one of these important criteria is the ability of a

model to make a priori testable and confirmable predictions. It is our observation that this is typically achievable only when there is a more or less direct mapping between the experimental manipulations and the parameters contained within the theory. The early stage of evolution for psychology as a science makes theories of this type fairly rare, but we believe the ROUSE model and these studies provide a good example of the way in which our field is advancing.

References

- Huber, D. E. (2002). *Computer simulations of the ROUSE model: An analytic method and a generally applicable technique for producing parameter confidence intervals*. Manuscript submitted for publication.
- Huber, D. E., & O'Reilly, R. C. (in press). Persistence and accommodation in short-term priming and other perceptual paradigms: Temporal segregation through synaptic depression. *Cognitive Science*.
- Huber, D. E., Shiffrin, R. M., Lyle, K. B., & Ruys, K. I. (2001). Perception and preference in short-term word priming. *Psychological Review*, 108, 149–182.
- Huber, D. E., Shiffrin, R. M., Quach, R., & Lyle, K. B. (in press). Mechanisms of source confusion and discounting in short-term priming 1: Effects of prime duration and prime recognition. *Memory & Cognition*.
- Kučera, H., & Francis, W. (1967). *Computational analysis of present day American English*. Providence, RI: Brown University Press.
- Myung, I. J., Forster, R. M., & Browne, W. M. (Eds.). (2000). Model selection [Special issue]. *Journal of Mathematical Psychology*, 44.
- Ratcliff R., & McKoon, G. (1997). A counter model for implicit priming in perceptual word identification. *Psychological Review*, 104, 319–343.
- Ratcliff, R., & McKoon, G. (2001). A multinomial model for short-term priming in word identification. *Psychological Review*, 108, 835–846.
- 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.
- Roberts, S., & Pashler, H. (2000). How persuasive is a good fit? A comment on theory testing. *Psychological Review*, 107, 358–367.

(Appendix follows)

Appendix

Observed and Predicted Two-Alternative, Forced-Choice Data

Target duration and priming condition	Passive priming					Active priming				
	Observed	ROUSE	Multinomial	SE	No. of observations	Observed	ROUSE	Multinomial	SE	No. of observations
Experiment 1										
Full										
Neither	.686	.681	.677	.024	816	.796	.767	.798	.027	784
Both	.625	.614	.633	.024	816	.624	.664	.658	.027	784
Target	.695	.703	.721	.022	816	.685	.655	.672	.029	784
Foil	.603	.609	.586	.029	816	.776	.790	.762	.029	784
Half										
Neither	.537	.553	.555	.016	816	.570	.580	.591	.022	784
Both	.556	.531	.536	.018	816	.569	.543	.537	.020	784
Target	.661	.640	.638	.021	816	.574	.573	.559	.025	784
Foil	.441	.449	.452	.026	816	.557	.561	.557	.024	784
Zero										
Neither	.472	.500	.500	.019	816	.505	.500	.500	.018	784
Both	.502	.500	.500	.016	816	.503	.500	.500	.018	784
Target	.614	.616	.604	.018	816	.515	.544	.521	.024	784
Foil	.406	.384	.395	.018	816	.480	.456	.479	.023	784
Experiment 2										
Prime similarity and priming condition										
Unprimed										
Neither	.630	.634	.635	.023	1590	.657	.645	.649	.020	1590
4 of 5										
Both	.606	.593	.611	.025	530	.585	.577	.599	.020	530
Target	.696	.669	.669	.029	530	.613	.629	.614	.029	530
Foil	.606	.577	.575	.028	530	.621	.628	.629	.028	530
3 of 5										
Both	.619	.600	.604	.021	530	.577	.589	.599	.024	530
Target	.647	.670	.680	.026	530	.642	.656	.645	.026	530
Foil	.551	.581	.556	.032	530	.581	.609	.598	.030	530
2 of 5										
Both	.657	.634	.627	.023	530	.615	.645	.636	.028	530
Target	.634	.637	.645	.028	530	.651	.660	.649	.027	530
Foil	.621	.631	.617	.031	530	.623	.630	.636	.029	530

Note. ROUSE = responding optimally with unknown sources of evidence.

Received June 26, 2001
Revision received April 24, 2002
Accepted April 24, 2002 ■