Available online at www.sciencedirect.com

ScienceDirect

Journal of Memory and Language 59 (2008) 495–523

Journal of Memory and Language

www.elsevier.com/locate/jml

Beyond mean response latency: Response time distributional analyses of semantic priming

David A. Balota [a,*], Melvin J. Yap [b], Michael J. Cortese [c], Jason M. Watson [d]

[a] Department of Psychology, Washington University, St. Louis, MO 63130, United States
[b] Department of Psychology, National University of Singapore, Block AS 4, #02-07, 9 Arts Link, Singapore 117570, Republic of Singapore
[c] Department of Psychology, University of Nebraska at Omaha, 6001 Dodge Street, Omaha, NE 68182, United States
[d] Department of Psychology, University of Utah, 380 South, 1530 East, Room 502, Salt Lake City, UT 84112, United States

Received 1 March 2007; revision received 3 October 2007
Available online 4 March 2008

Abstract

Chronometric studies of language and memory processing typically emphasize changes in mean response time (RT) performance across conditions. However, changes in mean performance (or the lack thereof) may reflect distinct patterns at the level of underlying RT distributions. In seven experiments, RT distributional analyses were used to better understand how distributions change across related and unrelated conditions in standard semantic priming paradigms. In contrast to most other lexical variables, semantic priming in standard conditions simply shifts the RT distribution, implicating a headstart mechanism. However, when targets are degraded, the priming effect increases across the RT distribution, a pattern more consistent with current computational models of semantic priming. Interestingly, priming effects also increase across the RT distribution when targets are degraded and primes are highly masked, supporting a memory retrieval account of priming under degraded conditions. Finally, strengths and limitations of alternative approaches for modeling RT distributions are discussed.

Keywords: Semantic priming; Masked priming; Distributional analysis; Lexical decision; Pronunciation

Breakthroughs in science often reflect improvements in the measurement tool investigators use to study a phenomenon. This can be most obviously seen in fields such as astronomy and biology, wherein the developments of higher magnification systems opened up new worlds for exploration. The recent advances in neuroimaging methods are another prime example of the power of measurement development. The present paper describes a step in this direction by increasing the magnification of the chronometric tools used to study psycholinguistic, and other response time (RT) dependent, phenomena.

Chronometric studies of language, memory, and attention have accumulated a vast amount of knowledge regarding the nature of representations, the processes engaged to tap such representations, and the time-course of the interactions between representations and processes. In order to better understand how one might increase the magnification of the standard chronometric approach, let us briefly consider the implicit assumptions researchers make.

In standard paradigms, researchers often manipulate a variable by including multiple observations (typically

* Corresponding author.
E-mail address: dbalota@artsci.wustl.edu (D.A. Balota).

0749-596X/$ - see front matter © 2008 Published by Elsevier Inc.
doi:10.1016/j.jml.2007.10.004

10–20) at each level of an independent variable (IV). A mean is then typically calculated for each level of an IV, and these means are submitted to inferential tests (most often analyses of variance) to estimate how reliable effects are across participants (and/or across items). Consider the classic semantic priming effect, which we will target in the present study. Here, the finding is that participants produce faster response latencies to a target, when the target word is related to a prime word (e.g., DOCTOR–NURSE), compared to when it is unrelated (e.g., FOREST–NURSE)[1]. The implicit assumption that researchers make is that the related and unrelated conditions produce symmetric RT distributions, and hence, the mean is a reasonably good estimate of the central tendency of these distributions. So, if one observes a 50 ms semantic priming effect, this indicates that the distribution of the related condition is shifted 50 ms away from the unrelated condition.

However, we all know that this implicit assumption is wrong. That is, RT distributions are rarely symmetrical around a mean, but are almost always positively skewed (see Luce, 1986, for a comprehensive review). Fig. 1 reflects an RT distribution from a single participant across approximately 2400 observations in lexical decision performance. Notice the strong positive skewing of the distribution. Hence, returning to the 50 ms semantic priming effect in the means, we are confronted with a number of first-order reasons why one might obtain such a difference: (a) The modal portion of the distribution may shift, without changing the tail; (b) The tail of the distribution may increase without changing the modal portion of the distribution; (c) Both the modal portion and tail may increase.

If researchers know that RT distributions are skewed, and that there are multiple ways in which an effect in means may be observed, then why does the field continue to use estimates of the mean to gain insights into the cognitive architecture? Clearly, there are many advantages in support of the mean. First, and probably most importantly, the mean is relatively easy to calculate and understand. Means are a fundamental summary statistic and dominate much of our common knowledge of the world (e.g., mean income, average miles per gallon, batting average, etc.). Second, the estimates are relatively stable. Why should one worry about the underlying distributions if the effects with means are replicable across studies? Third, and related to this, higher-order estimates of the RT distribution, such as skewness and kurtosis, are considerably less reliable (see Ratcliff, 1979). Why spend the additional

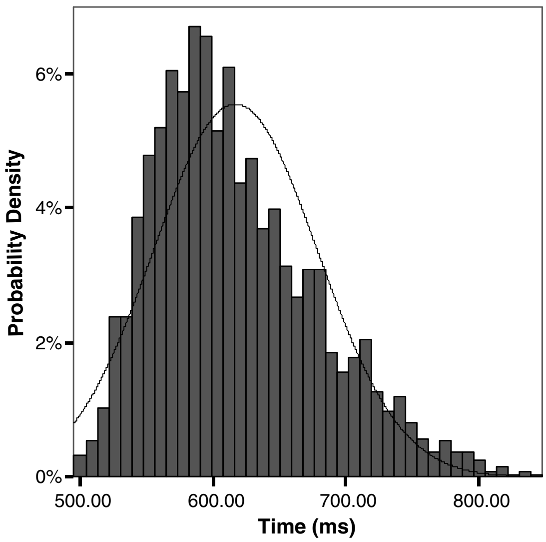

Fig. 1. Response time distribution for lexical decision performance across 2428 words taken from Balota et al. (2004).

effort to capture more subtle aspects of RT distributions if there is indeed a lack of stability in these estimates? In order to obtain stable estimates of higher order moments, one needs considerably more observations then the standard 10–20 observations per participant/cell. Does the added benefit justify the cost?

Although there are advantages to the mean, we, along with many others (e.g., Heathcote, Popiel, & Mewhort, 1991; Luce, 1986; Ratcliff, 1979; Rouder, Lu, Speckman, Sun, & Jiang, 2005; Van Zandt, 2002), believe that the *zeitgeist* is appropriate for researchers to move beyond the mean. The goal of the present paper is to provide a review of recent developments and extensions of RT distributional analyses to visual word recognition research. We should emphasize here that these arguments are not restricted to psycholinguistic variables, but indeed are relevant to all chronometric explorations of performance. However, in order to exemplify the power of this approach, we will focus on one of the most frequently studied effects in language and memory processing, i.e., the semantic priming effect.

Measuring aspects of the RT distribution: Beyond the mean

If it is time to move beyond the mean in estimating the influence of a variable or variables on RT distributions, how might one measure such influences? There are typically three major approaches that are used in the literature. First, one may have an explicit model that predicts how an underlying RT distribution may change as a function of a manipulation. Hence, one can simply fit the empirical data to the model's specific predictions

[1] Here we use the term "semantic" priming effect for simplicity; however, it should be noted that some, if not most, of the priming effects observed in these tasks may reflect associative relations, instead of semantic (see Hutchison, 2003, for a review).

regarding the RT distribution. An excellent example of this is the use of the diffusion model by Ratcliff and colleagues (see, for example, Ratcliff, Gomez, & McKoon, 2004). A second approach is to fit an empirical RT distribution to a theoretical function that captures important aspects of typical RT distributions. One can then make inferences from the estimated parameters of the theoretical function to determine the nature of an effect. This approach has been advocated by Luce (1986), among many others (e.g., Ratcliff, 1978; Rouder et al., 2005; Van Zandt, 2000, 2002) to better understand how variables influence RT distributions. Third, one may simply plot the data directly to determine if there are differential influences of a target variable on different portions of the RT distribution. For example, one may plot the mean of RTs across bins, called Vincentiles, or specific quantiles (e.g., 10%, 20%, 30%, etc.). Here, we will focus on the latter two approaches, but will have more to say about the first approach later in the paper.

Fitting an obtained RT distribution to an explicit mathematical function

There has been considerable work describing how best to capture empirical RT distributions (see Heathcote, Brown, & Mewhort, 2002; Luce, 1986; Rouder et al., 2005; Van Zandt, 2000, 2002). There are many functions available to fit RT functions, including the ex-Gaussian, ex-Wald, Weibull, Gamma, among many others. The advantages and disadvantages of the different approaches have been extensively reviewed by Van Zandt (2000). Although there may well be better functions available, for reasons described below, a number of researchers have used the ex-Gaussian function to capture aspects of RT functions. Indeed, it was Ratcliff's (1978, 1979) seminal work which demonstrated the stability of the ex-Gaussian estimates, and the power of this approach for testing specific predictions of models of memory retrieval. Here, we will attempt to demonstrate the utility of the ex-Gaussian approach for capturing visual word recognition performance.

The ex-Gaussian function conceptualizes RT distributions as the convolution of two underlying distributions: a Gaussian distribution and an exponential distribution. These are displayed in Fig. 2. The mean and standard deviation of the Gaussian component are captured by two parameters, μ and σ, while the exponential function is captured by a single parameter, τ, which reflects its mean and standard deviation. Importantly, ex-Gaussian analyses can be used as a descriptive model for capturing the influence of a variable on underlying RT distributions, with the parameters having a direct relation to the mean of a distribution. Specifically, the mean of an RT distribution is constrained so that it is the algebraic sum of the μ and τ estimates obtained by fitting that distribution. Hence, the ex-Gaussian function possesses an interesting

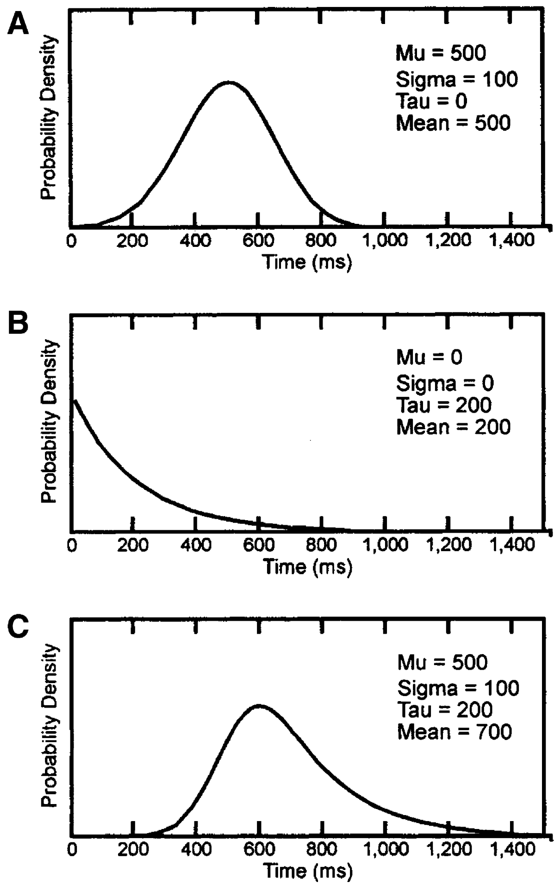

Fig. 2. Gaussian (A) and exponential distributions (B) and their convolution (C) for an ex-Gaussian distribution.

descriptive utility which provides an important connection to the extant mean-dominated literature.

Fig. 3 displays how a variable may influence the RT distribution and estimates of the ex-Gaussian parameters. For example, comparing Fig. 3A and B (taken from Balota & Spieler, 1999), a variable may primarily shift an RT distribution, which would be reflected in a change in the μ parameter. As noted earlier, this is the implicit assumption that researchers make. Alternatively, comparing Fig. 3A and C, a variable may have an isolated influence on the τ component, influencing the tail of the distribution. Finally, comparing Fig. 3A and D, one can see that a variable may actually have no effect on mean performance, but have opposing effects on the underlying components of the RT distributions. In fact, such a tradeoff in parameters was an important observation made by Heathcote et al. (1991), which was subsequently replicated by Spieler, Balota, and Faust (1996). Specifically, in a color naming Stroop task, the congruent condition, compared to the neutral condition, decreased μ but increased τ. Since the mean is the sum of these two parameters, there was no influence

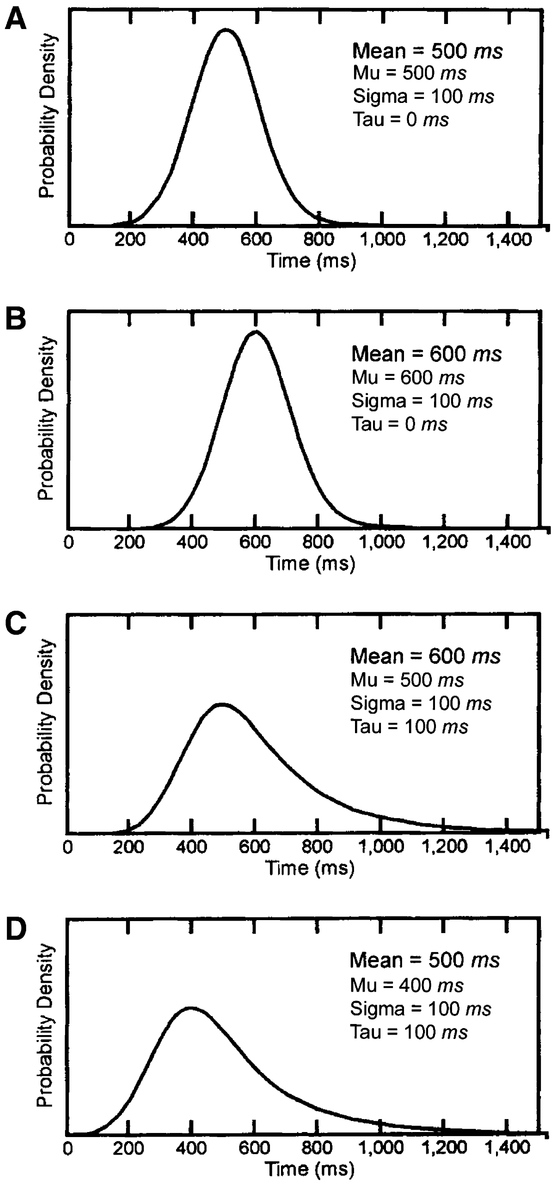

Fig. 3. Possible changes in distributions and the underlying influences on mean estimates and the parameter estimates from the ex-Gaussian analyses.

on the mean. Hence, it is possible that systematic tradeoffs in aspects of the RT distributions can mask differences in mean performance. Of course, such tradeoffs can have important implications for computational models (see Mewhort, Braun, & Heathcote, 1992).

Vincentile analyses

In order to more directly estimate the influence of a variable on RT distributions, parameter estimates from underlying functions such as the ex-Gaussian may be supplemented by analyses of Vincentiles (or Quantiles). Vincentile analyses provide mean estimates of ascending bins of RTs for each condition. In these analyses, one orders the RTs (from fastest to slowest) within each condition and then plots the mean of the first 10%, the second 10%, and so on. One can then plot the mean of the Vincentiles across participants to obtain a description of how the RT distribution is changing across conditions. Importantly, one can also plot the differences between two levels of a variable across Vincentiles to better understand how the influence of a variable may change as a function of the location in the RT distribution. These are functionally equivalent to delta plots (see Bub, Masson, & Lalonde, 2006).

Vincentile analyses should converge with the ex-Gaussian parameter estimates in systematic ways. Consider, for example, the idealized data in which a variable simply shifts the RT distribution, which is reflected by a change in μ. This is shown in Fig. 4A in the closed circles. On the other hand, consider how a variable that only changes the tail of the distribution (i.e., τ) would look in the Vincentiles. This is shown in the open circles in the same figure. Sigma can also influence the nature of the observed Vincentiles. Here, the change in the size of σ (assuming no influence in other parameters) will produce a set of functions that leverage at the midpoint. This is shown in Fig. 4B. Of course, variables do not simply influence one parameter, but typically influence multiple parameters. As shown below, the signature influence of a parameter change in the Vincentiles can be particularly helpful in further understanding how that variable is influencing the underlying RT distribution.

Distributional analyses of standard lexical variables

Now that we are armed with some preliminary tools for RT distributional analyses, let us consider the influence of variables on underlying RT distributions. There has already been work investigating how variables influence the underlying RT distributions in lexical decision and pronunciation performance (e.g., Andrews & Heathcote, 2001; Balota & Spieler, 1999; Plourde & Besner, 1997; Ratcliff et al., 2004; Yap & Balota, 2007; Yap, Balota, Cortese, & Watson, 2006). Interestingly, these studies typically show that variables often both shift *and* skew RT distributions. For example, Fig. 5 shows the influence of a set of standard lexical variables (word frequency, stimulus degradation, lexicality, and animacy) on lexical decision, pronunciation, and semantic classification performance. The top four panels are taken from the Yap et al. (2006) and the Yap and Balota (2007) studies, and the bottom four panels are from Andrews and Heathcote (2001). As one can see, the effect of these variables increases across Vincentiles, and is typically reflected by changes in both the μ and τ estimates.

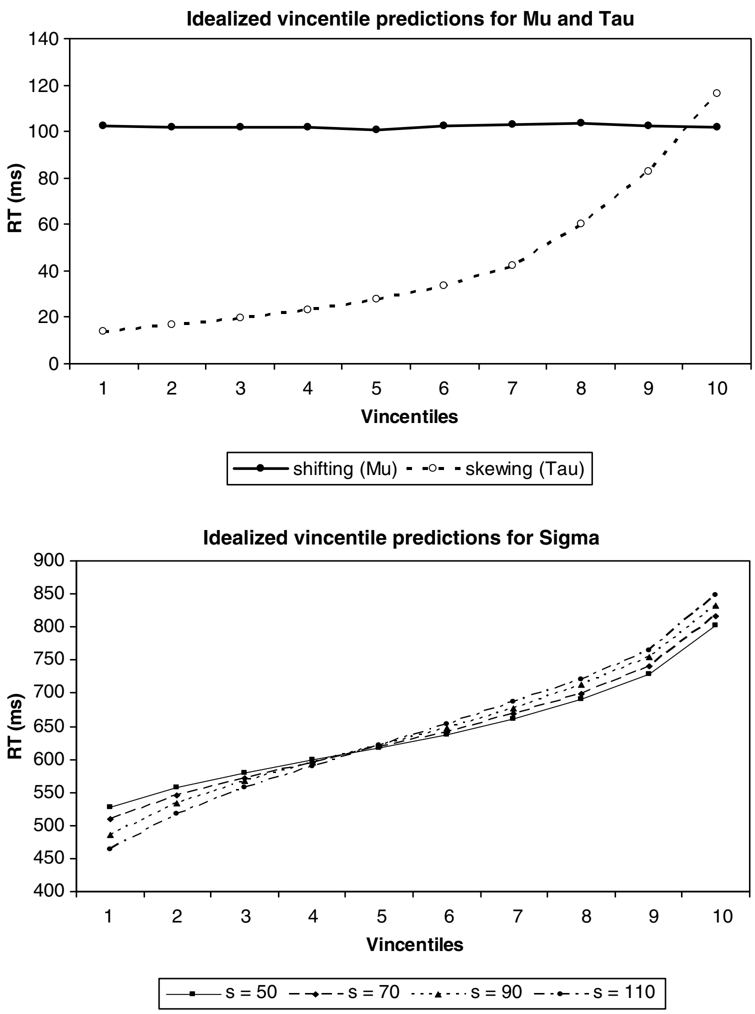

Fig. 4. Isolated effects of changes in the ex-Gaussian parameters on the underlying Vincentiles.

Extending distributional analyses to semantic priming

In the current study, we use distributional analyses to examine the semantic priming effect, which is one of the most widely studied effects in cognitive psychology (see Neely, 1991, for a review). As noted earlier, this effect simply reflects the facilitation of a speeded lexical decision or pronunciation response to a target that follows a related word, compared to when it follows an unrelated word. This effect has been central to computational models of memory retrieval (e.g., Masson, 1995; Plaut & Booth, 2000; Ratcliff & McKoon, 1988), distinctions between automatic and attentional processes (e.g., Balota, 1983; Neely, 1977), the nature of semantic/associative representations (e.g., Balota & Paul, 1996; Jones, Kintsch, & Mewhort, 2006; McRae, De Sa, & Seidenberg, 1997), and recent neuroimaging investigations (e.g., Gold et al., 2006; Martin, 2005).

Given the available evidence regarding how variables typically influence RT distributions (see Fig. 5), one might expect to find both a shift and an increase in the tail of the RT distribution as a function of semantic relatedness. This also appears to be most compatible with the predictions from the available computational models. For example, according to the compound cue model (Ratcliff & McKoon, 1988), priming influences the drift rate in a diffusion process. If a variable has an isolated effect on the drift rate, the most straightforward prediction would be a change in μ, σ, and τ in the distribution.[2] Simulations with Masson's (1995) feature

[2] However, it should also be noted that it is possible to produce isolated effects on some of the parameters when there is a change in drift rate from a simple random-walk process (a relative of the diffusion model) in lexical decision performance (e.g., Yap et al., 2006).

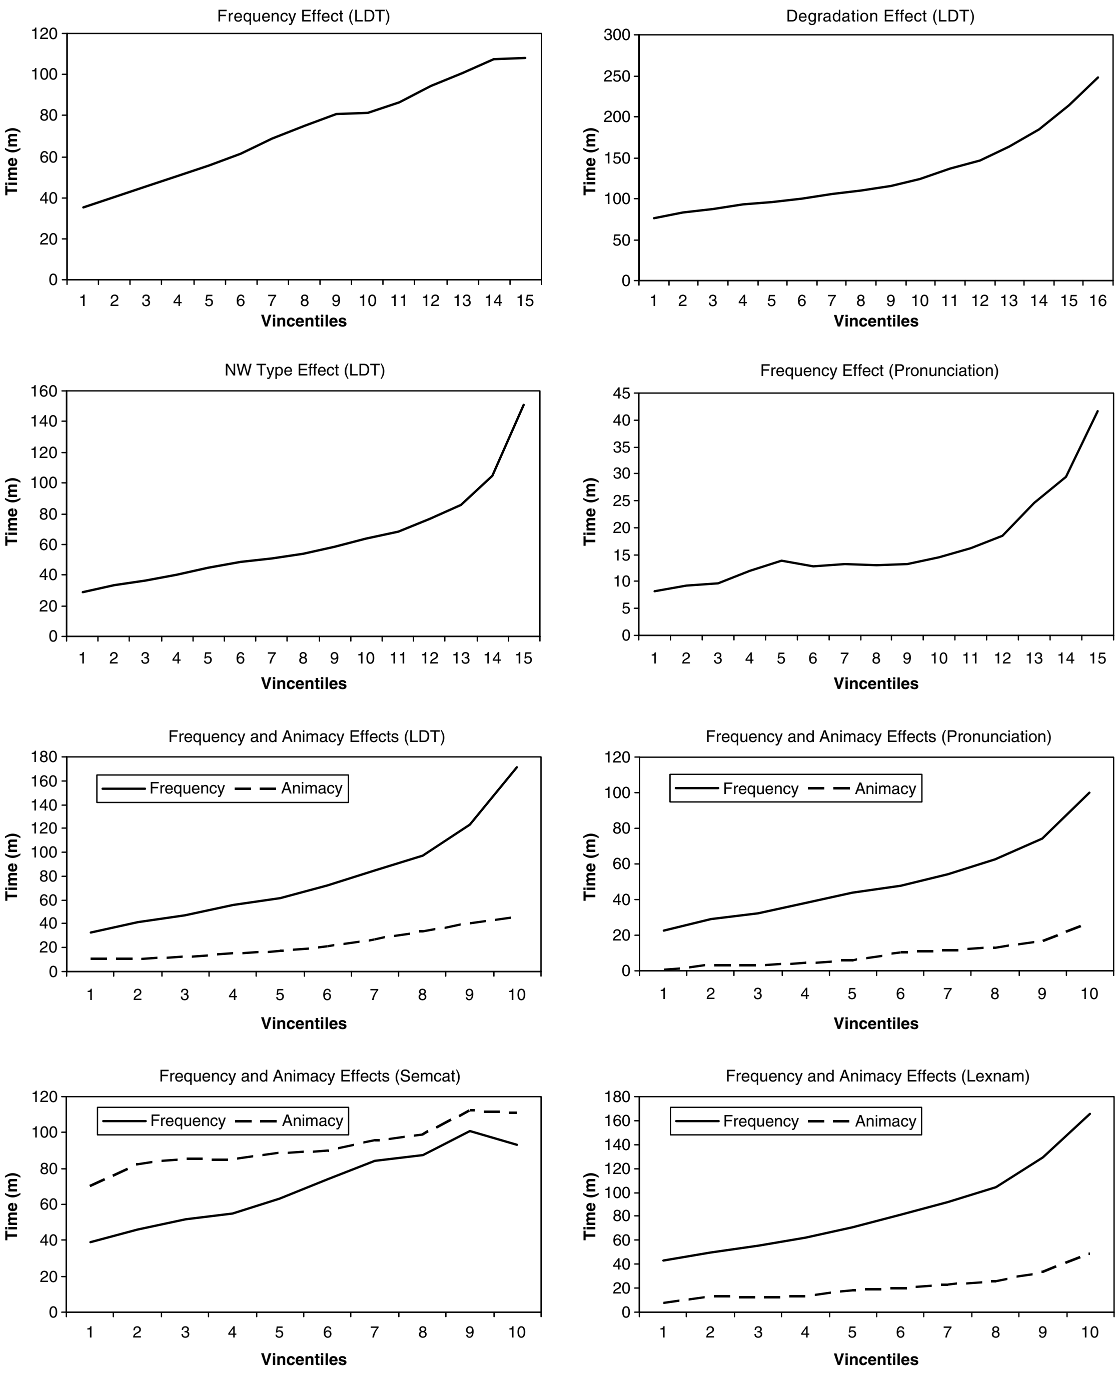

Fig. 5. Examples of the effects of standard word recognition variables as a function of Vincentiles.

overlap model would also appear to predict changes across the parameters (Spieler, 2000, personal communication). Finally, one might argue that the most straightforward prediction from the Plaut and Booth (2000) model would be a non-linear change in the RT distribution, because this model relies heavily on the non-linear logistic function relating settling times (a proxy for RT) to prime-target featural overlap. Where one is at on this function depends upon such variables as word frequency, reading skill, and stimulus degradation. Although possible, it is unlikely that one would hit the “sweet spot” within the Plaut and Booth model, and find a simple shift in the RT distribution as a function of prime-target relatedness.

In contrast to the available computational models, one might expect additive effects of prime relatedness based on metaphorical pre-activation (e.g., Neely, 1977) or headstart mechanisms (e.g., Forster, Mohan, & Hector, 2003). Specifically, consider the possibility that the prime produces some amount of activation for the target and this pre-activation, assuming sufficient time has passed, is completed before the target is presented. Such a simple pre-activation (headstart) account would predict a simple shift in the RT distributions as a function of prime-target relatedness.

Overview of the present experiments

In the first two experiments, we explored semantic priming effects across two dimensions that have been widely investigated in the priming literature. The first experiment used the speeded pronunciation task. One group of participants received the prime-target pairs at a relatively short stimulus onset asynchrony (SOA) of 250 ms, while a second group of participants received the prime-target pairs at a relatively long SOA of 1250 ms. This SOA manipulation has been well-studied since the seminal paper by Neely (1977). Based on the Posner and Snyder (1975) framework, Neely predicted that the short SOA should be more reflective of an automatic spreading activation process, whereas the long SOA should be more reflective of a limited capacity attentional component. Indeed, in an elegant demonstration of converging operations, Neely provided evidence for such an automatic/attentional dissociation across a set of variables. In fact, one could argue that SOA manipulations have been the central way of distinguishing between more automatic and more attentional processes (e.g., Balota, 1983; Balota, Black, & Cheney, 1992; Burke, White, & Diaz, 1987; den Heyer, Briand, & Dannenbring, 1983; Favreau & Segalowitz, 1983; Swinney, 1979).

The second experiment was identical to the first experiment, except that the lexical decision task (LDT) was used. There has been considerable interest in the locus of semantic priming in speeded pronunciation versus lexical decision, with some researchers arguing that the pronunciation task is a purer measure of pre-lexical influences of primes on target processing (see, for example, Balota & Lorch, 1986; Seidenberg, Waters, Sanders, & Langer, 1984). Neely (1991) has argued, and subsequently demonstrated, that priming in lexical decision performance reflects both a prelexical forward influence from the prime to the target and a postlexical retrieval process. This postlexical process reflects the possibility that participants can use the relation between the prime and target to bias the “word” response in lexical decision performance. Specifically, if the target is related to the prime, it must be a word, because nonwords are never related to primes. Given the possibility that the influence of this check process may not be involved on all trials, one might expect differences in the influence of semantic priming on the underlying RT distributions across lexical decision and speeded pronunciation performance. Experiment 3 provides a replication of the short SOA lexical decision results.

The final four experiments explore the utility of RT distributions in understanding the joint effects of multiple variables in both speeded pronunciation and lexical decision performance. Here, we target the robust interaction between stimulus degradation and semantic relatedness. These studies nicely extend and replicate the pattern observed in the first set of experiments and further demonstrate how RT distributional analyses can be particularly insightful for understanding the nature of the interactions across variables.

Experiment 1: Effects of relatedness and SOA in pronunciation

Method

Participants

All participants in the present experiments were recruited from the Washington University undergraduate psychology pool, had normal or corrected-to-normal vision, and participated for course credit. Forty-eight participants were in Experiment 1.

Stimuli

Three hundred words served as targets (see Table 1 for summary statistics for the primes and targets). Related targets were the primary associates of the primes according to the Nelson, McEvoy, and Schreiber (1998) norms, and unrelated prime-target pairs were not associates (i.e., forward and backward associate strength = .000). Stimulus pairs were constructed so that each target was paired with a related, unrelated, and

Table 1
Stimulus characteristics of the words used in the experiment

Factor	Means (*SD*s in parentheses)
Prime frequency	8.56 (2.0)
Prime length	5.44 (1.83)
Target frequency	10.09 (1.58)
Target length	4.83 (1.12)
Prime-target forward associative strength	.660 (.115)
Prime-target backward associative strength	.206 (.222)

Note: Frequency values = log HAL (Lund & Burgess, 1996) norms. Associative strength was determined according to the Nelson et al. (1998) norms.

neutral prime (i.e., the word "BLANK").[3] Three lists were constructed via random assignment of target word to prime condition. In each of these lists, 100 target words were preceded by a related prime, 100 target words were preceded by an unrelated prime, and 100 target words were preceded by a neutral prime. Lists were counterbalanced across participants such that each target word occurred equally often in each of the three prime contexts. Targets were initially randomly assigned to conditions for a given list. The lists were divided into four blocks, each consisting of 25 related pairs, 25 unrelated pairs, and 25 neutral pairs. Block order was also counterbalanced across participants such that each block of stimuli appeared equally often in the first, second, third, or fourth position throughout the experiment.

Procedure

A microcomputer with a 133 mHz processor running in DOS mode was used to control the experiment. A 17-in. monitor was set to 40-column mode for stimulus presentation. Vocal responses triggered a voice key (Gerbrands G1341T) connected to the PC's real-time clock, which recorded response latencies to the nearest ms.

Words were presented at the center of the computer screen individually in white uppercase letters against a black background. Within each block, the presentation order was random. Ten practice trials preceded the experimental trials. Participants were instructed to silently read the first word and to read aloud the second word as quickly and accurately as possible. Each trial began with a blank screen for 2000 ms followed by a fixation stimulus (+) appearing in the center of the screen for 1000 ms. After the fixation stimulus, the prime appeared either for 200 ms (short SOA) or 1000 ms (long SOA). The prime was followed by a blank screen for 50 ms (short SOA) or 250 ms (long SOA). The blank screen was replaced by the target, which remained on the screen until the vocal response triggered the voice key. After the pronunciation response, the experimenter coded the trial as correct, incorrect (mispronunciation), or noise (i.e., some extraneous noise triggered the voice key or it failed to be triggered by the pronunciation response). The coding of the response initiated the next trial sequence. A mandatory one-minute break occurred after each block of trials.

Design

Relatedness (related, unrelated, neutral) was manipulated within participants, and SOA (short, long) was manipulated between participants. The dependent variables were response latency and accuracy rate.

Results and discussion

Errors (3.1% across both conditions) and response latencies faster than 200 ms or slower than 1500 ms were first excluded from the analyses. Based on the remaining observations, the overall mean and *SD* of each participant's pronunciation latencies were computed. Response latencies 2.5 *SD*s above or below each participant's respective mean latency were removed. These criteria eliminated a further 2.1% of the responses. ANOVAs were then carried out on the mean, accuracy, and the ex-Gaussian parameters of the RT data. The mean response latencies, accuracies, and ex-Gaussian parameters are displayed in Table 2.[4]

[3] In the first two experiments, we included a neutral prime condition, but for the present purposes have decided not to include this estimate to make inferences about facilitation of the related and inhibition of the unrelated conditions, respectively. There are two reasons for this decision. First, Jonides and Mack (1984) have convincingly argued that finding an appropriate neutral baseline (equated on all dimensions with the other prime conditions) to measure facilitation and inhibition effects is nearly impossible. We are particularly concerned about the influence of differences in the RT distribution across neutral and word-type prime stimuli, since these stimuli will have different alerting characteristics across trials due to repetition of the neutral primes. Second, we did not include the neutral condition in the later experiments, and so for ease of comparison we do not include these data in the main tables. However, analyses of the neutral condition in Experiment 1 and 2 are provided here for interested readers. For short SOA pronunciation ($M_{\text{related}} = 474$ ms, $M_{\text{neutral}} = 486$ ms, $M_{\text{unrelated}} = 491$ ms), both facilitation ($p = .001$) and inhibition ($p = .017$) were significant. For long SOA pronunciation ($M_{\text{related}} = 467$ ms, $M_{\text{neutral}} = 503$ ms, $M_{\text{unrelated}} = 504$ ms), only facilitation was significant ($p < .001$). For short SOA LDT ($M_{\text{related}} = 569$ ms, $M_{\text{neutral}} = 614$ ms, $M_{\text{unrelated}} = 609$ ms), only facilitation was significant ($p < .001$). For long SOA LDT ($M_{\text{related}} = 592$ ms, $M_{\text{neutral}} = 638$ ms, $M_{\text{unrelated}} = 649$ ms), both facilitation ($p < .001$) and inhibition ($p = .039$) were significant. In general, with this neutral prime, the facilitatory effects appear more powerful than the inhibitory effects.

[4] In addition to examining RTs for correct trials, we also report RTs for *error* trials (along with standard errors) as a function of condition for each of the Experiments in the Appendix. These data are based only on participants who had at least one error in both the related and unrelated conditions. As shown in the Appendix this greatly reduced the number of participants in each experiment, and especially for pronunciation. Generally, error RTs were slightly longer than accurate RTs. Furthermore, the effect of relatedness on error RTs was not significant in any of the experiments, with the exception of the masked priming in lexical decision with degraded targets (Experiment 7), where the difference approached significance, $p < .10$. However, because of the paucity of data in these analyses, and the possibility that error trials may have multiple distinct causes, one needs to exert caution in interpreting these results.

Table 2
Mean response latency, percent error rates, and ex-Gaussian parameters as a function of Stimulus Onset Asynchrony, and Prime-Target Relatedness for pronunciation performance in Experiment 1

	Mean	%Errors	μ	σ	τ
Short SOA					
Related	474	1.5	438	49	36
Unrelated	491	2.5	459	49	32
Effect	17	1.0	21	0	−4
Long SOA					
Related	467	4.5	431	59	37
Unrelated	504	3.8	470	56	34
Effect	37	−.7	39	−3	−3
Interaction	20	−1.7	18	−3	1

Response latencies

The main effect of relatedness was significant by participants and items, $F_p(1,46) = 106.11$, $p < .001$, $MSE = 161.32$, $\eta^2 = .70$; $F_i(1,299) = 247.53$, $p < .001$, $\eta^2 = .45$. The main effect of SOA was not significant by participants ($F_p < 1$) or items, $p = .10$. The relatedness × SOA interaction, $F_p(1,46) = 15.53$, $p < .001$, $MSE = 161.32$, $\eta^2 = .25$; $F_i(1,299) = 32.04$, $p < .001$, $MSE = 898.97$, $\eta^2 = .10$, was also significant, with larger relatedness effects at the long SOA condition.

Percent correct

There was no main effect of relatedness in accuracy data by participants or by items, F_p and $F_i < 1$. The main effect of SOA was significant by participants and by items, $F_p(1,46) = 17.54$, $p < .001$, $MSE = .00061$, $\eta^2 = .28$; $F_i(1,299) = 34.00$, $p < .001$, $MSE = .0039$, $\eta^2 = .10$. The relatedness × SOA interaction was also significant, $F_p(1,46) = 6.34$, $p = .015$, $MSE = .00028$, $\eta^2 = .12$; $F_i(1,299) = 6.56$, $p = .011$, $MSE = .0033$, $\eta^2 = .02$; the relatedness effect (higher accuracy for related targets) was significant in the short ($p = .019$), but not long, SOA condition.

Ex-Gaussian analyses

Ex-Gaussian parameters (μ, σ, τ) were obtained for each participant using continuous maximum likelihood estimation (CMLE) in R (R Development Core Team, 2004). CMLE provides efficient and unbiased parameter estimates (Van Zandt, 2000) while using all the available raw data. Using Nelder and Mead's (1965) simplex algorithm, negative log-likelihood functions were minimized in the R statistics package (c.f., Speckman & Rouder, 2004), with all fits successfully converging within 500 iterations. An alternative approach is to fit a specific set of quantiles (e.g., Heathcote, Brown, & Cousineau, 2004). An excellent website for both continuous and quantile fitting functions is available at http://www.newcastle.edu.au/school/psychology/ncl/software_repository.html (see Brown & Heathcote, 2003, for further description).

For μ, the main effect of relatedness, $F(1,46) = 63.46$, $p < .001$, $MSE = 343.48$, $\eta^2 = .58$, and the interaction, $F(1,46) = 5.79$, $p = .020$, $MSE = 343.48$, $\eta^2 = .11$, were significant, with larger relatedness effects in the long SOA condition. However, only the main effect of SOA was significant, $F(1,46) = 6.31$, $p = .016$, $MSE = 260.60$, $\eta^2 = .12$ for σ. Turning to τ, none of the effects were significant.

In summary, Table 2 shows that the relatedness effects for both short and long SOA targets are mediated by the μ component, indicating that the semantic priming effect is largely reflected by distributional shifting. Interestingly, the priming × SOA interaction is also mediated by μ, suggesting that the larger priming effects observed at the long SOA primarily reflected greater shifting for related, compared to unrelated, targets.

Vincentile analysis

As noted, a converging procedure for distributional analysis is to plot the mean Vincentiles for the data. Vincentizing averages RT distributions across participants (Andrews & Heathcote, 2001; Ratcliff, 1979; Rouder & Speckman, 2004; Vincent, 1912) to produce the RT distribution for a *typical* participant. This approach does not make any distributional assumptions, and examines the raw data directly. In the present data, we first ordered the data from fastest RT to slowest RT for each subject within each condition. Then, we calculated the mean of the first 10%, the next 10%, etc. Vincentile plots are then computed by collapsing across the same bins across subjects.

The mean Vincentiles for the different experimental conditions are plotted in the top two-thirds of Fig. 6, with the bottom third of Fig. 6 being the mean relatedness effect as a function of Vincentiles and SOA. Note that for the top two panels, the empirical mean Vincentiles are represented by data points and standard error bars, while the estimated Vincentiles for the respective best-fitting ex-Gaussian distribution are represented by lines. Presenting the data in this manner is useful because it allows one to visually assess the extent to which empirical and estimated Vincentiles overlap, providing a measure of goodness of fit. Clearly, the data are fitted well by the ex-Gaussian distribution, and the divergence between mean Vincentiles and theoretical ex-Gaussian Vincentiles is typically smaller than one standard error in most cases. The bottom panel, which presents difference scores, depicts only empirical Vincentiles.

In agreement with the ex-Gaussian analysis, it is clear from Fig. 6 that the semantic priming effect in speeded pronunciation is mediated by distributional shifting at both the short and long SOAs, since, within each SOA condition, the magnitude of the priming effect is approx-

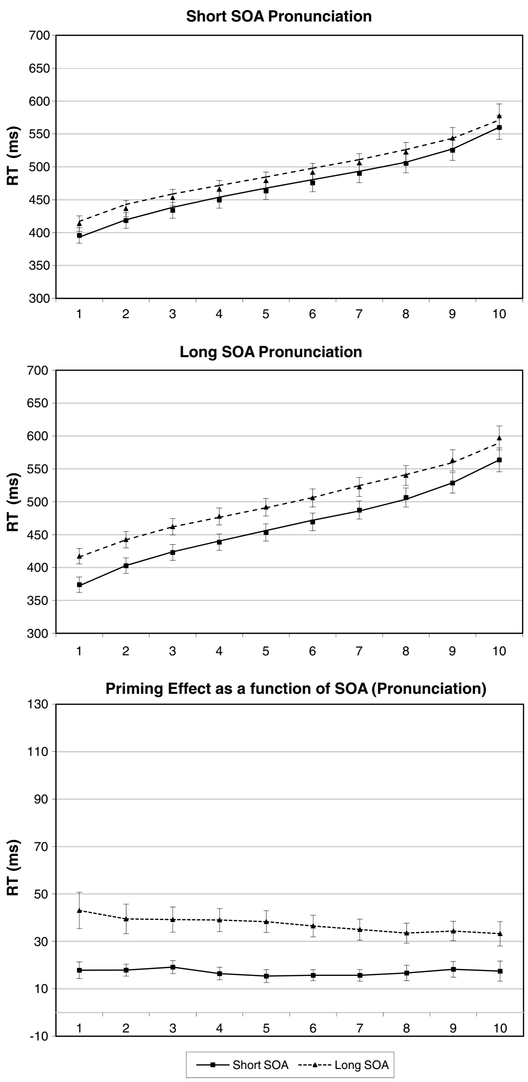

Fig. 6. Pronunciation performance from Experiment 1 as a function of prime relatedness and Vincentiles in the short SOA (top panel) and long SOA (middle panel) conditions, along with the priming effects as a function of Vincentiles (bottom panel). In the top and middle panels, participants' mean Vincentiles (■ = related, ▲ = unrelated) are represented by data points and standard error bars. Best-fitting ex-Gaussian Vincentiles are represented by lines (solid line = related, dashed line = unrelated).

imately the same across the Vincentiles. Likewise, as discussed earlier, the relatedness × SOA interaction seems to be reflected largely by more pronounced shifting for the long SOA targets. In order to explore the reliability of this pattern, we conducted an ANOVA with Vincentile as a factor. Importantly, in the present and subsequent analyses of the Vincentiles, we used the Greenhouse–Geisser correction for potential violations of sphericity note that all Greenhouse–Geisser corrected dfs were rounded to the nearest whole number. The results from this analysis indicated that neither the relatedness by Vincentile ($p = .31$) nor the relatedness × SOA × Vincentile interaction ($F < 1$) approached significance, confirming that the effect of relatedness is relatively constant across the RT distribution, i.e., reflecting a simple shift.

In summary, the results from Experiment 1 indicate that in speeded pronunciation performance, the influence of semantic priming is a shift in the RT distribution. This pattern occurred at both the short and long SOAs, even though there was evidence of larger relatedness effects at long SOAs. Given how other variables affect RT distributions in word recognition experiments (see Introduction), and the predictions from computational models, this pattern is surprising. These results appear most consistent with simple pre-activation (headstart) metaphors of priming in which the prime pre-activates (provides a headstart in processing) the target's lexical representation by some constant amount. Before drawing inferences from these results, it is important to determine if a similar pattern exists in the LDT, which has been the primary target for the computational models of semantic priming.

Experiment 2: Effects of relatedness and SOA in lexical decision

Method

Participants

Sixty undergraduates participated in Experiment 2.

Stimuli

Words were those employed in Experiment 1. Pronounceable nonwords served as distracters and were constructed by changing one or two letters in the target words. There were four blocks of trials each consisting of 75 prime-nonword pairs intermixed with 75 prime-word pairs. Otherwise, the block composition was the same as Experiment 1.

Procedure

The procedure for Experiment 2 was the same as that employed in Experiment 1, with the following exceptions: First, in Experiment 2, there were four blocks of 150 trials. Second, in Experiment 2, participants responded to each target by pressing either a key labeled "YES" (the *slash* key) for a word decision, or one labeled "NO" (the *Z* key) for a nonword decision. A 1500 ms blank screen followed correct responses. For incorrect responses, a 200 Hz sound occurred for 750 ms while the message "incorrect response" appeared. A blank screen lasting 750 ms followed this message. Third, ten lexical decision (5 word and 5 nonword) trials preceded the test trials.

Design

Relatedness (related, unrelated, neutral) was manipulated within participants, and SOA (short, long) was manipulated between participants. The dependent variables were response latency and accuracy rate.

Results and discussion

Errors (3.0% across both conditions) and response latencies faster than 200 ms or slower than 3000 ms were first excluded from the analyses. Using the trimming criteria described in Experiment 1, a further 3.1% of the responses were removed. The mean RT, accuracy, and the ex-Gaussian parameters are displayed in Table 3.

Response latencies

For mean response latencies, the main effect of relatedness was significant, $F_p(1,58) = 126.50$, $p < .001$, $MSE = 525.77$, $\eta^2 = .69$; $F_i(1,299) = 79.42$, $p < .001$, $MSE = 9845.14$, $\eta^2 = .21$. The main effect of SOA was significant by items but not by participants, $F_p < 1$; $F_i(1,299) = 80.86$, $p < .001$, $MSE = 3686.77$, $\eta^2 = .21$. The interaction was significant by items and approached significance by participants, $F_p(1,58) = 3.98$, $p = .051$, $MSE = 525.77$, $\eta^2 = .06$; $F_i(1,299) = 6.03$, $p = .015$, $MSE = 3789.10$, $\eta^2 = .02$, with larger relatedness effects in the long SOA condition.

Percent correct

The accuracy data yielded a main effect of relatedness, $F_p(1,58) = 32.68$, $p < .001$, $MSE = .00052$, $\eta^2 = .36$; $F_i(1,299) = 45.44$, $p < .001$, $MSE = .0040$, $\eta^2 = .13$. The main effect of SOA was not significant by participants, $F_p < 1$, and approached significance by items, $p = .073$. The interaction was not significant by participants or by items.

Table 3
Mean response latency, percent error rates, and ex-Gaussian parameters as a function of Stimulus Onset Asynchrony, and Prime-Target Relatedness for lexical decision performance in Experiment 2

	Mean	%Errors	μ	σ	τ
Short SOA					
Related	569	2.2	433	42	136
Unrelated	609	4.6	478	62	131
Effect	40	2.4	45	20	−5
Long SOA					
Related	592	1.6	454	66	138
Unrelated	649	4.0	503	63	145
Effect	57	2.4	49	−3	7
Interaction	17	0	4	−23	12

Ex-Gaussian analyses

For μ, only the main effect of relatedness was significant, $F(1,58) = 21.85$, $p < .001$, $MSE = 2869.68$, $\eta^2 = .27$. Turning to σ, the relatedness × SOA interaction approached significance, $F(1,58) = 3.94$, $p = .052$, $MSE = 983.08$, $\eta^2 = .06$, with larger relatedness effects in the short SOA condition. For τ, none of the effects were significant.

In sum, the results from the ex-Gaussian analyses show that consistent with the pronunciation results from Experiment 1, the relatedness effects at the long SOA condition in Experiment 2 primarily reflect distributional shifting, wherein there is only a change in the μ parameter as a function of prime relatedness. However, when the SOA is short, the parameter estimates provided a slightly different story. Here, both μ and σ are larger for unrelated, compared to related, targets. We now turn to the Vincentile analyses to determine if there is convergence with these parameter estimates.

Vincentile analysis

The mean Vincentiles for the different experimental conditions are plotted in Fig. 7, along with the best fitting ex-Gaussian distribution. Fig. 7 (bottom panel) shows that for the long SOA condition, the semantic relatedness effect is mediated mainly by distributional shifting. In contrast, in the short SOA condition, the magnitude of the relatedness effect increases monotonically across Vincentiles. Relatedness effects are smallest in the fastest Vincentiles, and increase as the Vincentiles become slower. Statistical support for this observation was provided by a Vincentile by relatedness analysis which indicated that at the short SOA, the interaction between Vincentile and relatedness approached significance, $F(2,45) = 2.91$, $p = .065$, $MSE = 1454.41$, $\eta^2 = .11$, whereas there was no hint of such an interaction at the long SOA, $F < 1$. To summarize, for the long SOA condition, priming reflects mainly shifting (μ), but for the short SOA condition, priming involves both shifting (μ) and some influence in σ.

Experiment 3: A replication of short SOA priming in lexical decision

Overall, the results from the first two experiments indicate that semantic priming primarily reflects distributional shifting. The only discrepant pattern was found at the short SOA lexical decision results, wherein there was evidence that the relatedness effect increased systematically across Vincentiles, and this was primarily reflected in a change in σ in the ex-Gaussian analysis. Before discussing the implications of this pattern, an attempt was made to replicate the pattern observed at the short SOA condition in Experiment 2. Such a repli-

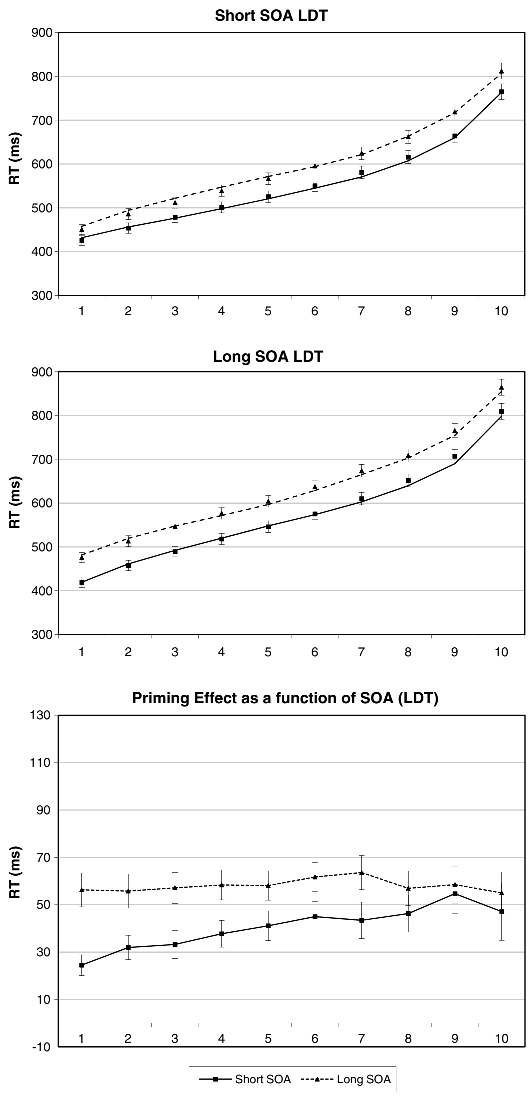

Fig. 7. Lexical decision performance from Experiment 2 as a function of prime relatedness and Vincentiles in the short SOA (top panel) and long SOA (middle panel), along with the priming effect as a function of Vincentiles (bottom panel). In the top and middle panels, participants' mean Vincentiles (■ = related, ▲ = unrelated) are represented by data points and standard error bars. Best-fitting ex-Gaussian Vincentiles are represented by lines (solid line = related, dashed line = unrelated).

cation would also provide further support for the stability of RT distributional analyses.

Method

Participants

Sixteen undergraduates participated in Experiment 3.

Procedure and design

The design was identical to Experiment 2, with the following exceptions. First, in Experiment 3, only the short SOA condition was included, and the neutral condition was omitted, so there were 150 observations per cell. Second, in Experiment 3, participants responded to word targets by pressing the *apostrophe* key and to nonword targets by pressing the *A* key. Finally, each trial began with a fixation mark (+) appearing on the center of the screen for 2000 ms, followed by the prime for 150 ms, then by a blank screen for 100 ms. The blank screen was replaced by the target, which remained on the screen until a button press was detected. For incorrect responses, a 170 ms tone was presented simultaneously with "Incorrect" displayed for 450 ms slightly below the fixation point.

Results and discussion

Errors (6.2% across both conditions) and response latencies faster than 200 ms or slower than 3000 ms were first excluded from the analyses. Using the trimming criteria described in Experiment 1, a further 2.9% of the responses were removed. The mean RT, accuracy, and ex-Gaussian parameters are displayed in Table 4.

For mean response latencies, the main effect of relatedness was significant by participants and by items, $t_p(15) = 5.96$, $p < .001$; $t_i(299) = 6.27$, $p < .001$. For accuracy, the main effect of relatedness was not significant by participants or by items. For μ and σ, the main effects of relatedness were highly significant, $t(15) = 6.72$, $p < .001$, $t(15) = 2.77$, $p = .014$ respectively. Turning to τ, the relatedness effect was not significant, $t < 1$. Table 4 shows that for the short SOA used in Experiment 3, μ and σ, but not τ, are larger for unrelated targets. This is a clear replication of the short SOA condition in Experiment 2.

Vincentile analysis

The mean Vincentiles for the related and unrelated conditions, along with the best fitting ex-Gaussian distribution are displayed in the top two panels of Fig. 8. The difference scores across related and unrelated conditions are plotted in the bottom panel of Fig. 8. As shown at

Table 4
Mean response latency, percent error rates, and ex-Gaussian parameters as a function of Prime-Target Relatedness for lexical decision performance in Experiment 3

	Mean	%Errors	μ	σ	τ
Short SOA					
Related	538	5.3	402	35	136
Unrelated	568	6.2	430	47	138
Effect	30	0.9	28	12	2

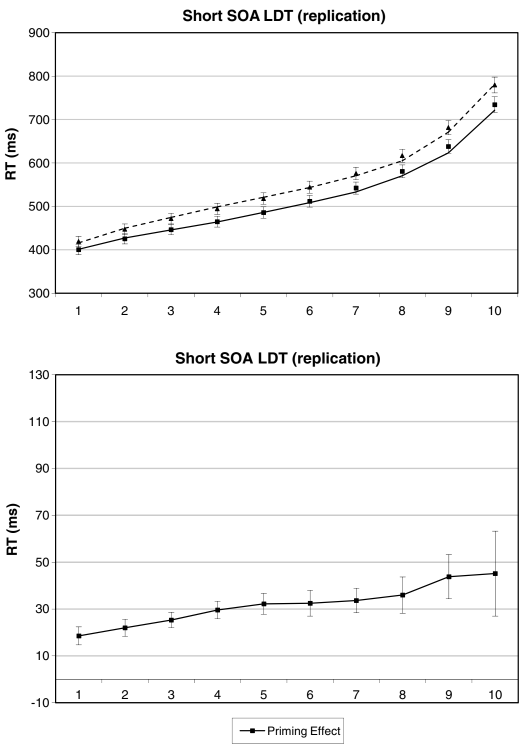

Fig. 8. Lexical decision performance for the related and unrelated conditions and Vincentiles for the Experiment 3 (top panel), and the priming effect as a function of Vincentiles (bottom panel). In the top panel, participants' mean Vincentiles (■ = related, ▲ = unrelated) are represented by data points and standard error bars. Best-fitting ex-Gaussian Vincentiles are represented by lines (solid line = related, dashed line = unrelated).

the bottom, the semantic relatedness effect increases monotonically across Vincentiles, and, as noted in the ex-Gaussian analyses, is mediated by reliable effects of shifting (μ) and an increase in σ. It should be noted, however, that the Vincentile by Relatedness interaction did not reach significance in this experiment, $p = .174$, likely due to the smaller number of participants in this study, compared to Experiment 2.

In summary, the results from Experiment 3 overall provide a replication of the short SOA results of Experiment 2. This pattern is quite distinct from the pronunciation data at both the short and long SOAs in Experiment 1, and also the long SOA lexical decision data in Experiment 2, wherein there is a simple shift in the RT distribution as a function of prime relatedness. Of course, the important question is why one might find the increasing relatedness effect across the RT distribution in the short SOA lexical decision experiments. One possibility is that when the SOA is short, there is insufficient time for the prime to be fully utilized before participants make their decision. Because this pattern was not found at the short SOA condition in speeded pronunciation, it appears specific to the operations in lexical decision. Consider the possibility that at the short SOA, there is a race between the word recognition processes that drive lexical decisions for the target in the unrelated condition, and the influence of the prime. The prime's effect could include both a forward-acting influence from the prime and a postlexical check process that is specific to lexical decision performance (see Neely, 1991). If this were the case, then the words that produce faster response latencies in the LDT in the unrelated condition will produce smaller relatedness effects, and as response latencies increase, there will be more time for the prime to influence target processing. This would produce the signature increasing relatedness effect across the Vincentiles displayed in Fig. 8. At this point, we will defer further discussion of this intriguing pattern until the General discussion.

Semantic priming and target degradation: Implications from RT distributions for interactive effects

We now turn to the utility of RT distributional analyses in understanding how multiple variables combine to influence visual word recognition performance. Although the first two experiments produced interactions, these effects included between-participant manipulations, and the larger priming effects at longer SOAs may reflect the influence of an additional predictive attentional mechanism (see Neely, 1991).

Distributional analyses can be particularly instructive regarding the stage where variables interact (see, for example, Roberts & Sternberg, 1993; Yap & Balota, 2007). Regarding semantic priming, one of the standard findings in the visual word recognition literature is that semantic priming effects increase when targets are visually degraded (see Becker & Killion, 1977; Borowsky & Besner, 1993; Meyer, Schvaneveldt, & Ruddy, 1975). Because degrading a stimulus is typically viewed as influencing a relatively early process in the visual word recognition flow, this pattern has been taken as evidence for interactions between top-down semantic support from related primes and early visual processing (e.g., Becker & Killion, 1977). This interaction has received considerable attention in recent discussions, because of the additional pattern that word frequency produces additive effects with stimulus degradation, but interactive effects with semantic priming, in lexical decision performance (see, for example, Borowsky & Besner, 2006; Plaut & Booth, 2006). Using additive factors logic, Borowsky and Besner (1993) argued that this pattern indicates that semantic priming has both an early influence reflected by the interaction between context and degradation, and a later influence reflected by the interaction between word frequency and semantic context. Here, we will focus on the early influence of context, i.e., its interaction with stimulus degradation (see Plourde & Besner, 1997; Yap & Balota, 2007, for a discussion of distributional analysis of the joint effects of word frequency and degradation).

In Experiments 4 and 5, we report two experiments that manipulate stimulus quality and semantic context in speeded pronunciation and lexical decision perfor-

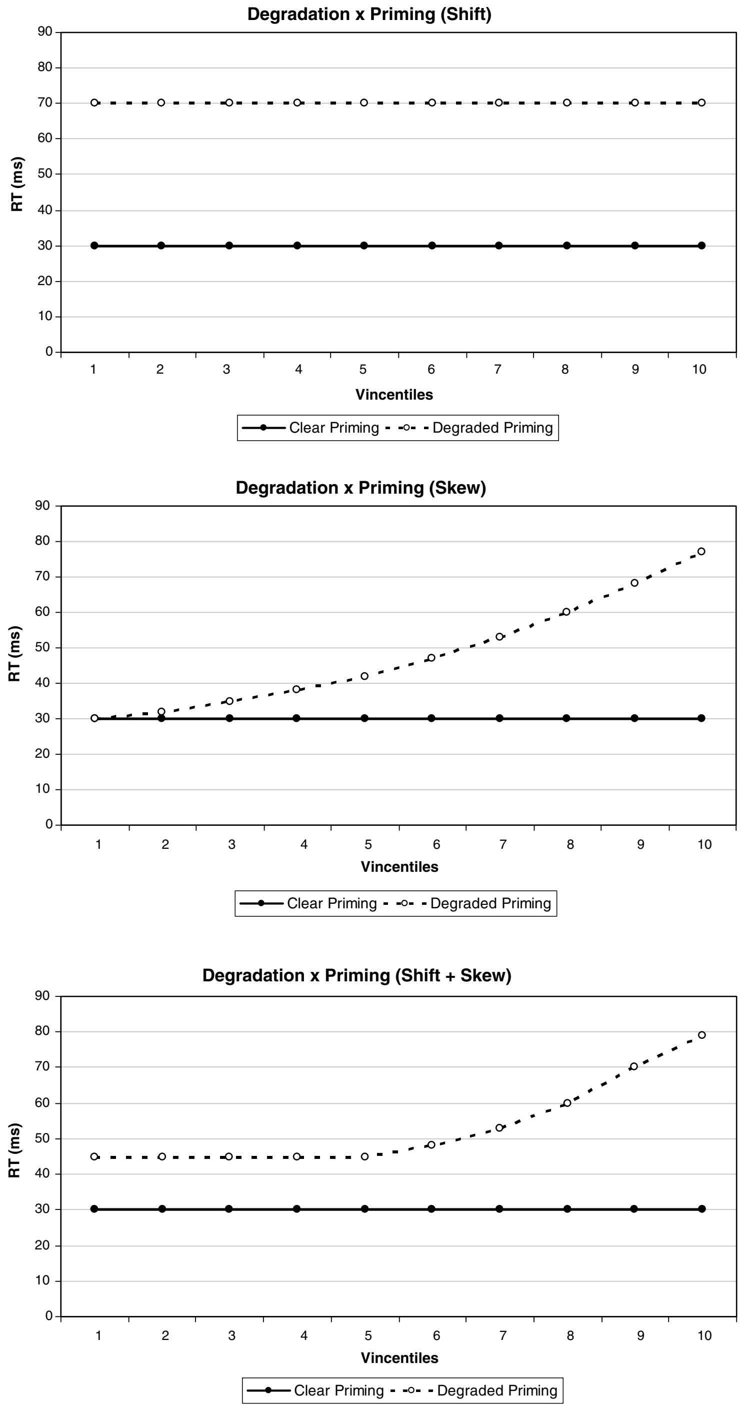

Fig. 9. Hypothetical priming effects as a function of target degradation. The top panel indicates that the priming by degradation interaction reflects an increase in priming (μ) across the entire RT distribution. The middle panel suggests that the priming by degradation interaction reflects an increase in skewing (τ) in the RT distribution. The bottom panel indicates that the priming by degradation interaction reflects both a shift and an increase in skewing.

mance, respectively. Because of the discrepant results of pronunciation and lexical decision performance at the short SOA, we decided to use a sufficiently long SOA (800 ms) in these experiments to determine if there is a task-independent pattern of interactive effects at the level of the RT distributions. Importantly, Experiments 4 and 5 also afford an opportunity to replicate the shift in RT distributions as a function of relatedness for the clear target conditions.

What should the nature of the stimulus quality by semantic context interaction in mean response latencies look like at the level of the RT distribution? If both stimulus degradation and semantic context only influence the same processing stage, the simplest of models might predict multiplicative effects of degradation. Because the first two experiments show that semantic priming reflects a shift in the RT distribution (at least at long SOAs), one might expect the pattern depicted at the top of Fig. 9, where the interaction is fully mediated by μ. On the other hand, it is possible that degradation of the target may produce an increase in response latency across the bins due to the increased difficulty of target processing, because of increased reliance on the primes for the more difficult targets. This pattern is most consistent with the pattern depicted in the middle panel of Fig. 9, where the interaction is fully mediated by τ. Finally, one might expect both a shift in the priming effect and also an additional effect for the particularly slow target words. This would be reflected in the pattern at the bottom panel of Fig. 9, where the interaction is mediated by both μ and τ.

Experiment 4: Effects of relatedness and stimulus quality in pronunciation

Method

Participants

Thirty-two undergraduates participated in Experiment 4.

Apparatus

An IBM-compatible computer was used to control stimulus presentation and to collect data. The stimuli were displayed on a 17-inch Super VGA monitor, and participants' pronunciation responses were detected by an Audio-Technica microphone connected to a PST serial response box with an integrated voice key.

Stimuli

The stimuli were the same set of 300 prime-target pairs used in the previous experiments. Across each group of four participants, targets were counterbalanced across related and unrelated conditions and degraded and clear conditions. No prime or target was repeated within a participant.

Procedure

Participants first received 10 practice trials followed by 4 experimental blocks of 75 trials, with mandatory breaks occurring between blocks. The presentation sequence was the same for both clear and visually degraded stimuli. Stimuli were presented in 14 point Courier font. For the degraded condition, letter strings were rapidly alternated with a randomly generated mask of the same length. For example, the mask &?# was presented for 14 ms, followed by DOG for 28 ms, and the two repeatedly alternated until the participant responded. The mask was generated from random permutations of the following symbols (@#$%&?*), with the proviso that the symbols were not repeated within a mask. Although masks across trials were uniquely randomly generated, the alternating masks within a trial were always the same. Each trial consisted of the following order of events: (a) a fixation point (+) at the center of the monitor for 2000 ms, (b) a prime for 150 ms, (c) a blank screen for 650 ms, and (d) the stimulus at the fixation point's location. The stimulus word remained on the screen until a pronunciation response was detected. Participants then coded their responses by pressing the left mouse button for a correct response and the right mouse button for an incorrect response.

Design

A 2×2 factorial design was used: both stimulus quality (clear vs. degraded) and relatedness (related vs. unrelated target) were manipulated within-participants.

Results and discussion

Errors (1.8% across both conditions) and response latencies faster than 200 ms or slower than 3000 ms were first excluded from the analyses. Any response latencies beyond 2.5 *SD*s from the mean were then excluded. A total of 2.9% of the responses were removed. The mean RT, accuracy, and ex-Gaussian parameters are displayed in Table 5.

Response latencies

For mean response latencies, the main effects of stimulus quality, $F_p(1,31) = 62.31$, $p < .001$, $MSE = 2437.67$, $\eta^2 = .67$; $F_i(1,299) = 242.14$, $p < .001$, $MSE = 6582.01$, $\eta^2 = .45$, and relatedness, $F_p(1,31) = 42.37$, $p < .001$, $MSE = 815.18$, $\eta^2 = .58$; $F_i(1,299) = 87.54$, $p < .001$, $MSE = 4572.57$, $\eta^2 = .23$, were significant. The stimulus quality $\times$ relatedness interaction was highly significant, $F_p(1,31) = 11.26$, $p = .002$, $MSE = 367.49$, $\eta^2 = .27$; $F_i(1,299) = 10.21$, $p = .002$, $MSE = 7189.72$, $\eta^2 = .03$. As shown in Table 5, there were larger relatedness effects for degraded targets compared to clear targets.

Percent correct

The accuracy data yielded main effects of stimulus quality, $F_p(1,31) = 7.50$, $p = .010$, $MSE = .0022$,

Table 5
Mean response latency, percent error rates, and ex-Gaussian parameters as a function of Target Degradation, and Prime-Target Relatedness for pronunciation performance in Experiment 4

	Mean	%Errors	μ	σ	τ
Clear targets					
Related	528	0.4	464	67	64
Unrelated	550	0.8	486	65	64
Effect	22	0.4	22	−2	0
Degraded targets					
Related	586	1.6	477	65	108
Unrelated	630	4.2	508	67	122
Effect	44	2.6	31	−2	14
Interaction	22	2.2	9	0	14

$\eta^2 = .20$; $F_i(1,299) = 37.54$, $p < .001$, $MSE = .0041$, $\eta^2 = .11$, and relatedness, $F_p(1,31) = 25.78$, $p < .001$, $MSE = .00027$, $\eta^2 = .45$; $F_i(1,299) = 31.76$, $p < .001$, $MSE = .0021$, $\eta^2 = .10$. The interaction between stimulus quality and relatedness was also significant, $F_p(1,31) = 10.78$, $p = .003$, $MSE = .00034$, $\eta^2 = .26$; $F_i(1,299) = 18.21$, $p < .001$, $MSE = .0019$, $\eta^2 = .06$, with larger relatedness effects for degraded targets.

Ex-Gaussian analyses

For μ, the main effects of stimulus quality, $F(1,31) = 9.61$, $p = .004$, $MSE = 1032.45$, $\eta^2 = .24$, and relatedness, $F(1,31) = 22.00$, $p < .001$, $MSE = 996.48$, $\eta^2 = .42$, were significant. The stimulus quality × relatedness interaction did not reach significance, $p = .22$. In contrast, none of the effects were significant for σ, Fs < 1. Turning to τ, the main effect of stimulus quality was significant, $F(1,31) = 23.27$, $p < .001$, $MSE = 3616.73$, $\eta^2 = .43$, and the stimulus quality x relatedness interaction approached significance, $F(1,31) = 3.40$, $p = .075$, $MSE = 455.42$, $\eta^2 = .10$. Separate analyses on the degraded and clear target conditions indicated that there was a reliable relatedness effect in τ for the degraded targets, $t(31) = 2.08$, $p = .045$, but not for the clear targets, $t < 1$. Table 5 shows that the semantic relatedness effect for clear words is mediated mainly by μ (distributional shifting), and that the stimulus quality × relatedness interaction is mediated primarily by τ (distributional skewing), and to a lesser extent, by μ (distributional shifting).

Vincentile analysis

The mean Vincentiles are plotted in Fig. 10, along with the best fitting ex-Gaussian function. Replicating the pattern of data from Experiment 1, for *clear* targets, semantic relatedness effects were mediated primarily by distributional shifting, since the magnitude of the relatedness effect was approximately the same across the entire RT distribution. For *degraded* targets, the larger relatedness effect, relative to clear targets, remained relatively invariant across the early and middle Vincentiles. However, for the later slow Vincentiles, the relatedness effect for degraded targets steadily grew as RTs became longer. In other words, although degradation increased relatedness effects for *all* the items, it had disproportionately large

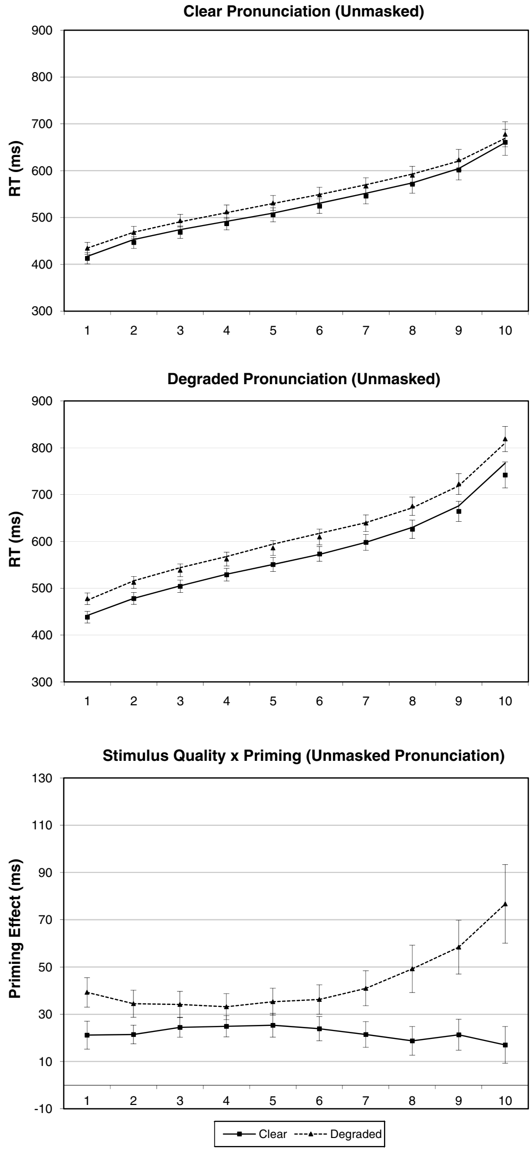

Fig. 10. Pronunciation performance from Experiment 4 as a function of prime relatedness and Vincentiles in the clear target (top panel) and degraded target (middle panel) conditions, along with the priming effects as a function of Vincentiles (bottom panel). In the top and middle panels, participants' mean Vincentiles (■ = related, ▲ = unrelated) are represented by data points and standard error bars. Best-fitting ex-Gaussian Vincentiles are represented by lines (solid line = related, dashed line = unrelated).

effects on the slowest items. This is consistent with a reliable stimulus quality × relatedness × Vincentile interaction, $F(2,72) = 7.15$, $p = .001$, $MSE = 1149.85$, $\eta^2 = .19$. Separate analyses on the clear and degraded items indicated that there was a highly reliable interaction between relatedness and Vincentile for the degraded items, $F(2,59) = 6.94$, $p = .002$, $MSE = 2147.99$, $\eta^2 = .18$, but not for the clear items, $F < 1$.

In sum, the pronunciation results from Experiment 4 replicate and extend the pronunciation results from Experiment 1. Specifically, for clear targets there was a simple shift in the RT distributions, precisely the same pattern found in Experiment 1. However, for degraded targets there is a clear increase in priming, and this occurs most strongly for the later portions of the RT distribution. This pattern does not reflect a simple multiplicative influence of degradation on semantic priming, but instead reflects primarily an increase in skewing and some more modest shifting of the distribution, as shown in the bottom panel of Fig. 10. We now turn to lexical decision performance to determine if this pattern is task-specific.

Experiment 5: Effects of relatedness and stimulus quality in lexical decision

Method

Participants

Thirty-two undergraduates participated in Experiment 5.

Apparatus

Identical to Experiment 4 except that participants' responses were made on a computer keyboard.

Stimuli

Includes all the stimuli used in Experiment 4, together with another 300 primes that were matched to the original 300 primes in terms of length, word frequency, and initial letter. These 300 new primes were then randomly paired with nonwords; the nonwords were formed by rearranging the letters of the new primes to form pronounceable nonwords. The same counterbalancing scheme used in Experiment 4 was used in Experiment 5.

Procedure

Participants pressed the *apostrophe* key for words and the *A* key for nonwords. Responses were followed by a 1600 ms delay. If the response was incorrect, 450 ms of that 1600ms was consumed by a 170 ms tone that was presented simultaneously with "Incorrect" displayed slightly below the fixation point. Participants were presented with 20 practice trials, followed by eight experimental blocks of 75 trials, with mandatory breaks occurring between blocks.

Results and discussion

Errors (5.8% across both conditions) and response latencies faster than 200 ms or slower than 3000 ms were first excluded from the analyses. Using the same criteria used in Experiment 4, a further 2.3% of the responses were excluded. The mean RT, accuracy, and the ex-Gaussian parameters are displayed in Table 6.

Response latencies

For mean response latencies, the main effects of stimulus quality, $F_p(1,31) = 117.68$, $p < .001$, $MSE = 3123.91$, $\eta^2 = .79$; $F_i(1,299) = 362.90$, $p < .001$, $MSE = 9786.72$, $\eta^2 = .55$, and relatedness, $F_p(1,31) = 47.66$, $p < .001$, $MSE = 1807.89$, $\eta^2 = .61$; $F_i(1,299) = 172.81$, $p < .001$, $MSE = 5121.53$, $\eta^2 = .37$, were significant. The stimulus quality × relatedness interaction was also significant, $F_p(1,31) = 22.33$, $p < .001$, $MSE = 531.11$, $\eta^2 = .42$; $F_i(1,299) = 23.95$, $p < .001$, $MSE = 4855.60$, $\eta^2 = .07$, with larger relatedness effects for degraded targets.

Percent correct

The accuracy data again yielded main effects of stimulus quality, $F_p(1,31) = 38.07$, $p < .001$, $MSE = .0012$, $\eta^2 = .55$; $F_i(1,299) = 35.36$, $p < .001$, $MSE = .012$, $\eta^2 = .11$, and relatedness, $F_p(1,31) = 33.59$, $p < .001$, $MSE = .00086$, $\eta^2 = .52$; $F_i(1,299) = 47.52$, $p < .001$, $MSE = .0057$, $\eta^2 = .14$. The interaction between stimulus quality and relatedness was also significant, $F_p(1,31) = 6.18$, $p = .019$, $MSE = .00056$, $\eta^2 = .17$; $F_i(1,299) = 6.78$, $p = .010$, $MSE = .0048$, $\eta^2 = .02$, with larger relatedness effects for degraded targets.

Table 6
Mean response latency, percent error rates, and ex-Gaussian parameters as a function of Target Degradation, and Prime-Target Relatedness for lexical decision performance in Experiment 5

	Mean	%Errors	μ	σ	τ
Clear targets					
Related	550	2.0	428	47	122
Unrelated	583	4.0	460	49	123
Effect	33	2.0	32	2	1
Degraded targets					
Related	638	4.7	479	57	159
Unrelated	709	8.8	529	48	181
Effect	71	4.1	50	−9	22
Interaction	38	2.1	18	−11	21

Ex-Gaussian analyses

For μ, the main effects of stimulus quality, $F(1,31) = 192.42$, $p < .001$, $MSE = 601.12$, $\eta^2 = .86$, and relatedness, $F(1,31) = 46.10$, $p < .001$, $MSE = 1154.10$, $\eta^2 = .60$, were significant. The stimulus quality × relatedness interaction was also significant, $F(1,31) = 4.88$, $p = .035$, $MSE = 536.80$, $\eta^2 = .14$, reflecting larger relatedness effects in μ for degraded targets than clear targets. None of the effects were significant for σ. Turning to τ, the main effect of stimulus quality was significant, $F(1,31) = 26.61$, $p < .001$, $MSE = 2664.24$, $\eta^2 = .46$. The main effect of relatedness was not significant ($p = .11$), but the stimulus quality × relatedness interaction approached significance, $F(1,31) = 3.48$, $p = .072$, $MSE = 957.55$, $\eta^2 = .10$, with larger relatedness effects for degraded targets. Separate analyses of the degraded and clear targets indicated that there was a reliable relatedness effect in τ for the degraded targets, $t(31) = 1.93$, $p < .05$, one tailed, but not for the clear targets, $t < 1$. Table 6 shows that the semantic relatedness effect for clear words is mediated primarily by μ (distributional shifting), whereas the stimulus quality × relatedness interaction is mediated by a mixture of μ and τ (distributional shifting and skewing). This converges nicely with the pronunciation data from Experiment 4 (see Table 5).

Vincentile analysis

The mean Vincentiles are plotted in Fig. 11, along with the best fitting ex-Gaussian function. As shown in Fig. 11, the semantic relatedness effects for clear targets were again primarily mediated by distributional shifting, whereas the relatedness effect for the degraded targets increases and appears to be relatively invariant across the early to middle Vincentiles, but then grows steadily at the later Vincentiles. This is consistent with the significant stimulus quality × relatedness × Vincentile interaction, $F(2,59) = 3.70$, $p = .033$, $MSE = 2543.50$, $\eta^2 = .11$, which indicates that the shape of the semantic relatedness effect differs for clear and degraded targets. Indeed, separate analyses on the clear and degraded target conditions indicated that there was no evidence of a Vincentile by relatedness interaction for the clear conditions, $F < 1$, replicating the long SOA condition from Experiment 2, but a reliable interaction for the degraded conditions, $F(2,47) = 3.99$, $p = .036$, $MSE = 6000.09$, $\eta^2 = .11$. This pattern of results is identical with the interactive effects found in Experiment 4 for the pronunciation task.

The results from both Experiments 4 and 5 do not suggest that stimulus degradation produces a simple multiplicative influence on semantic priming at the level of the underlying RT distributions. Specifically, because the long SOA data from Experiments 1 and 2 indicate that priming primarily produces distributional shifting, a simple multiplicative model would simply predict greater distributional shifting in the degraded condition (see Fig. 9). However, this is clearly not the pattern obtained. Specifically, we replicated the distributional shifting in the clear conditions, but in the degraded conditions, there was evidence of both a multiplicative influence at the short to medium Vincentiles, but for very slow RTs, there was clear evidence of exaggerated relat-

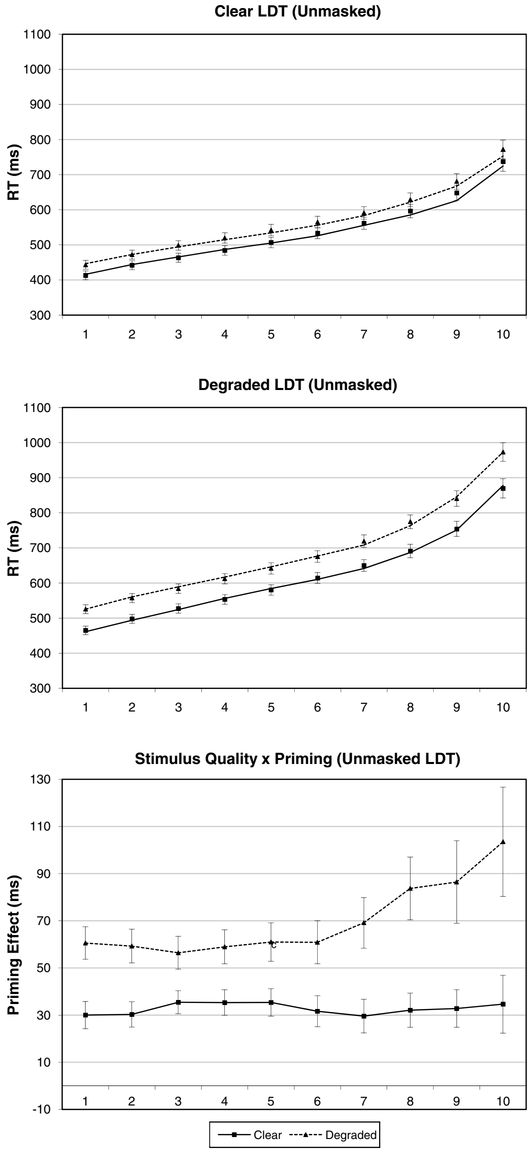

Fig. 11. Lexical decision performance from Experiment 5 as a function of prime relatedness and Vincentiles in the clear target (top panel) and degraded target (middle panel) conditions, along with the priming effects as a function of Vincentiles (bottom panel). In the top and middle panels, participants' mean Vincentiles (■ = related, ▲ = unrelated) are represented by data points and standard error bars. Best-fitting ex-Gaussian Vincentiles are represented by lines (solid line = related, dashed line = unrelated).

edness effects. Moreover, the pattern observed for the degraded conditions is more typical of the influence of other variables on RT distributions in the visual word recognition literature, as shown earlier in Fig. 5.

Why might the relatively slow components of the RT distribution benefit considerably more from semantic priming? One possibility is that participants may use a more controlled retrieval process when targets become particularly difficult to resolve, i.e., the targets that are in the slow end of the RT distribution. The notion here is that this would be quite distinct from the standard forward priming influence that may produce the shift in the RT distribution. If this were indeed the case, then one might eliminate the exaggerated influence of prime type at the slow end of the RT distribution when the primes are no longer available for explicit retrieval. Hence, Experiments 6 and 7 are a replication of Experiments 4 and 5 with highly masked primes. The prediction is that if the primes are unavailable for conscious processing, one may eliminate the exaggerated influence of priming at the tail of the distribution, thereby leaving only the forward influence of the prime. If this is indeed the case, then one might not expect the relatedness by stimulus quality interaction when primes are unavailable for conscious processing.

Experiment 6: Effects of relatedness and stimulus quality in pronunciation performance with highly masked primes

Method

Participants

Forty undergraduates participated in Experiment 6.

Apparatus, stimuli, procedure, and design

Same as Experiment 4 except that primes were masked. Each trial consisted of the following order of events: (a) a blank screen for 500 ms, (b) a forward mask of hashes (e.g., ######) that is length-matched to the prime for 56 ms, (c) a prime for 42 ms, and (d) the target stimulus at the fixation location, which served as the backward mask.

Results and discussion

Errors (3.8% across both conditions) and response latencies faster than 200 ms or slower than 3000 ms were first excluded from the analyses. A further 3.4% of the responses were identified as outliers and were excluded. The mean RT, accuracy, and ex-Gaussian parameters are displayed in Table 7.

Table 7
Mean response latency, percent error rates, and ex-Gaussian parameters as a function of Target Degradation, and Prime-Target Relatedness for pronunciation performance with masked primes in Experiment 6

	Mean	%Errors	μ	σ	τ
Clear targets					
Related	584	1.8	529	84	56
Unrelated	584	1.5	534	85	50
Effect	0	−.3	5	1	−6
Degraded targets					
Related	670	5.8	558	93	111
Unrelated	699	6.1	583	100	117
Effect	29	0.3	25	7	6
Interaction	29	0.6	20	6	12

Response latencies

For mean response latencies, the main effects of stimulus quality, $F_p(1,39) = 172.82$, $p < .001$, $MSE = 2330.07$, $\eta^2 = .82$; $F_i(1,299) = 739.51$, $p < .001$, $MSE = 4304.52$, $\eta^2 = .71$, relatedness, $F_p(1,39) = 17.60$, $p < .001$, $MSE = 507.98$, $\eta^2 = .31$; $F_i(1,299) = 14.29$, $p < .001$, $MSE = 3944.19$, $\eta^2 = .05$, and the stimulus quality × relatedness interaction were significant, $F_p(1,39) = 10.58$, $p = .002$, $MSE = 826.63$, $\eta^2 = .21$; $F_i(1,299) = 13.02$, $p < .001$, $MSE = 4856.77$, $\eta^2 = .04$. It is particularly noteworthy that, as shown in Table 7, there was absolutely no evidence of relatedness (0 ms) under the masked conditions with clear targets; however, for degraded targets, there was clear evidence for reliable relatedness effects, $t(39) = 3.90$, $p < .001$.

Percent correct

The accuracy data only yielded a main effect of stimulus quality, $F_p(1,39) = 21.40$, $p < .001$, $MSE = .0034$, $\eta^2 = .35$; $F_i(1,299) = 114.68$, $p < .001$, $MSE = .0048$, $\eta^2 = .28$.

Ex-Gaussian analyses

For μ, the main effects of stimulus quality, $F(1,39) = 33.29$, $p < .001$, $MSE = 1830.09$, $\eta^2 = .46$, and relatedness, $F(1,39) = 8.04$, $p = .007$, $MSE = 1103.99$, $\eta^2 = .17$, were significant. The stimulus quality × relatedness interaction approached significance, $F(1,39) = 3.04$, $p = .089$, $MSE = 1113.98$, $\eta^2 = .07$, with the relatedness effect only reaching significance for degraded targets, $t(39) = 2.79$, $p = .008$, and not approaching significance for clear targets, $t < 1$. Only the main effect of stimulus quality reached significance for σ, $F(1,39) = 6.49$, $p = .015$, $MSE = 889.51$, $\eta^2 = .14$. Turning to τ, the main effect of stimulus quality was again significant, $F(1,39) = 35.39$, $p < .001$, $MSE = 4247.39$, $\eta^2 = .48$. Although neither σ nor τ produced a reliable stimulus quality by relatedness interaction, as shown in Table 7, both components produced an increasing relatedness effect in the degraded condition, which is consistent with the Vincentile analyses below.

In sum, there is no masked semantic priming for clear targets, and the masked priming effect for degraded words is reflected predominantly in a change in μ, and to a lesser extent, by a change in σ and in τ.

Vincentile analysis

The mean Vincentiles are plotted in Fig. 12, along with the best fitting ex-Gaussian function. As shown, the masked priming effect for clear targets is absent across Vincentiles. For degraded targets, the masked priming effect is smallest at the fastest Vincentiles and grows steadily as the Vincentiles become slower. This increase across Vincentiles is consistent with the numerical increases in σ and τ, mentioned above. Importantly, the stimulus quality × relatedness × Vincentiles interaction was reliable, $F(1,55) = 3.84$, $p = .042$, $MSE = 8508.48$, $\eta^2 = .09$. Further exploration of this interaction indicated that the Vincentile by relatedness interaction did not approach significance for the clear targets, $F < 1$, but was reliable for the degraded targets, $F(1,54) = 3.80$, $p = .044$, $MSE = 13238.01$, $\eta^2 = .09$. Hence, masked priming does not simply shift the RT distribution when primes are degraded but it appears to increase across the Vincentiles, as in the previous experiments with stimulus degradation.

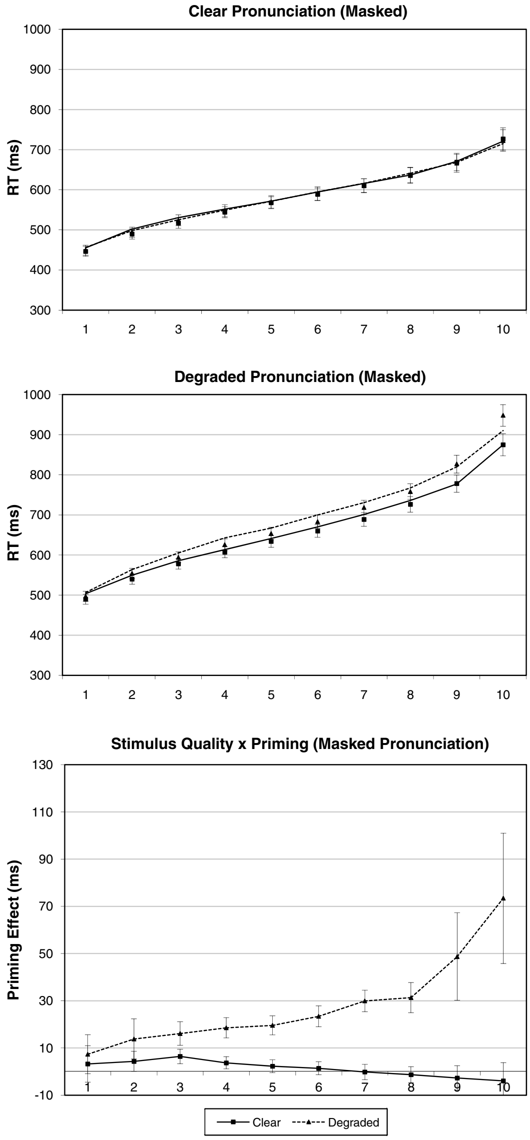

Fig. 12. Pronunciation performance from Experiment 6 as a function of masked prime relatedness and Vincentiles in the clear target (top panel) and degraded target (middle panel) conditions, along with the priming effects as a function of Vincentiles (bottom panel). In the top and middle panels, participants' mean Vincentiles (■ = related, ▲ = unrelated) are represented by data points and standard error bars. Best-fitting ex-Gaussian Vincentiles are represented by lines (solid line = related, dashed line = unrelated).

Experiment 7: Effects of relatedness and stimulus quality in lexical decision performance with highly masked primes

Method

Participants

Thirty-two undergraduates participated in Experiment 7.

Apparatus, stimuli, procedure, and design

Identical to Experiment 5, but using the masking priming conditions described in the Method section for Experiment 6.

Results and discussion

Errors (7.2% across both conditions) and response latencies faster than 200 ms or slower than 3000 ms were first excluded from the analyses. Using the trimming criteria described earlier, a further 2.9% of the responses were excluded. The mean RT, accuracy, and ex-Gaussian parameters are displayed in Table 8.

Response latencies

For mean response latencies, the main effects of stimulus quality, $F_p(1,31) = 78.10$, $p < .001$, $MSE = 7121.60$, $\eta^2 = .72$; $F_i(1,299) = 1105.42$, $p < .001$, $MSE = 4852.90$, $\eta^2 = .79$, relatedness, $F_p(1,31) = 27.84$, $p < .001$, $MSE = 505.40$, $\eta^2 = .47$; $F_i(1,299) = 12.97$, $p < .001$, $MSE = 11721.26$, $\eta^2 = .04$, and the stimulus quality × relatedness interaction were significant, $F_p(1,31) = 4.89$, $p = .034$, $MSE = 471.18$, $\eta^2 = .14$; $F_i(1,299) = 5.59$, $p = .019$, $MSE = 5260.04$, $\eta^2 = .02$. As shown in Table 6, there were larger masked priming effects for degraded targets compared to clear targets. Although

Table 8
Mean response latency, percent error rates, and ex-Gaussian parameters as a function of Target Degradation, and Prime-Target Relatedness for lexical decision performance with masked primes in Experiment 7

	Mean	%Errors	μ	σ	τ
Clear targets					
Related	562	4.1	460	42	103
Unrelated	575	5.4	469	45	106
Effect	13	1.3	9	3	3
Degraded targets					
Related	686	7.3	519	47	167
Unrelated	715	8.7	530	51	185
Effect	29	1.4	11	4	18
Interaction	16	0.1	2	1	15

the effects were larger for degraded than clear targets, in contrast to the pronunciation results of Experiment 6, there was evidence in lexical decision of reliable masked priming even for the clear targets, $t(31) = 3.00$, $p = .005$.

Percent correct

The accuracy data yielded main effects of stimulus quality, $F_p(1,31) = 35.47$, $p < .001$, $MSE = .00094$, $\eta^2 = .53$; $F_i(1,299) = 42.69$, $p < .001$, $MSE = .0073$, $\eta^2 = .13$, and relatedness, $F_p(1,31) = 9.59$, $p = .004$, $MSE = .00061$, $\eta^2 = .24$; $F_i(1,299) = 7.38$, $p = .007$, $MSE = .0075$, $\eta^2 = .02$. The interaction was not significant by participants or by items, F_p and $F_i < 1$.

Ex-Gaussian analyses

For μ, the main effects of stimulus quality, $F(1,31) = 123.54$, $p < .001$, $MSE = 937.94$, $\eta^2 = .80$, and relatedness, $F(1,31) = 8.63$, $p = .006$, $MSE = 362.64$, $\eta^2 = .22$, were significant, but the stimulus quality × relatedness interaction was not significant, $F < 1$. In contrast, none of the effects were significant for σ. Turning to τ, the main effects of stimulus quality, $F(1,31) = 41.95$, $p < .001$, $MSE = 3917.03$, $\eta^2 = .58$, and relatedness, $F(1,31) = 5.25$, $p = .029$, $MSE = 748.69$, $\eta^2 = .15$, were significant. The stimulus quality × relatedness interaction was also significant, $F(1,31) = 4.50$, $p = .042$, $MSE = 423.54$, $\eta^2 = .13$, with larger relatedness effects in the degraded condition. Table 6 shows that the semantic priming effect for clear words is mediated primarily by μ (distributional shifting), while the stimulus quality × relatedness interaction is mediated by a mixture of μ and τ (distributional shifting and skewing).

Vincentile analysis

The mean Vincentiles are displayed in Fig. 13, along with the best fitting ex-Gaussian functions. Masked semantic priming effects for clear targets were smaller at the fast Vincentiles, and increased slightly towards the slowest Vincentiles. For degraded targets, the masked priming effect is smallest at the fastest Vincentiles and grows much more strongly across Vincentiles. Although the stimulus quality × relatedness × Vincentiles interaction was not significant, $p = .128$, separate

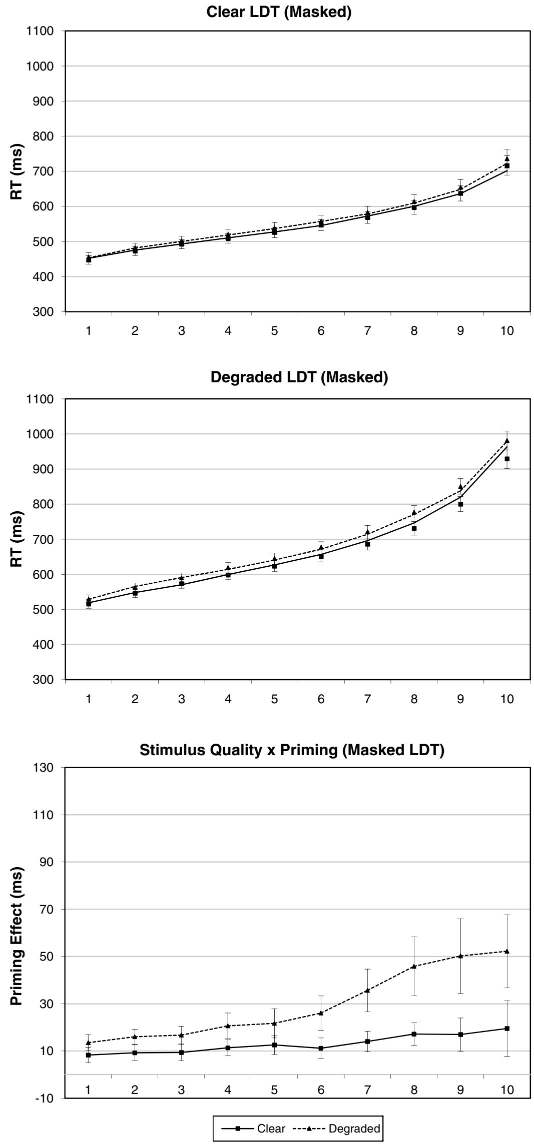

Fig. 13. Lexical decision performance from Experiment 7 as a function of masked prime relatedness and Vincentiles in the clear target (top panel) and degraded target (middle panel) conditions, along with the priming effects as a function of Vincentiles (bottom panel). In the top and middle panels, participants' mean Vincentiles (■ = related, ▲ = unrelated) are represented by data points and standard error bars. Best-fitting ex-Gaussian Vincentiles are represented by lines (solid line = -related, dashed line = unrelated).

analyses of the clear and degraded target conditions yielded a reliable interaction between Vincentile and masked priming for the degraded targets, $F(2,64) = 4.84$, $p = .010$, $MSE = 3219.77$, $\eta^2 = .14$, but not for the clear targets, $F < 1$. This pattern indicates that the shape of the relatedness effect is modulated by the stimulus quality of the target.

The results from Experiment 6 and 7 are clear. In both pronunciation and lexical decision performance, there was clear evidence for a priming by degradation interaction even when primes were highly masked and unavailable for explicit conscious processing. Hence, the pattern observed in the RT distributions, reflecting the degradation by relatedness interaction in Experiments 4 and 5, cannot be attributed to explicit retrieval of a clearly presented prime. Of course, one can always question the quality of the masking. However, because the clear conditions in the pronunciation performance of Experiment 6 provided no evidence of masked priming, the masking was sufficient to eliminate the strong effects from these same primes that were observed in Experiments 1 and 4. Interestingly, there was also a change in the overall magnitude of the masked priming effects as a function of target degradation. Specifically, in both Experiments 6 and 7, there was evidence that the masked priming effects increased when the target was degraded. Hence, instead of decreasing reliance on highly masked primes under target degraded conditions, it appears that participants *increase* their reliance. This is important because it suggests that the threshold at which there is an influence of the prime is a reflection of the utility of the prime information, as suggested by Whittlesea and Jacoby (1990) and Bodner and Masson (2001), and described further below.

General discussion

The goal of the present experiments was to demonstrate the power of RT distributional analyses to develop a better understanding of a fundamental finding in the psycholinguistics literature, the semantic priming effect. If semantic priming behaved as other standard variables in the visual word recognition literature, one would expect an increasing effect of the variable across the RT distribution. Such a pattern would appear to be most compatible with predictions from computational models of priming (e.g., the compound cue model of Ratcliff & McKoon, 1988). However, this was not the pattern observed in the present series of experiments. Specifically, when the target was presented in a clear non-degraded fashion, distributional shifting was observed in the results of Experiment 1 (short SOA and long SOA), Experiment 2 (long SOA), and Experiments 4, 5, and 7. The only conditions where distributional shifting did *not* occur for clearly presented targets was when the prime was presented for a short SOA, and the task was lexical decision (Experiment 2 and replicated in Experiment 3). In Experiment 6, where masked primed pronunciation was used, the clear target conditions were uninformative, since there was no evidence of priming in these conditions.

In contrast to the nature of the semantic priming effect under clear target conditions, a markedly different pattern was observed when the target was degraded. Specifically, in Experiments 4, 5, 6, and 7, there was evidence of an increased influence of the prime across the RT distribution. This qualitative change in the priming effects (i.e., from a shift with clear targets to an increase across the RT distribution for degraded targets) would appear to be most consistent with qualitatively different processes being engaged for clear versus degraded targets. Interestingly, this general pattern occurred both when the primes were clearly available (Experiments 4 and 5) and when the primes were highly masked (Experiments 6 and 7).

In considering the implications of the present results, we will focus on two general issues: First, we will discuss the utility of RT distributional analyses, noting some limitations in the present approach. Second, we discuss the theoretical implications of the present results for models of semantic priming.

RT distributional analyses

In the current paper, we have attempted to demonstrate the utility of RT distributional analyses for better understanding the influence of variables in response latency studies. We have utilized two different techniques, ex-Gaussian analyses and Vincentile analyses, to capture the influence of variables. We believe the convergence of the techniques is particularly helpful. However, it is also important to note limitations of this approach, and discuss possible alternative approaches.

Are RT distributions sufficiently stable to make strong inferences?

As one increases the power of a measurement device, it is possible that one may be measuring noise. Returning to the microscope metaphor, one may be looking at a fleck of dust instead of a targeted cellular component. So, how stable are the RT distributional patterns? The present results highlight the stability of such measures. As noted above, the shift in distributions for clear targets as a function of prime relatedness and the increase in the priming effect across the distribution for degraded targets is a very consistent pattern in the present results. Moreover, as noted in the Introduction (see Fig. 5), the available literature also seems quite consistent regarding the influence of word frequency, lexicality, and degradation across experiments, both within and across laboratories. Indeed, there are now replicable tradeoffs in the ex-Gaussian parameters that mask an effect in means in both Stroop performance (e.g., Heathcote et al., 1991; Spieler et al., 1996), and in visual word recognition (see Yap, Balota, Tse, & Besner, in press).

Another question regarding stability concerns how well distributional effects generalize to new sets of stimuli. This would appear to be a more fundamental question, because inherent item differences would place items at different points within the RT distribution. For example, consider the word frequency effect. One typically finds an increase in the word frequency effect across the RT distribution, with disproportionate effects at the tail of the distribution in the LDT (see Andrews & Heathcote, 2001; Balota & Spieler, 1999; Yap & Balota, 2007). However, one may be concerned that these distributional changes may reflect the range of frequencies within the high-frequency and low-frequency words selected for comparisons. Specifically, because there is a systematic relationship between log frequency and response latency in the LDT, when one manipulates word frequency, one might change the nature of the influence of word frequency on the RT distributions by explicitly manipulating the frequency ranges within the band of high- and low-frequency words. So, when considering RT distributions, it is important to be cautious about the ability to replicate across different stimulus sets. With respect to some variables, such as lexicality and word frequency, the effects appear to be quite stable across different stimulus sets.

In the present study, although we have sampled from the same large set of prime-target trials across experiments, because of counterbalancing procedures, different subsets of items were selected across participants and experiments, and yet they produced similar patterns of results. Although this does not eliminate the problem of item selection influences on the RT distributions, it does minimize it. Such item influences on RT distributions are further minimized in other domains of research such as attentional selection work (e.g., Stroop performance), wherein a relatively small set of stimuli is repeated across trials.

In RT distributional analyses, one also needs to consider how practice and fatigue effects may modulate the mixture of components in the distributions across time (see Cousineau, Brown, & Heathcote, 2004, for a discussion). For example, the RTs in the slow tail of the RT distribution may reflect earlier trials, where participants have not yet fully adapted to the task requirements. The RTs in the faster bins may be more reflective of trials after participants have adapted to the tasks. This is indeed an important issue that also needs to be explored with larger datasets.

So, what is the stability of the RT distributional estimates across trials? We have recently explored this question in a large data set in which each subject produced responses to over 2400 monosyllabic words (see Balota, Cortese, Sergent-Marshall, Spieler, & Yap, 2004). We separated the data into two halves of lexical decision and pronunciation data, most of which were collected on separate testing sessions between one day and one week apart. We then computed partial correlations between the parameter estimates taken from the two halves. Importantly, we partialed out overall mean response latency for each participant to insure that such relationships were not simply a reflection of general slowing (i.e., a relationship between overall mean performance and standard deviations, see Faust, Balota, Spieler, & Ferraro, 1999). The correlations of the parameter estimates were surprisingly high, supporting the stability of these estimates. Specifically, the partial correlations between the testing sessions for μ, σ, and τ for lexical decision were .98, .83, and .98, and in pronunciation were .60, .61, and .91, respectively. The relatively smaller correlations in pronunciation may be due to two factors. First, the RT distribution for pronunciation performance is more Gaussian, and hence, by partialing out overall RT, one is simply partialing out the natural relationship between means and *SD*s, reflected in large part by differences in μ and σ. Second, it is possible that the decrease in correlation is due to the large idiosyncratic influence of voice onsets (over 30% of the variance in these data) that are likely to vary across different sets of stimuli across the two sessions.

Finally, regarding the stability of RT distributional estimates, it is also interesting to note that Schmiedek, Oberauer, Wilhelm, Süß, and Wittmann (2007) have recently found common latent factors related to the three ex-Gaussian parameters in a set of RT tasks. Importantly, Schmiedek et al. found τ (along with a measure of drift rate in the diffusion model) was the strongest predictor of working memory, reasoning, and psychometric speed. It is indeed quite interesting that the τ parameter can be used as a reliable individual difference marker that is related to control measures, further supporting its potential utility as a useful descriptive measure.

Why (or why not) the ex-Gaussian?

As noted in the Introduction, there are many formal models of RT distributions that have been used to capture the characteristics of empirical RT distributions, including the Weibull, ex-Gaussian, Gamma, Poisson Race, Wald, and ex-Wald models (see Luce, 1986; Van Zandt, 2002, for excellent reviews). In large part, we have chosen to fit our empirical results to the ex-Gaussian model in order to make direct connections with the mean-dominated extant literature, since the sum of two of the three parameters approximates the mean of the distribution. We believe that this is a useful intermediate step for making connections with available literature, where factorial designs and ANOVAs on means are the typical levels of analyses. However, there are clearly alternative models that may be preferable in some cases. For example, the Weibull function nicely affords estimates of location, scale, and shape of the RT distribution, and some have argued (e.g., Rouder, Tuerlinckx, Speckman, Lu, & Gomez, submitted for publication) that this is a more appropriate function for capturing how variables change an empirical distribution.

Ultimately, the best approach may be to use a model that more directly maps onto models of response latency. For example, Schwarz (2001) has argued that the ex-Wald (which is the convolution of an exponential and a Wald distribution) may have this advantage. For example, Schwarz argues that there is good evidence for an additive exponential RT component of RT distributions (reflecting the "ex"), and that the Wald distribution is useful because it maps onto a diffusion process to reach a fixed criterion, reflected by two parameters of the Wald. Hence, instead of simply being a descriptive model of the RT distribution, the ex-Wald makes connections with a general process model of response latency. Indeed Schwarz nicely demonstrates how the model can be applied to a go/no-go RT study of mental digit comparison, showing that numerical distance and prior probability of a go response have isolated effects on the parameters of the model. It will be particularly useful in the future to (a) extend this approach to additional tasks with well-studied variables (e.g., word frequency effects in lexical decision performance), and (b) consider the implications of this approach for tasks that do not include a single or binary response criterion, such as speeded pronunciation.

Although there are many models of RT distributions available, Cousineau et al. (2004) have recently argued that there is no consensus regarding which model one should choose. However, Cousineau et al. also point out that there are some guiding principles for considering an appropriate model for distributional analyses. First, one should consider the informative utility across conditions as an important metric. Do the parameters from a model capture the manipulations in a useful manner? It is unclear if a model of RT distributions is providing any new information if all parameters are consistently changing across manipulations. The clear dissociation observed in the present study between semantic priming effects in clear and degraded conditions provides a nice example of specificity in the parameters. Second, of course, the model should be parsimonious in the number of parameters that are used to describe the RT distribution. The relatively simple three-parameter ex-Gaussian would appear efficient by this measure (see Myung, 2000, for comparing models with varying number of parameters). Ultimately, we agree that it is most useful to map the characteristics of empirical distributions onto a specific model of response latencies within a given task. However, it may indeed be the case that there will be no general model of response latencies, and that models will need to specifically capture characteristics within a given task. In this light, we believe the ex-Gaussian is a useful descriptive model that has sufficient generality to make connections with predictions from the available models (see, for example, Balota & Spieler, 1999; Ratcliff, 1978). Finally, it is also important to remember to provide converging evidence from Vincentile (or Quantile) analyses to provide direct links between the fits from the model and the empirical distributions.

Semantic priming

The present results indicate that when targets are clearly presented, and participants have sufficient time to process the prime, the influence of semantic priming is primarily to shift the RT distribution. This pattern is intriguing because it appears to suggest that the prime information affords a headstart on target processing, and is inconsistent with the effects of most other variables in the word recognition literature. Importantly, such a simple shift would appear to be inconsistent with simple models of how response latency should change as a function of priming. Specifically, one might expect that the distribution of priming effects should reflect the convolution of two distributions, one which reflects the difficulty of the target and a second which reflects the differences across items in prime-target associative strength. At the very least, the convolution of two such distributions should reflect both a change in means and in variances, but the current results appear to primarily provide evidence for a change in means, as if each of the words in the unrelated condition received a speedup of N ms by the presence of a related prime. In general, shifts in RT distributions, without changes in variance, are quite intriguing within models of RT.

Of course, a simple headstart model is indeed consistent with classic metaphorical models of lexical processing and priming. For example, within Morton's (1969) classic logogen model, if the primes produced a constant amount of pre-activation for the logogen, one might expect a simple shift. The serial search model of Becker (1980) would appear to predict such a shift, since the prime would have the effect of restricting the search set for the target. Also, if one extends the entry-opening account of masked repetition priming (see Forster et al., 2003) to semantic priming, one would also predict a distributional shift, since the influence of the prime in this model affords a headstart on target processing. However, such a pattern does not seem to easily fall from current computational models (e.g., the compound cue model or the feature overlap model), wherein the influence of the prime is on a diffusion process or on the settling rate. These models do *not* predict a simple shift in the RT distribution, but an increase in the scale of the RT distribution which would produce an increase in the additional parameters. Of course, the advantage of metaphorical headstart frameworks in accommodating the present shift results largely reflects the flexibility of such descriptions, which are unencumbered by the necessary machinery used to predict trial-by-trial response latencies for a given task. Hence, until a detailed implementation is available that can capture the shape of RT distributions for specific tasks, one needs to be cautious about accepting the headstart account.

The influence of stimulus degradation is of particular interest here. Specifically, when targets are degraded, the influence of relatedness increases the priming effect on the relatively fast RTs, and produces an increasing effect across the RT distribution (see however, Brown & Besner, 2002, for evidence of additive effects of degradation and masked priming in the LDT). One might argue here that when targets are degraded, the RT distributions better reflect the predictions from current computational models, i.e., producing larger priming effects for the items in the tail of the distribution. One way to consider the implications of these results is that for highly fluent lexical processors (i.e., Washington University students), and under conditions of clear target presentation, there is relatively little need to retrieve the prime information, and so the priming effect may be reflective of a relatively modular lexical processing system, thereby producing the simple shift. When the target is degraded, however, the system uses any available information available to better resolve the target, and hence, one finds the expected increase in effect size across Vincentiles. Consequently, the more difficult items, i.e., those at the slowest Vincentiles, will be associated with more reliance on the prime information.

Another way of considering these results is within the Plaut and Booth (2000) single-mechanism connectionist account of semantic priming. Because this model assumes a non-linear logistic activation function relating input activation to behavioral output, different manipulations may produce additive or interactive effects depending on where the effects are located on the logistic function. Within this framework, one would need to argue that for highly skilled readers processing clear targets, one is in the relatively gradual portion of the logistic function, whereas, when one degrades the target, one moves to a steeper, more sensitive area of the activation function, thereby producing larger priming effects for the slower items. However, consider what happens were we to assume that the clear targets are at the steep linear portion of the activation function. In this case, one might actually predict *smaller* priming effects for the degraded targets (associated with lower input activation), which is of course *not* what is typically found. Thus, although such a logistic activation function has the potential flexibility to accommodate the present results, it will be particularly important to make *specific* predictions that are constrained by the variance in reading skill, item difficulty and variability, size of the degradation manipulation, and variability in associative strength to accommodate distributional results (see Besner & Borowsky, 2006; Besner, Wartak, & Robidoux, in press; Plaut & Booth, 2006, for a discussion of some potential difficulties when such specific constraints are assumed).

Masked priming

Turning to the results from the masked priming experiments, there are some important implications. Specifically, in both lexical decision and speeded pronunciation, there were larger masked priming effects when the target was degraded compared to when the target was presented in a clear fashion. This ran counter to our expectation that the slower tail of the RT distribution under degraded conditions might reflect more of a prime retrieval process when the prime was consciously available. Of course, it is quite possible that the prime masking procedure used in the present study was not sufficiently powerful to minimize conscious processing of the prime. Hence, one might argue that the results simply reflect the standard priming by degradation interaction that is well established in the literature and is replicated in the present Experiments 4 and 5. However, the results from Experiment 6 are again informative here. Specifically, there was no evidence of a priming effect when the targets were presented in a clear fashion and the primes were masked. Hence, the mask was indeed sufficiently effective to eliminate any evidence for priming in this experiment. However, on trials when the target was degraded, the masked priming effect returned in this experiment.

Why does target degradation increase priming effects under conditions of highly masked semantic primes? There is already some evidence in the literature that addresses this issue, albeit from a different paradigm. Specifically, Whittlesea and Jacoby (1990), using a three-event priming paradigm, found that the speed to name a third target stimulus (e.g., GREEN) was influenced more by a 60 ms repetition prime (e.g., GREEN) if a 150 ms interpolated word was degraded (pLaNt), compared to when it was non-degraded (e.g., PLANT). Whittlesea and Jacoby argued that degrading the interpolated word increased the emphasis on retrieving the first masked word, thereby producing greater repetition priming for the third word, the target. There was little direct evidence in the Whittlesea and Jacoby study regarding how strongly the primes were masked and so there may have been some leakage of conscious processing of the primes. Moreover, it is unclear if the implications from the three-prime mixture of repetition and semantic priming would extend to the more standard priming paradigm. However, the interesting inference from this study is that under masked prime conditions, participants appear to rely more on prime information when there is a degraded, compared to clear, interpolated target.

More recently, the notion that participants rely on episodic retrieval in masked priming situations has been further advocated by Bodner and Masson (1997, 2001, 2003, 2004; Masson & Bodner 2003) in their memory recruitment account. For example, in one study, Bodner and Masson (2001) varied the proportion of masked repetition primes in a LDT from .2 to .8. The relatedness proportion in the semantic priming domain has been typically viewed as reflecting controlled processing (e.g., den Heyer et al., 1983). However, Bodner and Masson found that one obtains a relatedness proportion effect even when the primes are highly masked. Bodner and Masson (2003) have also extended the relatedness proportion effect to masked semantic priming paradigms. They have argued from these results that the cognitive system is sensitive to the utility of the prime information across trials. Because repetition or related primes help resolve the targets, it is useful to retrieve the prime information. In the .8 relatedness proportion condition, participants are more likely to recruit the masked primes in helping to identify the targets then in the .2 proportion condition because the payoff for such recruitment is higher. As Bodner, Masson, and Richard (2006) have recently argued, this would be quite consistent with Anderson and Milson's (1989) rational analysis of memory account for semantic priming effects.

The present results did not include a relatedness proportion manipulation. Rather, we increased the utility of the prime information by degrading the target. Hence, like Whittlesea and Jacoby (1990), we found that the influence of a masked prime increases when the target is degraded. Importantly, the RT distributional analyses indicated that the increased reliance on the masked prime was greatest for the more difficult items in the tail of the RT distribution, precisely as one might expect. This is yet another example of the useful converging evidence that distributional analyses afford.

The present results also have implications for threshold priming effects. In particular, these results question the notion that there is an absolute threshold that one must achieve to demonstrate unconscious priming. Such a threshold is typically demonstrated in some other task (e.g., forced choice prime identification or presence/absence prime detection), and so has been susceptible to a number of criticisms (see, for example, Holender, 1986). The present results indicate that at a level of masking in which there is no influence of the prime for clear targets, there is a strong influence of such primes for degraded targets. This pattern suggests that the system is sensitive to the utility of available information and flexibly recruits such information depending on the task demands (see Balota & Yap, 2006, for a recent discussion of a flexible lexical processor).

Conclusions

There are three major conclusions from the present results. First, when primes are clearly presented and targets are clearly presented, the effect of semantically related primes compared to unrelated primes is primarily to shift the RT distribution, instead of shifting and increasing the tail of the distribution, the pattern found with the vast majority of other variables in the visual word recognition domain. These priming effects were viewed as most consistent with a simple headstart metaphor of prime-to-target processing. Second, when the target is degraded, there is increased reliance on the prime information and this reliance produces both a shift and an increase in priming in the tail of the RT distribution. This pattern suggests that individuals increase their reliance on prime information when targets are degraded, particularly for the difficult items in the tail of the distribution. Third, even under highly masked prime conditions, degrading the target increases reliance on prime information. This has been viewed as being most consistent with an episodic retrieval account of masked priming.

Most importantly, the major goal of this paper was to demonstrate and evaluate the utility of RT distributional analyses for better understanding the influence of a standard manipulation in the psycholinguistic domain, i.e., semantic priming. Although there are clearly constraints in the interpretation of the influences of variables on RT distributions, and how best to measure such effects, we believe the accumulating evidence in the literature indicates that it is indeed time for researchers using chronometric methods to move beyond estimates of central tendency and to increase the magnification of their measurement tool.

Acknowledgments

Portions of this research were presented at the 47th (November, 2006, Houston) annual meeting of the Psychonomics Society. This research was supported by NIA Grant AGO3991 and NSF Grant BCS 0001801 to D.A. Balota. We thank Andrew Heathcote and Jeffrey Rouder for helpful discussions regarding RT distributional analyses.

Appendix A

		N	M_{related}	SE_{related}	$M_{\text{unrelated}}$	$SE_{\text{unrelated}}$	Effect
E1: Speeded pronunciation	Short SOA	8	739	151	605	125	134
	Long SOA	15	626	88	655	79	−29
E2: Lexical decision	Short SOA	16	654	46	637	35	17
	Long SOA	16	593	47	586	28	7
E3: LDT replication	Short SOA	15	566	28	573	18	−7
E4: Speeded pronunciation	Clear	4	898	212	534	69	364
	Degraded	13	691	86	754	80	−63
E5: Lexical decision	Clear	16	549	40	617	55	−68
	Degraded	29	808	54	836	61	−28
E6: Masked priming pronunciation	Clear	8	663	52	611	46	52
	Degraded	23	863	89	859	103	4
E7: Masked priming lexical decision	Clear	26	558	30	583	28	−25
	Degraded	29	691	44	764	44	73*

* $p < .10$.

References

Anderson, J. R., & Milson, R. (1989). Human memory: An adaptive perspective. *Psychological Review, 96*, 703–719.

Andrews, S., & Heathcote, A. (2001). Distinguishing common and task-specific processes in word identification: A matter of some moment. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 27*, 514–544.

Balota, D. A. (1983). Automatic semantic activation and episodic memory encoding. *Journal of Verbal Learning & Verbal Behavior, 22*, 88–104.

Balota, D. A., Black, S. R., & Cheney, M. (1992). Automatic and attentional priming in young and older adults: Reevaluation of the two-process model. *Journal of Experimental Psychology: Human Perception and Performance, 18*, 485–502.

Balota, D. A., Cortese, M. J., Sergent-Marshall, S. D., Spieler, D. H., & Yap, M. J. (2004). Visual word recognition of single-syllable words. *Journal of Experimental Psychology: General, 133*, 283–316.

Balota, D. A., & Lorch, R. F. (1986). Depth of automatic spreading activation: Mediated priming effects in pronunciation but not in lexical decision. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12*, 336–345.

Balota, D. A., & Paul, S. T. (1996). Summation of activation: Evidence from multiple primes that converge and diverge within semantic memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22*, 827–845.

Balota, D. A., & Spieler, D. H. (1999). Word frequency, repetition, and lexicality effects in word recognition tasks: Beyond measures of central tendency. *Journal of Experimental Psychology: General, 128*, 32–55.

Balota, D. A., & Yap, M. J. (2006). Attentional control and flexible lexical processing: Explorations of the magic moment of word recognition. In S. Andrews (Ed.), *From inkmarks to ideas: Current issues in lexical processing* (pp. 229–258). Hove, England: Psychology Press.

Becker, C. A. (1980). Semantic context effects in visual word recognition: An analysis of semantic strategies. *Memory & Cognition, 8*, 493–512.

Becker, C. A., & Killion, T. H. (1977). Interaction of visual and cognitive effects in word recognition. *Journal of Experimental Psychology: Human Perception and Performance, 3*, 389–401.

Besner, D., & Borowsky, R. (2006). Postscript: Plaut and Booth's (2006) new simulations–What have we learned? *Psychological Review, 113*, 194–195.

Besner, D., Wartak, S., & Robidoux, S. (in press). Constraints on computational accounts of basic processes in reading. *Journal of Experimental Psychology: Human Perception & Performance*.

Bodner, G. E., & Masson, M. E. J. (1997). Masked repetition priming of words and nonwords: Evidence for a nonlexical basis for priming. *Journal of Memory & Language, 37*, 268–293.

Bodner, G. E., & Masson, M. E. J. (2001). Prime validity affects masked repetition priming: Evidence for an episodic resource account of priming. *Journal of Memory & Language, 45*, 616–647.

Bodner, G. E., & Masson, M. E. J. (2003). Beyond spreading activation: An influence of relatedness proportion on masked semantic priming. *Psychonomic Bulletin & Review, 10*, 645–652.

Bodner, G. E., & Masson, M. E. J. (2004). Beyond binary judgments: Prime validity modulates masked repetition priming in the naming task. *Memory & Cognition, 32*, 1–11.

Bodner, G. E., Masson, M. E. J., & Richard, N. T. (2006). Repetition proportion biases masked priming of lexical decisions. *Memory & Cognition, 34*, 1298–1311.

Borowsky, R., & Besner, D. (1993). Visual word recognition: A multistage activation model. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19*, 813–840.

Borowsky, R., & Besner, D. (2006). Parallel distributed processing and lexical-semantic effects in visual word recognition. Are a few stages necessary? *Psychological Review, 113*, 181–195.

Brown, M., & Besner, D. (2002). Semantic priming: On the role of awareness in visual word recognition in the absence of an expectancy. *Consciousness and Cognition, 11*, 402–422.

Brown, S., & Heathcote, A. (2003). QMLE: Fast, robust, and efficient estimation of distribution functions based on quantiles. *Behavior Research Methods, Instruments, & Computers, 35*, 485–492.

Bub, D. N., Masson, M. E. J., & Lalonde, C. (2006). Cognitive control in children: Stroop interference and suppression of word reading. *Psychological Science, 17*, 351–357.

Burke, D. M., White, H., & Diaz, D. L. (1987). Semantic priming in young and older adults: Evidence for age constancy in automatic and attentional processes. *Journal of Experimental Psychology: Human Perception and Performance, 13*, 79–88.

Cousineau, D., Brown, S., & Heathcote, A. (2004). Fitting distributions using maximum likelihood: Methods and packages. *Behavior Research Methods, Instruments, & Computers, 36*, 742–756.

den Heyer, K., Briand, K., & Dannenbring, G. L. (1983). Strategic factors in a lexical-decision task: Evidence for automatic and attention-driven processes. *Memory & Cognition, 11*, 374–381.

Faust, M. E., Balota, D. A., Spieler, D. H., & Ferraro, F. R. (1999). Individual differences in information-processing rate and amount: Implications for group differences in response latency. *Psychological Bulletin, 125*, 777–799.

Favreau, M., & Segalowitz, N. S. (1983). Automatic and controlled processes in the first- and second-language reading of fluent bilinguals. *Memory & Cognition, 11*, 565–574.

Forster, K. I., Mohan, K., & Hector, J. (2003). The mechanics of masked priming. In S. Kinoshita & S. J. Lupker (Eds.), *Masked priming: State of the art* (pp. 3–37). Hove, UK: Psychology Press.

Gold, B. T., Balota, D. A., Jones, S. J., Powell, D. K., Smith, C. D., & Andersen, A. H. (2006). Dissociation of automatic and strategic lexical-semantics: Functional magnetic resonance imaging evidence for differing roles of multiple frontotemporal regions. *Journal of Neuroscience, 26*, 6523–6532.

Heathcote, A., Brown, S., & Cousineau, D. (2004). QMPE: Estimating Lognormal, Wald and Weibull RT distributions with a parameter dependent lower bound. *Behavior Research Methods, Instruments, & Computers, 36*, 277–290.

Heathcote, A., Brown, S., & Mewhort, D. J. K. (2002). Quantile maximum likelihood estimation of response time distributions. *Psychonomic Bulletin & Review, 9*, 394–401.

Heathcote, A., Popiel, S. J., & Mewhort, D. J. K. (1991). Analysis of response time distributions: An example using the Stroop task. *Psychological Bulletin, 109*, 340–347.

Holender, D. (1986). Semantic activation without conscious identification in dichotic listening, parafoveal vision, and visual masking: A survey and appraisal. *Behavioral & Brain Sciences, 9*, 1–66.

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

Jones, M. N., Kintsch, W., & Mewhort, D. J. K. (2006). High-dimensional semantic space accounts of priming. *Journal of Memory & Language, 55*, 534–552.

Jonides, J., & Mack, R. (1984). On the cost and benefit of cost and benefit. *Psychological Bulletin, 96*, 29–44.

Luce, R. D. (1986). *Response times: Their role in inferring elementary mental organization*. New York: Oxford University Press.

Lund, K., & Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods, Instruments & Computers, 28*, 203–208.

Martin, A. (2005). Functional neuroimaging of semantic memory. In R. Cabeza & A. Kingstone (Eds.), *Handbook of functional neuroimaging of cognition* (pp. 153–186). Cambridge, MA: MIT Press.

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

Masson, M. E. J., & Bodner, G. E. (2003). A retrospective view of masked priming: Toward a unified account of masked and long-term repetition priming. In S. Kinoshita & S. Lupker (Eds.), *Masked priming: The state of the art* (pp. 57–94). New York: Psychology Press.

McRae, K., De Sa, V. R., & Seidenberg, M. S. (1997). On the nature and scope of featural representations of word meaning. *Journal of Experimental Psychology: General, 126*, 99–130.

Mewhort, D. J., Braun, J. G., & Heathcote, A. (1992). Response time distributions and the Stroop task: A test of the Cohen, Dunbar, & McClelland (1990) model. *Journal of Experimental Psychology: Human Perception and Performance, 18*, 872–882.

Meyer, D. E., Schvaneveldt, R. W., & Ruddy, M. G. (1975). Loci of contextual effects on visual word-recognition. In P. M. A. Rabbitt (Ed.), *Attention and performance V* (pp. 98–118). London: Academic Press.

Morton, J. (1969). The interaction of information in word recognition. *Psychological Review, 76*, 165–178.

Myung, J. I. (2000). The importance of complexity in model selection. *Journal of Mathematical Psychology, 44*, 190–204.

Neely, J. H. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology: General, 106*, 226–254.

Neely, J. H. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. In D. Besner & G. Humphreys (Eds.), *Basic processes in reading: Visual word recognition* (pp. 236–264). Hillsdale, NJ: Erlbaum.

Nelder, J. A., & Mead, R. (1965). A simplex algorithm for function minimization. *Computer Journal, 7*, 308–313.

Nelson, D. L., McEvoy, C. L., & Schreiber, T. A. (1998). The University of South Florida word association, rhyme, and word fragment norms. Available from http://www.usf.edu/FreeAssociation/.

Plaut, D. C., & Booth, J. R. (2000). Individual and developmental differences in semantic priming: Empirical and computational support for a single-mechanism account of lexical processing. *Psychological Review, 107*, 786–823.

Plaut, D. C., & Booth, J. R. (2006). More modeling but still no stages: Reply to Borowsky and Besner. *Psychological Review, 113*, 196–200.

Plourde, C. E., & Besner, D. (1997). On the locus of the word frequency effect in visual word recognition. *Canadian Journal of Experimental Psychology, 51*, 181–194.

Posner, M. I., & Snyder, C. R. R. (1975). Attention and cognitive control. In R. Solso (Ed.), *Information processing and cognition: The Loyola symposium* (pp. 55–85). Hillsdale, NJ: Erlbaum.

R Development Core Team (2004). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review, 85*, 59–108.

Ratcliff, R. (1979). Group reaction time distributions and an analysis of distribution statistics. *Psychological Review, 86*, 446–461.

Ratcliff, R., Gomez, P., & McKoon, G. (2004). A diffusion model account of the LDT. *Psychological Review, 111*, 159–182.

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

Roberts, S., & Sternberg, S. (1993). The meaning of additive reaction-time effects: Tests of three alternatives. In D. E. Meyer & S. Kornblum (Eds.). *Attention and performance 14: Synergies in experimental psychology, artificial intelligence, and cognitive neuroscience* (pp. 611–653). Cambridge, MA: MIT Press.

Rouder, J. N., Lu, J., Speckman, P., Sun, D., & Jiang, Y. (2005). A hierarchical model for estimating response time distributions. *Psychonomic Bulletin & Review, 12*, 195–223.

Rouder, J. N., & Speckman, P. L. (2004). An evaluation of the Vincentizing method of forming group-level response time distributions. *Psychonomic Bulletin & Review, 11*, 419–427.

Rouder, J. N., Tuerlinckx, F., Speckman, P. L., Lu, J., & Gomez, P. (submitted for publication). Modeling parametric relations with response time: The case of word frequency. *Psychological Review*.

Schmiedek, F., Oberauer, K., Wilhelm, O., Süß, H.-M., & Wittmann, W. W. (2007). Individual differences in components of response time distributions and their relations to working memory and intelligence. *Journal of Experimental Psychology: General, 136*, 414–429.

Schwarz, W. (2001). The ex-Wald distribution as a descriptive model of response times. *Behavior Research Methods, Instruments, & Computers, 33*, 457–469.

Seidenberg, M. S., Waters, G. S., Sanders, M., & Langer, P. (1984). Pre- and post-lexical loci of contextual effects on word recognition. *Memory & Cognition, 12*, 315–328.

Speckman, P. L., & Rouder, J. N. (2004). A comment on Heathcote, Brown, and Mewhort's QMLE method for response time distributions. *Psychonomic Bulletin & Review, 11*, 574–576.

Spieler, D. H., Balota, D. A., & Faust, M. E. (1996). Stroop performance in healthy young and older adults and individuals with dementia of the Alzheimer's type. *Journal of Experimental Psychology: Human Perception and Performance, 22*, 461–479.

Swinney, D. A. (1979). Lexical access during sentence comprehension. (Re)consideration of context effects. *Journal of Verbal Learning & Verbal Behavior, 18*, 645–659.

Van Zandt, T. (2000). How to fit a response time distribution. *Psychonomic Bulletin & Review, 7*, 424–465.

Van Zandt, T. (2002). Analysis of response time distributions. In J. T. Wixted (Vol. Ed.) & H. Pashler (Series Ed.) *Stevens' Handbook of Experimental Psychology* (3rd Edition), *Volume 4: Methodology in Experimental Psychology* (pp. 461–516). New York: Wiley Press.

Vincent, S. B. (1912). The function of vibrissae in the behavior of the white rat. *Behavioral Monographs, 1* (Whole No. 5).

Whittlesea, B. W., & Jacoby, L. L. (1990). Interaction of prime repetition with visual degradation: Is priming a retrieval phenomenon? *Journal of Memory & Language, 29*, 546–565.

Yap, M. J., & Balota, D. A. (2007). Additive and interactive effects on response time distributions in visual word recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 33*, 274–296.

Yap, M. J., Balota, D. A., Cortese, M. J., & Watson, J. M. (2006). Single- versus dual-process models of lexical decision performance: Insights from response time distributional analysis. *Journal of Experimental Psychology: Human Perception and Performance, 32*, 1324–1344.

Yap, M. J., Balota, D. A., Tse, C.-S., & Besner, D. (in press). On the additive effects of stimulus quality and word frequency in lexical decision: Evidence for opposing interactive influences revealed by RT distributional analyses. *Journal of Experimental Psychology: Learning, Memory and Cognition*.