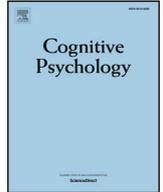


Contents lists available at [ScienceDirect](#)

Cognitive Psychology

journal homepage: www.elsevier.com/locate/cogpsych

The elimination of positive priming with increasing prime duration reflects a transition from perceptual fluency to disfluency rather than bias against primed words

Kevin W. Potter^{a,*}, Chris Donkin^b, David E. Huber^a^a University of Massachusetts Amherst, Amherst, 01003 MA, USA^b University of New South Wales, Sydney, 2052 New South Wales, Australia

ARTICLE INFO

Keywords:

Repetition priming
 Perceptual fluency
 Response time modeling
 Perceptual identification
 Perceptual habituation

ABSTRACT

With immediate repetition priming of forced choice perceptual identification, short prime durations produce positive priming (i.e., priming the target leads to higher accuracy, while priming the foil leads to lower accuracy). Many theories explain positive priming following short duration primes as reflecting increased perceptual fluency for the primed target (i.e., decreased identification latency). However, most studies only examine either accuracy or response times, rather than considering the joint constraints of response times and accuracy to properly address the role of decision biases and response caution. This is a critical oversight because several theories propose that the transition to negative priming following a long duration prime reflects a decision strategy to compensate for the effect of increased perceptual fluency. In contrast, the nROUSE model of Huber and O'Reilly (2003) explains this transition as reflecting perceptual habituation, and thus a change to perceptual disfluency. We confirmed this prediction by applying a sequential sampling model (the diffusion race model) to accuracy and response time distributions from a new single item same-different version of the priming task. In this way, we measured strategic biases and perceptual fluency in each condition for each subject. The nROUSE model was only applied to accuracy from the original forced-choice version of the priming task. This application of nROUSE produced separate predictions for each subject regarding the degree of fluency and disfluency in each condition, and these predictions were confirmed by the drift rate parameters (i.e., fluency) from the response time model in contrast to the threshold parameters (i.e., bias).

1. Introduction

Fluency, the ease/speed with which we process information, can strongly influence our judgments and decisions (Alter & Oppenheimer, 2009). In other words, people can base their decisions not only on the content of a set of information, but also on the evaluation of how easily and/or quickly they processed that information. Research indicates that fluency provides an adaptive heuristic for judgments, capturing relevant statistics from the environment (e.g., Hertwig, Herzog, Schooler, & Reimer, 2008). Indeed, decisions can often be based predominantly on fluency rather than cognitive content (Schwarz et al., 1991). In this article, we focus on perceptual fluency, an internal measure of the time taken to process perceptual stimuli. We examine the role perceptual fluency plays in a word identification task, in particular its role in explaining priming benefits following short duration primes and the

* Corresponding author.

E-mail address: kevin.w.potter@gmail.com (K.W. Potter).

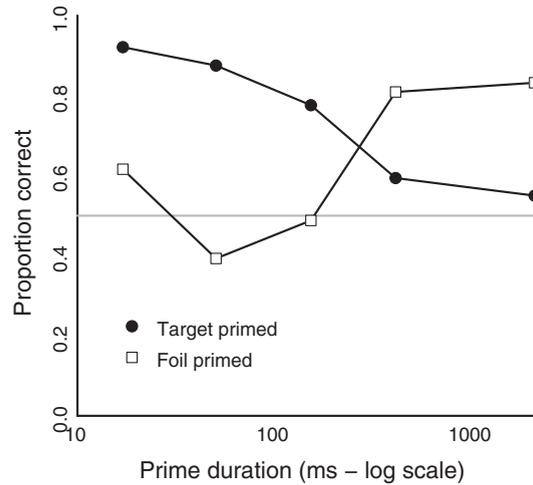


Fig. 1. Proportion correct in a forced-choice perceptual identification task (Huber, 2008). Filled points represent conditions for which the target word was primed; Empty points represent conditions for which the foil word was primed. The x-axis shows the duration for which the prime was displayed. Prime durations ranged from 17 to 2000 ms.

elimination of these benefits following longer duration primes.

Jacoby and Whitehouse (1989) provided one of the first demonstrations of this prime duration effect. After studying a list of words, subjects performed a yes/no episodic recognition task. Subliminal presentation of the test word immediately prior to a test trial produced a “memory illusion”, with priming increasing both hits and false alarms. This effect was explained as an attribution error: subliminal priming increasing the perceptual fluency of the test word, and this fluency was mistakenly attributed to a sense of episodic familiarity. Support for this conclusion was found by examining a condition with a long duration easily seen prime, which produced decreased hits and false alarms (i.e., negative priming rather than positive priming). Jacoby and White argued that when subjects were aware of the potential source of perceptual fluency, decision processes were able to correct for this misattribution (the fluency was attributed to the prime rather than episodic familiarity). Subsequent work demonstrated that this prime duration effect can be found even when people are aware of the short duration prime (Huber, Clark, Curran, & Winkielman, 2008). However, despite attributions to perceptual fluency, decision biases, and response caution, the aforementioned studies did not directly measure these constructs.

Immediate word priming in perceptual identification also produces differing effects based on the duration of the prime. For instance, Huber, Shiffrin, Lyle, and Ruys (2001) used perceptual identification with forced-choice testing, comparing a condition where the target word was primed versus a condition where the foil word was primed. The authors found that short duration primes produced positive priming, as priming the target word improved accuracy and priming the foil word worsened accuracy. However, longer prime durations led to negative priming, as priming the target worsened performance and priming the foil improved performance.

This crossover interaction has been replicated many times; Fig. 1 presents an example from a study by Huber (2008). The y-axis denotes average proportion correct, and the x-axis indicates the duration (in log ms) that the prime was shown. The filled points show performance when the target word was primed, whereas the empty points show performance when the foil word was primed. The study demonstrates that as prime durations exceed 150 ms, accuracy performance improves substantially for the foil-primed condition, yet dips almost to chance performance for the target-word condition. The transition is distinctly non-linear as well.

The nROUSE model (Huber & O’Reilly, 2003) is a dynamic neural network that successfully accounts for this transition from positive to negative priming with increasing prime duration. The nROUSE model captures the time course of activation in early visual processing for the representations underlying the target and foil choice words. The model tracks the time it takes to process the onscreen choice alternatives, providing an explicit measure of perceptual fluency for each alternative. The subject chooses the most quickly processed alternative (i.e., the alternative whose neural representation reaches its peak activation soonest). Psychologically, this represents a decision process in which an observer uses a heuristic, picking the more fluent choice. This is a rational heuristic in a situation where the target is not overtly identified: the briefly flashed target imparts lingering activation for the target choice, reducing the perceptual identification latency of that item at the time when it reappears as a choice alternative.

In the case of the nROUSE model, perceptual fluency is an effective heuristic, but this opens the door for priming effects, with the prime influencing decisions in either direction depending on whether the prime provides lingering activation or lingering habituation. Lingering activation from short duration primes makes the primed choice more fluent, boosting accuracy when the target is primed, but reducing accuracy when the foil is primed. In contrast, after presentation of a long duration prime, lingering habituation makes it difficult to perceptually identify the primed choice. This disfluency has the opposite effect, reducing accuracy when the target is primed, but boosting accuracy when the foil is primed.

The nROUSE model is relevant across multiple domains of psychology, providing insights into how neural mechanisms like inhibition and synaptic depression produce priming effects, how meta-cognitive evaluations of speed of processing can drive

decisions, and how people separate or integrate moment-to-moment visual stimuli in an optimal manner. Demonstrating the generality of its accounts, the nROUSE model has successfully explained forced-choice accuracy data in a wide variety of situations including word repetitions (Huber, 2008), face repetitions (Rieth & Huber, 2010), immediate change detection (Davelaar, Tian, Weidemann, & Huber, 2011), episodic recognition priming (Huber, Clark, et al., 2008), semantic satiation (Tian & Huber, 2010), continual-flash suppression (Huber, 2015), spatial attention cueing (Rieth & Huber, 2013) and temporal attention (Rusconi & Huber, 2017). In addition, the nROUSE model has explained EEG/MEG neural data for several of these tasks (Davelaar et al., 2011; Huber, Tian, Curran, O'Reilly, & Woroch, 2008; Tian & Huber, 2013).

Although the nROUSE model successfully explained and predicted perceptual dynamics over many paradigms, its assumption that the decision process is based on perceptual fluency remains untested. This assumption is a critical aspect of the model, and without this assumption, the nROUSE model would not produce the observed transition from positive to negative priming. For instance, if the decision was based on the maximum achieved activation for each choice alternative, the nROUSE model would only produce negative priming owing to habituation, with the magnitude of the negative priming growing with increasing prime duration.

We sought to closely examine the nROUSE model's assumption that perceptual fluency drives the decision process, with changes in perceptual fluency explaining the transition from positive to negative priming. Many theories assume that brief or subliminal primes influence behavior via a fluency heuristic (e.g., Jacoby, Kelley, & Dywan, 1989; Whittlesea & Williams, 1998, 2000). In addition, some theories have been applied to the prime duration effect, explaining the elimination of positive priming as instead reflecting strategic discounting (e.g., Jacoby & Whitehouse, 1989; Ratcliff & McKoon, 2001), rather than any change in fluency. For instance, following a long duration prime, the subject may first determine whether the test word is or is not primed, and if it has been primed, they may adopt a more conservative evidence threshold to counteract the prime's influence on perceptual fluency (i.e., a bias against primed words).

Past work (e.g. Jacoby & Whitehouse, 1989; Whittlesea & Williams, 2000) has focused on measuring either on accuracy or response times. However, contrasting the strategic discounting explanation against the nROUSE's disfluency explanation of the prime duration effect requires examining both fluency and decision biases. This is because fluency and decisional biases simultaneously impact accuracy and response times in a complex fashion, and therefore an examination of either measure separately will not suffice. We instead applied a sequential sampling model to the joint distribution of accuracy and response time. To gain the data needed for this test, we ran a priming study with forced-choice testing, which is ideally suited for application of a sequential sampling response time model. We applied the nROUSE model to the forced-choice accuracy data separately for each subject to generate predicted identification latencies (i.e., fluencies) in each condition. To test these predictions, we then applied a sequential sampling model to the choice and response time distributions from same-different testing.

Sequential sampling models are measurement tools for untangling the joint effects of response caution, bias, and rate of evidence accumulation that underlie any decision making process. For instance, because people may tradeoff speed for accuracy, performance cannot be determined by only examining speed or only examining accuracy. Sequential sampling models consider the shapes of correct and error response time distributions to provide psychologically meaningful parameter values. Thus, if the prime duration effect reflects strategic discounting, this should be revealed as a change in the threshold level of evidence (aka, bias) necessary to give a response. However, if the prime duration effect reflects disfluency, as assumed by the nROUSE model, this should be revealed as a reduction in the evidence accumulation rate (aka, 'drift rate').

To contrast the strategic discounting explanation versus the perceptual disfluency explanation of the nROUSE model, we focused on two questions. First, would subjects strategically discount the primed choice with long prime durations in the same-different task, as evidenced by higher thresholds (i.e., greater bias)? The nROUSE model, in contrast, would predict a reduction in drift rates. Second, would perceptual identification speeds from the nROUSE model, its predictions on fluency, correlate with drift rates estimated from the same-different task? The nROUSE model predicts that perceptual fluency varies in systematic ways across all conditions. Therefore, we would expect a strong correlation between the estimated speed with which subjects identify choice alternatives in the forced-choice task and the rate of perceptual evidence accumulation in the same-different task. However, under a strategic discounting approach, one would not expect any perceptual disfluency at longer prime durations, leading to a prediction of no correlation between the identification speeds and the drift rates. Determining the correlation between the identification speeds from nROUSE and drift rates allowed us to also examine how well the nROUSE model captures individual differences. To give an example, if a particular subject has strong negative priming as measured with forced choice accuracy, nROUSE will explain this as slow perceptual identification latencies for primed words, and these assumed latencies will predict that the individual should also be slow to respond 'same' to a test word that matches a long duration prime in same-different testing.

To summarize, we sought to determine whether subjects use perceptual fluency (an internal measure of the speed with which the onscreen choices are identified) to make a decision in threshold word identification and whether the change from short to long duration primes reflects a change in fluency versus bias against primed words. In Section 2, we give a more detailed description of the nROUSE model, our choice of sequential sampling model, and their respective predictions. Section 3 describes our experimental methods, while Section 4 details the estimation approaches we used to fit the two models of interest. Section 5 presents our empirical findings, and Section 6 describes our subsequent modeling results. In the final section we discuss our findings, focusing on their implications for both the priming literature and our understanding of the same-different task. For the reader interested in the psychological 'take-home' message, we suggest focusing on Sections 3, 5, and 6.3.

2. The models

In this section, we elaborate in greater detail on the nROUSE model and its mechanisms, in particular focusing on its assumptions regarding perceptual fluency and the decision process. We then discuss our choice of sequential sampling models: the diffusion race model.

2.1. The nROUSE model

The nROUSE model is a neural instantiation of a Bayesian model proposed by Huber et al. (2001) to explain the pattern of costs and benefits over different prime durations for immediate word priming. The Bayesian model, termed ‘Responding Optimally with Unknown Sources of Evidence’ (ROUSE), assumes the perceptual system attempts to correct for the influence of the prime, applying too little discounting to primed features following a short duration prime but too much discounting to primed features following a long duration prime. However, the ROUSE model is a static model, making no specific predictions regarding the time course of priming or perceptual identification. Instead, different model parameters are needed to explain each prime duration. Huber and O’Reilly (2003) developed the nROUSE model to overcome this limitation and relate discounting to a neural implementation.

The original specification of the nROUSE model is reported by Huber and O’Reilly (2003), and we only briefly summarize the perceptual hierarchy of the neural model, along with the model parameters and their interpretations. The nROUSE model is a simplified point neuron model combined with a rate-coded equation for synaptic depression, the temporary deficit in a neuron’s ability to excite other neurons. Synaptic depression can occur, for instance, when the sending neuron suffers from a lack of available neurotransmitter owing to recent activity (e.g., Tsodyks & Markram, 1997). In other words, the model has a specific mechanism for neural habituation.

The success of the nROUSE model in a wide variety of rapid serial visual presentation situations suggests that neural habituation exists to temporally parse the stream of information in a rapidly changing environment. By relating neural habituation to Bayesian discounting, the gradual accrual of habituation in response to a perceptual object can be viewed as approximating the increasing probability that any active primed representation should be considered a hold-over from the prime, rather than something new. In this manner, neural habituation makes it easier to rapidly perceive a new perceptual object with minimal source confusion from previous presentations. However, this mechanism for parsing the stream of information comes at a cost, producing a form of repetition blindness (Kanwisher, 1987) whenever perceptual representations immediately repeat.

2.1.1. The perceptual hierarchy

The perceptual hierarchy for the nROUSE model has three levels. The lowest level includes visual features for each spatial region, the middle level consists of a general pool of orthographic features, and the top level is a general pool of lexical-semantic features. In this regard, nROUSE is similar to the interactive-activation model of McClelland and Rumelhart (IAC; 1981). However, unlike the IAC model, nROUSE does not attempt to explain the distributed nature of lexical representations, instead using a small set of fixed localist nodes in its attempt to describe the dynamic time course of perception. Fig. 2 provides a diagram showing the set of localist nodes in each layer and the connections between them. The membrane potential of nodes within a level are updated each millisecond of time based on the excitatory and inhibitory inputs of bottom-up and top-down connections (the black and grey arrows in the diagram, respectively), adjusted by lateral inhibition (the dashed lines in the diagram) from the nodes within the current level and a leak term.

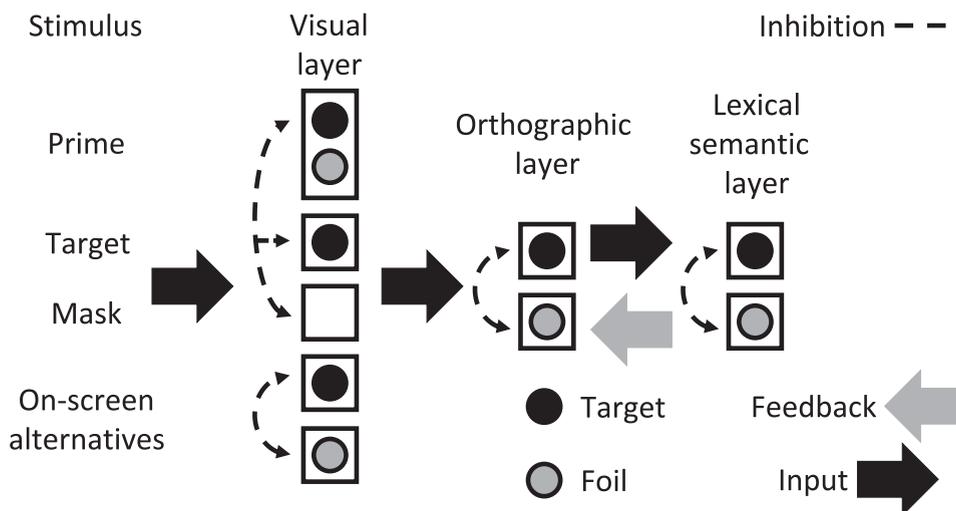


Fig. 2. The localist nodes making up the three-layer perceptual hierarchy of the nROUSE model (Huber & O’Reilly, 2003). Dark arrows denote bottom-up excitatory and inhibitory input, while grey arrows denote top-down feedback. Dashed lines indicate lateral inhibitory connections. Black circles indicate a node for the target word representation, while grey circles indicate a node for the foil word representation.

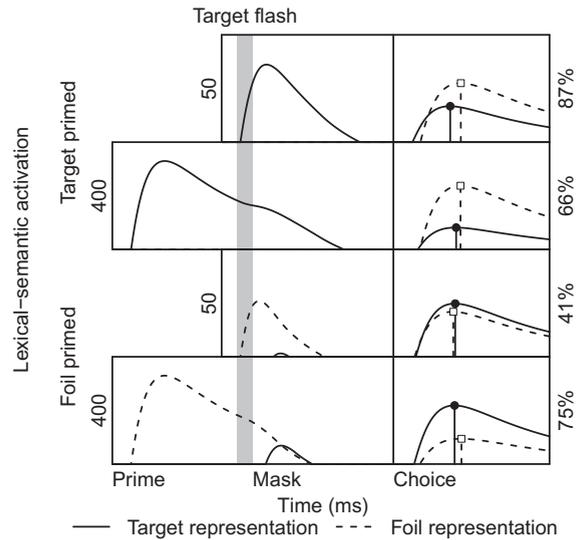


Fig. 3. Activation at the lexical-semantic layer for target and foil word representations across the presentation prime, target, mask, and choice alternatives for four different priming conditions. Solid lines indicate target representations, dashed lines indicate foil representations. The gray window denotes the brief display of the target word. Points indicate the peak activation for each representation during the choice phase. The predicted accuracy (based on the difference between the times to peak activation) is shown to the right.

Lateral inhibition limits the amount of total activation and sets up winner-takes-all attractor dynamics.

The nROUSE model has a complex functional form with a large number of parameters. However, most of the parameters were fixed; in the following we list these parameters and their associated values. Three separate parameters are required to specify the differing rates of integration for each level: S_1 (visual) = .0294, S_2 (orthographic) = .0609, and S_3 (lexical-semantic) = .015. Additionally, parameters are needed for the rates of depletion ($D = .324$) and recovery ($R = .022$) of synaptic resources. The depletion/recovery rates, along with a constant leak current ($L = .15$), within layer inhibition, and a firing rate threshold ($\theta = .15$), are set to the same values for all three levels (i.e., these are general properties of neural processing). In addition to these parameters, the nROUSE model includes feedback from the third level to the second level, with a magnitude of $F = .25$, which can be contrasted with the default feedforward weights set to 1.0. We fixed the values of these parameters based on work from previous studies. As was done by Huber and O'Reilly, the parameters for the firing threshold, feedback, and the lexical-semantic integration constant were set to sensible values approximating the electrophysiological recordings of Tsodyks and Markram (1997). Additionally, the values for the depression and recovery rates were based on the best-fitting results to word priming data from Huber and O'Reilly. Finally, the values for the visual and orthographic integration constants were based on the best-fitting results to further word priming data from Rieth and Huber (2017).

2.2. Perceptual fluency and the decision process

Given the model parameters, along with the durations of the presentation of the primes, masks, and targets, the nROUSE model is simulated, generating activation levels for the target and foil word representations every millisecond. Fig. 3 provides a visual demonstration of how activation at the lexical-semantic layer changes across the different stimulus presentations (the prime, target, mask, and choice alternatives) under four different priming conditions. The top two panels show activation when the target word has been primed for 50 or 400 ms. The bottom two panels show activation when the foil word has been primed for 50 or 400 ms. Solid lines represent activation for the target word, and dashed lines represent activation for the foil word. The grayed out window indicates the interval for the brief display of the target word. For easy comparison, we have marked the peak activation for both the target and foil representations. Finally, we present the predicted accuracies for each of the four conditions on the right-most axis. Hence, starting from the presentation of the on-screen alternatives, the time taken to reach peak activation at the lexical-semantic layer can be determined for each of the choice words (the target or foil). This time to peak activation represents the nROUSE model's prediction of the time it takes the visual system to process and identify the onscreen choice alternatives. In other words, the time to peak activation can be taken as a predicted perceptual identification latency. This predicted identification speed can then be used to implement a decisional heuristic based on perceptual fluencies: subjects pick the choice whose representation reaches peak activation first, or in other words, the choice that the visual system identified more quickly.

Fig. 3 emphasizes why the perceptual identification latencies form the basis of decision rule for the nROUSE model. As seen in the figure, the unprimed representation always has a higher amount of activation compared to the primed representation due to neural habituation. In other words, as noted earlier, if the decision rule used maximum activation, the model would predict that subjects would always pick the unprimed representation. Therefore, for reasonable predictions, the decision process must be based on perceptual fluencies (i.e., the time to peak activation).

Note that, for the nROUSE model, simulations with a given set of parameters and stimulus durations always produce the same identification latencies (i.e., the same choice outcome); some form of decisional variability is required to make accuracy predictions. We adopt the variability assumptions made by Rieth and Huber (2017), which involve a race between normally distributed perceptual identification latencies. The simulation-determined identification latencies specify the means of these two distributions and it is assumed that the variance of each distribution is proportional to its mean latency. Assuming that the two distributions are independent of each other, accuracy is determined from a difference distribution, with a mean equal to the mean difference and variance equal to the sum of the variances. The variances are weighted by a constant of proportionality relating mean latency to the variance of the latencies (henceforth termed the ‘noise multiplier’).¹

Rather than assuming a race process with normal distributions, the original version of the nROUSE model used a logistic function to determine accuracy from the difference between the target and foil perceptual identification latencies. Although the change to a race between normal distributions is negligible when using typical parameters that fit accuracy averaged across individuals, it has important consequences when using parameters that produce greatly different perceptual identification latencies (i.e., such as might occur when comparing individuals who require greatly different target durations to achieve threshold performance). More specifically, with the normal distribution assumptions made by Rieth and Huber, accuracy depends not only on the difference between the target and foil perceptual identification latencies but also the absolute magnitude of these latencies. Thus, the decisional assumptions of Rieth and Huber are more readily applied to comparisons across individuals, and we adopted their default parameters from their Experiment 1 (as noted earlier).

2.2.1. Individual differences and nROUSE

The nROUSE model can make a priori predictions about group-level performance using the default values listed earlier for the parameters. However, in the current study, we sought to examine how the perceptual identification latencies predicted by the nROUSE model related to parameters from a sequential sampling model. When we collected data for this study, subjects varied substantially in their performance. Because this variation would be reflected in the parameter estimates of the sequential sampling model, we chose to examine individual differences in estimates for the nROUSE model.

The nROUSE model can be fit to forced-choice accuracy data using maximum likelihood estimation to find the best fitting parameters for each subject (see Appendix A in the online supplementary materials for the likelihood equations). However, the nROUSE model has never been fit to subject-level data before, only group-level data aggregated over subjects. Therefore, to capture individual differences, we allowed three parameters to vary: the noise multiplier, the inhibition parameter, and a new parameter representing the ability to attend at the time the briefly flashed target appeared (henceforth termed the ‘temporal attention’ parameter). In the following, we detail the specific aspects of the individual differences these parameters are meant to capture.

In our experimental paradigm (discussed in subsequent sections), for each subject we selected in adaptive fashion a duration for the target flash to maintain a threshold accuracy of 70%. Subjects varied greatly in the duration needed to achieve said threshold accuracy. We hypothesized that the temporal attention parameter caused these large individual differences. Alternative explanations of threshold target duration, such as differences in reading ability or visual acuity, did not seem warranted for a population of healthy undergraduates with normal or corrected to normal vision. When simulating nROUSE, visual input is set to 1.0 when the stimulus preferred by a level one unit is presented, and set to 0 otherwise. We implemented individual differences in temporal attention by setting visual input for the briefly flashed target word to a subject-based free parameter (potentially above or below 1.0), whereas the visual inputs for all other presentations were set to 1.0. In this manner, the model captured any relative differences in attention to the briefly flashed target word as compared to the rest of the display sequence.

Individual differences in target duration also impacted nROUSE parameter estimation in another fashion. Activation in the nROUSE model is updated each millisecond based on the current visual inputs. Hence, the duration of the visual input for the target presentation was different for each subject. To reflect this, nROUSE simulations for each individual used the actual threshold target durations assigned to that individual to maintain performance at 70%. More specifically, simulations used the average threshold target duration across all of the experimental trials for a given individual, bearing in mind that threshold target durations were potentially adjusted every 80 trials. In light of this, the free parameter for temporal attention was not so much used to fit the pattern of data across the four conditions so much as it was used to normalize behavior of the model in light of these different threshold target durations (i.e., in some sense, the model was fitting more than just four conditions considering that target duration was another observation that was imposed on the model).

Subjects also differed substantially on the magnitude of negative and positive priming over different conditions. We hypothesized that the noise multiplier and inhibition parameters would help capture these individual differences. The noise multiplier parameter does not affect model simulations at all. However, it does rescale the results by transforming simulation-determined latencies into accuracy values, thus magnifying or reducing differences between conditions.

The inhibition parameter, in turn, plays a critical role for the magnitude of negative priming at longer prime durations. As will be seen, subjects ranged from exhibiting strong negative priming to merely having a reduction in positive priming. Because these differences lie on either side of zero, they can only be explained by a change in the dynamics of the nROUSE model. We hypothesize that these differences reflect difference in strength of inhibition, which might correspond to a different form of attention (i.e., the extent to which each subject attempts to identify just a single word at any moment, versus letting multiple words remain concurrently

¹ This is a common assumption in modeling work. For instance, consider a situation in which the normal distribution is an approximation of a binomial with large N . As N increases, so does the variance.

active). Differences in inhibition affect the dynamics of the model by modulating the effectiveness of habituation. More specifically, because each unit also experiences self-inhibition, stronger inhibition serves to limit the extent to which habituation can accrue owing to an initially strong activation response. As a consequence, with weak inhibition, there is greater habituation and thus greater repetition blindness for the briefly flashed target word. This greatly reduces target primed performance with habituation, resulting in negative priming following a long duration prime (i.e., target primed accuracy lower than foil primed accuracy). Therefore, the noise multiplier, inhibition, and temporal ability parameters should all be necessary to explain individual differences.

2.3. The diffusion race model

The nROUSE model is an account of perceptual dynamics. As such, it is not well-suited to explain the complexities of decision making such as response biases (a preference for one response key versus the other) and the tradeoff between speed and accuracy. Nevertheless, the nROUSE model makes some basic assumptions about the decision process. The nROUSE model assumes that forced-choice accuracy reflects a relative comparison between the speed of perceptual identification for the target versus the speed of perceptual identification for the foil. In other words, perceptual fluency underlies the decision process. Simulations with the nROUSE model present the choice words simultaneously, with both words processed in parallel. Thus, simulations with nROUSE correspond to a parallel race process between the choices. Nevertheless, nearly identical accuracy predictions would arise if it was assumed that the observer first attended to one choice, noting the latency to identify that word, and then turned to the alternative before deciding which word was identified more quickly. Building off of this fact, the use of perceptual fluency in the decision process should hold true even if subjects carry out same-different instead of forced-choice testing.

The nROUSE model posits that the positive and negative priming effects reflect the interaction of lingering activation and neural habituation. Due to the nature of neural habituation, perceptual fluency (i.e., perceptual identification latencies) must drive the decision process for the model to work. Thus, while the nROUSE model is agnostic about how evidence accumulates for a ‘different’ response, perceptual identification latencies should still underly ‘same’ responses. Hence, we can use both forced-choice and same-different testing to examine the roles of perceptual fluency versus strategic discounting.

Subjects completed both forced-choice and same-different testing within the immediate word priming paradigm. However, we still faced several challenges. The same-different task provided us with accuracy and response time data, but because we sought to determine how perceptual fluency and strategic discounting impacted the priming effects, we could not consider accuracy or response times separately. To do so would ignore the complicated relationship between response caution, biases, and how well subjects extracted information from the perceptual stimuli.

Ideally, we would like to decompose observed response times into latent measures that better represent the cognitive processes of interest. Thus we drew upon decades of research with sequential sampling models, a set of models that have proven quite successful in jointly modeling simple choices and response times (e.g., Luce, 1986). Sequential sampling models assume that decisions are based on information sampled from noisy stimulus representations in the central nervous system (Smith, 2000). We focus here on race models which assume a decision is based on two or more stochastic accumulators racing toward separate decision thresholds, closely matching the architecture of the nROUSE model. The accumulator that reaches its threshold first wins, determining the choice and response time. An attractive feature of these models is their neural plausibility: the accumulators can be likened to the firing rates for separate neural populations processing evidence for each alternative.

We selected our choice of sequential sampling model based on a combination of neural plausibility and computational tractability. As a contrast, consider the leaky competing accumulator model (LCA; Usher & McClelland, 2001), a good example of a race model built to be neurally plausible. The LCA model has additional mechanisms for the decay and inhibition typically seen in neuron populations. Variants of this type of model have been successfully applied to both behavioral data and single cell firing rates (e.g., Ratcliff, Cherian, & Segraves, 2003; Ratcliff, Hasegawa, Hasegawa, Smith, & Segraves, 2007; Ratcliff et al., 2011). In these models, the inclusion of the decay and inhibition mechanisms make it computationally challenging to estimate parameters. We therefore focused on the diffusion race model (Logan, Van Zandt, Verbruggen, & Wagenmakers, 2014), where evidence for each racer accumulates according to a one-boundary Wiener process and the racers are assumed to be independent. Therefore, the finishing times for the accumulators follow an inverse Gaussian or Wald distribution. The assumption of independence makes the model tractable and easy to compute, allowing for fast and stable parameter estimation (Appendix A in the online supplementary materials details the equations for the likelihood function).

Fig. 4 provides a visual representation of how the two stochastic accumulators in this sequential sampling model race other. Evidence accumulates at an average rate of ξ_1 and ξ_0 for the respective racers. In this case, the accumulator for choice 1 reaches its threshold κ_1 first, winning the race and determining the choice and decision time. The non-decision component τ is added to the decision time, producing the observed response time. The thresholds for the two racers are set to be equal in the figure. However, if a subject is biased towards a particular choice, this bias can be represented in the model by setting the corresponding threshold to be lower relative to the threshold of the opposing choice. An increase in response caution would in turn be represented by an equal increase in both thresholds. For the likelihood function of the diffusion race model, consult Appendix A.

Applying the diffusion race model to the same-different task allows us to convert the observed choice and response times into drift rates and threshold estimates. The nROUSE model posits that perceptual fluency, specifically the perceptual identification latencies, drives the evidence for ‘same’ responses. This assumption therefore implies that the predicted perceptual identification latencies extracted from the forced-choice task should be able to predict drift rates for the ‘same’ response, but should be unrelated to the threshold estimates. In other words, if speed-of-processing drives the decision process, the nROUSE model’s estimate of said speed should relate to the rate of evidence accumulation, not the total amount of evidence needed for a decision.

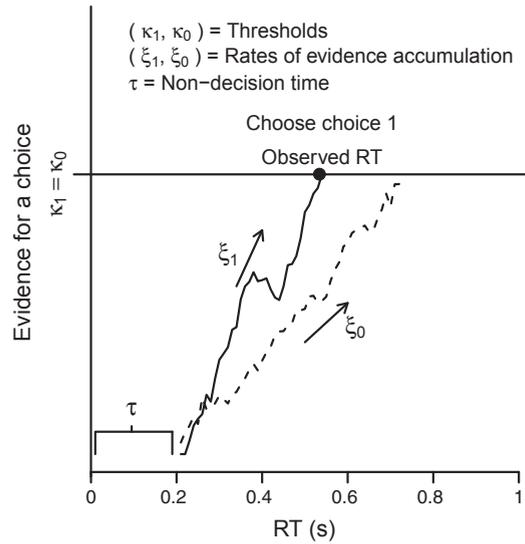


Fig. 4. A visual representation of the diffusion race model. The solid line shows the path of the stochastic accumulation of evidence for picking choice 1 within a single trial. The dashed line shows the evidence accumulation for picking choice 0. Accumulation stops once one of the thresholds is reached. Because the evidence for picking choice 1 reached its threshold first, the subject in this case would pick choice 1, and the observed response time is equal to the finishing time for the racer associated with choice 1 plus the non-decision component τ .

We specified separate drift rates for ‘same’ and ‘different’ responses over each condition in our study. Unlike drift rates, which reflect the quality/quantity of evidence, we assume thresholds are under strategic control and can only be adjusted based on consciously perceived properties of the experimental situation. Thus, the threshold for trials when the correct answer is ‘same’ cannot differ from trials when the correct answer is ‘different’ because such variation implies that the participant already knows the correct answer. However, participants are likely aware of whether the prime (particularly in the case of long duration priming). Finally, participants may set different evidence thresholds for giving ‘same’ responses as compared to ‘different’ responses (i.e., a simple response bias). This resulted in eight different thresholds, as seen in Table 1.

Note that the assumption that subjects have differing thresholds based on whether the test word was primed or not reflects an overly idealistic scenario. We do not actually believe subjects can determine whether a test word was primed or not with 100% accuracy. However, the benefits of such an assumption outweigh the costs. For instance, subjects come close to this degree of accuracy in the case of long duration primes (e.g., Jacoby & Whitehouse, 1989). As for short duration primes, due to the structure of our experimental design (outlined below), the primes were not masked as efficiently as the target flash. Hence, it is still possible that subjects could determine whether the test word was primed or not even with short duration primes. Therefore, it is useful to allow separate thresholds to capture biases for prime and test word match even with short prime durations.

2.4. Summary

The nROUSE model, when simulated, outputs a pair of perceptual identification latencies, a measure of the speed with which the perceptual system processed the target and foil word representations. Hence, the model provides explicit predictions for perceptual fluency. Furthermore, the nROUSE model assumes that during forced-choice testing a subject makes a decision based on the relative comparison between these identification latencies. Due to how the mechanisms for lingering activation and neural habituation

Table 1
The parameter mappings for the diffusion race model.

Correct choice	Prime	Duration	Same racer			Different racer		
Same	Foil	50 ms	θ_{t_1}	κ_1	ξ_1	θ_{t_2}	κ_7	ξ_{13}
		400 ms	θ_{t_3}	κ_2	ξ_2	θ_{t_4}	κ_8	ξ_{14}
	Target	50 ms	θ_{t_5}	κ_3	ξ_3	θ_{t_6}	κ_5	ξ_{15}
		400 ms	θ_{t_7}	κ_4	ξ_4	θ_{t_8}	κ_6	ξ_{16}
Different	Foil	50 ms	θ_{t_9}	κ_3	ξ_5	$\theta_{t_{10}}$	κ_5	ξ_9
		400 ms	$\theta_{t_{11}}$	κ_4	ξ_6	$\theta_{t_{12}}$	κ_6	ξ_{10}
	Target	50 ms	$\theta_{t_{13}}$	κ_1	ξ_7	$\theta_{t_{14}}$	κ_7	ξ_{11}
		400 ms	$\theta_{t_{15}}$	κ_2	ξ_8	$\theta_{t_{16}}$	κ_8	ξ_{12}

interact with immediate word priming, the decision must be based on a comparison of identification latencies in order for the nROUSE model to generate correct predictions on accuracy. This fact then implies that if one uses same-different testing with immediate word priming, perceptual identification latencies should still underly the decision process. In other words, we should still see positive priming for short prime durations and negative priming for long prime durations, reflecting the influence of perceptual fluency and disfluency, respectively. However, the strategic discounting approach posits that the negative priming instead reflects a conscious bias against the primed choice on the part of the subject.

We contrasted the fluency assumptions of the nROUSE model against the strategic discounting approach by having subjects complete both forced-choice and same-different testing with immediate word priming. We fit the nROUSE model to the forced-choice data of each subject to generate predicted perceptual identification latencies. We used the diffusion race model to convert the choice and response time data from same-different testing into drift rates and threshold parameters, representing the rate of evidence accumulation and the total amount of evidence needed for a decision. If perceptual fluency underlies the decision process for both short and long prime durations, then the perceptual identification latencies should strongly correlate with the appropriate drift rates. If strategic discounting plays a role instead, we would not expect a correlation between drift rates and identification speeds, as there should be no transition from fluency to disfluency as prime duration increases. Instead, at long prime durations subjects should become biased against the primed choice.

We now present our experimental methods, followed by our approach to estimate the parameters of the two models.

3. Method

3.1. Participants

Forty-two University of California San Diego undergraduates participated in the first session of this two-session experiment. Six participants failed to return one week later for the second session. The data files from an additional eight participants contained errors owing to a bug in the timing routine and were unusable. This left 28 complete and accurate data sets. The average accuracy of one participant was greater than 80%, which did not allow for sufficient error trials for response time modeling. The average accuracy of another participant was less than 60%, indicating that the participant was not adequately performing the task and/or had been assigned a target flash duration that was below threshold. A third participant was eliminated from further consideration because of excessively long response times (greater than one second on average). These elimination criteria were determined in advance of any data analyses, and after excluding the data from these three participants, this left 25 datasets that were included in subsequent analyses. Participants provided written informed consent before the experiment, and all participants received course extra-credit for their participation.

3.2. Materials

One thousand five-letter words were used for all word stimuli. These words had a minimum written language frequency of four as measured by Kucera and Francis (1967). Randomly generated letter-like pattern masks were used to avoid pattern-mask habituation (see Fig. 5 for examples of the pattern masks). All words were displayed in capitalized Times Roman 22-point font. Letters were

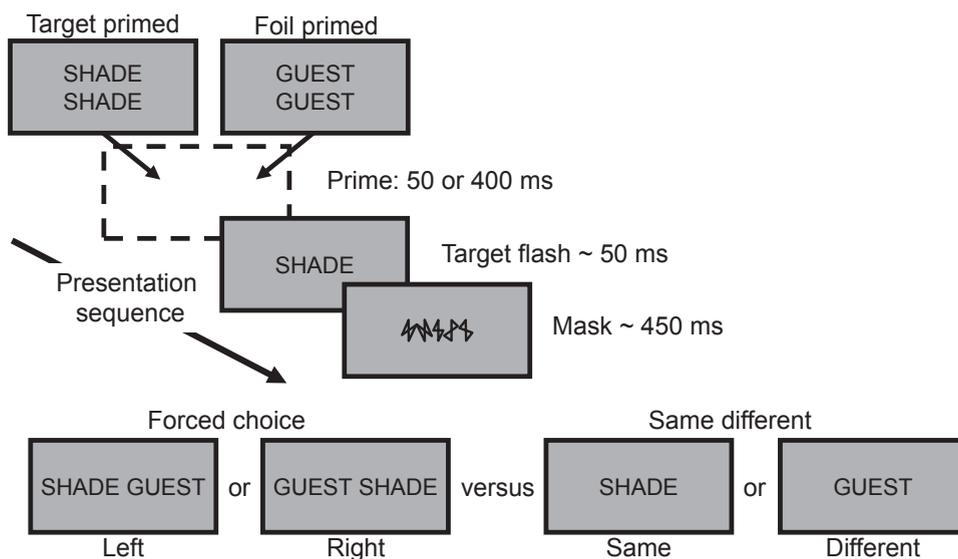


Fig. 5. Presentation sequence with example words, showing all 16 conditions (all combinations of forced-choice vs. same-different testing, target vs. foil primed, 50 vs. 400 ms prime duration, and the correct response: left vs. right or same vs. different, depending on the type of testing).

approximately one degree of visual angle in the vertical direction and less than three degrees of visual angle in the horizontal direction across the 5 letters of each word. Stimuli were gray against a black background and presented on an LCD monitor set to 640 by 480 resolution, with a 60 Hz refresh rate. All presentations were synchronized with the vertical screen refresh, yielding display durations in increments of 16.7 ms. All responses were collected through the computer keyboard.

3.3. Procedure

The experimental design included 16 conditions, as seen in [Table 1](#) and [Fig. 5](#). Eight of these conditions used forced-choice testing whereas the other eight used same-different testing. For forced-choice testing, the eight conditions represented all combinations of screen position for the correct choice word ('left' versus 'right'), repetition priming of the target versus the foil (target primed versus foil primed), and a prime duration that was either short (50 ms) or long (400 ms). For same-different testing, the eight conditions represented all combinations of the correct answer ('same' versus 'different') as well as the same type of priming and prime duration conditions as forced-choice testing. In other words, the same-different conditions were constructed to match the forced-choice conditions. For instance, suppose the target primed condition with forced-choice testing consisted of word A as a prime, followed by word A flashed as the target, and finally a choice between word A versus word B. Then analogously, the target primed condition with same-different testing would consist of word A as a prime, followed by word A flashed as the target, and finally either a single test word that was word A (the same condition) or word B (the different condition). If the foil primed condition with forced-choice testing consisted of word A as a prime, followed by word B flashed as the target, and finally a choice between word A versus word B, then, correspondingly, the foil primed condition with same-different testing consisted of word A as a prime, followed by a word B flashed as the target, and finally either a single test word that was word B (the same condition) or word A (the different condition). Thus, the foil primed condition with same-different testing did not actually involve any priming considering that the prime never reappeared during the trial. However, this is an appropriate label for this condition because it attempts to measure the separate perceptual response to the target word for the foil primed condition with forced-choice testing.

Forced-choice and same-different testing were run in separate blocks of 80 trials. During each of the two experimental sessions, participants alternated between a block of same-different testing and a block of forced-choice testing four different times (eight experimental blocks per session). Each block consisted of 10 repetitions of the eight conditions for each type of testing. Thus, across the two sessions, there were 80 repetitions for each of the 16 conditions. Every trial used two different words. At the start of each block, the allocation of words to conditions was randomized and words were assigned to conditions without replacement such that no word appeared on more than one trial within a block. In addition, the order of conditions within each block was randomized.

Each trial was initiated with a key press (self-paced). A central fixation point (*) then appeared for 250 ms, which was replaced by 5 hyphens (- - - -). The hyphens remained onscreen either for 450 ms or 100 ms depending on the prime duration, such that the duration between the appearance of the hyphens and the appearance of the flashed target word was fixed at 500 ms (thus, subjects knew when to expect the target). The hyphens were replaced by a prime word presented either for 50 ms or 400 ms. The prime word was presented in a doubled up fashion (two copies of the prime word, with one immediately above the other, as seen in [Fig. 5](#)). The prime was subsequently replaced by the briefly flashed target word, which, in turn, was replaced by a mask consisting of five pattern mask characters. The duration between onset of the target and onset of the test display was 500 ms. For same-different testing, the test word appeared in the same central location as the target and mask. For forced-choice testing, the two choice words appeared along the midline, with their centers at 3/7 and 4/7 of the way along the horizontal axis (i.e., 91 pixels apart). With a character spacing of 13 pixels, this corresponds to a gap of 52 pixels between the choice words (well within one degree of visual angle). The 'f' key served as the left response and the different response, while the 'j' key served as the right response and the same response. Immediately after entering a response, subjects received onscreen feedback indicating whether they were correct and how quickly they had responded.

Subjects were informed that the prime word was just as likely to match the correct answer as the incorrect answer, making it clear that there could be no effective strategy based on the prime. The accuracy feedback following each response helped emphasize this fact. The first session began with eight practice trials on the forced-choice task (each of the eight forced-choice conditions occurring once). For these practice trials, the target was presented for 200 ms, which is well above threshold. This target duration was chosen to make it easy to appreciate the nature of the task. After practice, subjects received a block of 80 trials with the forced-choice task that used a staircase estimation procedure to determine a target duration producing threshold accuracy of 70%. This block contained 10 repetitions of each of the eight forced-choice conditions. The instructions prior to this block of trials indicated that response latency was important and that subjects should attempt to be as fast and as accurate as possible. The staircase estimation procedure adjusted the target duration every 20 trials based on performance during the last 20 trials. At the start of this block of trials, the target duration was set at 100 ms. Target duration decreased if accuracy was 80% or greater and it increased if accuracy was 60% or lower. The first adjustment after 20 trials could change the target duration by 33.3 ms in either direction. All subsequent adjustments could change the target duration by 16.7 ms in either direction. Following this target threshold block of trials, the experimental trials began with a block of same-different trials using the target duration that produced threshold performance with forced-choice testing. Throughout both sessions, after each block of forced-choice testing a target duration adjustment was made based on accuracy across all 80 trials within the block. This was done to maintain performance at threshold in the event that the subject continued to improve with the task. These adjustments were made if accuracy across the entire block was greater than 75% or less than 65%.

4. Estimation

This section details our approach for estimating the parameters of the nROUSE model, trimming response time outliers, and estimating the parameters of the diffusion race model. We note that our data, analysis, and estimation scripts are freely available at https://github.com/retopniviek/Analyses/tree/master/2017/SD_v_FC_reanalysis. The custom software for estimating the nROUSE model is available at <https://github.com/retopniviek/nROUSE> and the code used for the diffusion race model can be found at <https://github.com/retopniviek/seqmodels>.

4.1. Estimation for nROUSE model

In order to estimate the parameters of the nROUSE model, we used the statistical software environment *R* (version 3.2.3; [R Core Team, 2016](#)), adapting the equations used by [Huber and O'Reilly \(2003\)](#). For speed, we implemented the underlying algorithms using the packages *Rcpp* (version 0.12.2; [Eddelbuettel & François, 2011](#)) and *RcppArmadillo* (version 0.6.400.2.2; [Eddelbuettel & Sanderson, 2014](#)), a C++ interface for *R* and an associated linear algebra library, respectively.

We fixed the nROUSE parameters to the default values specified by [Rieth and Huber \(2017\)](#), except for the inhibition parameter, noise multiplier, and the new parameter for temporal attention. For each subject, we fit the model to four frequencies, the sum of correct responses for each combination of prime type and duration (collapsing over position correct). We computed the likelihood of these four frequencies under the binomial distribution. We set the parameter governing proportion correct equal to the predicted accuracy generated by simulating the nROUSE model with its three free parameters (see [Appendix A](#) for the likelihood equations). We obtained maximum likelihood estimates using a Nelder-Mead simplex routine ([Nelder & Mead, 1965](#)). To improve the optimization algorithm's exploration of the parameter space, all three free parameters were log-transformed to avoid inadmissible values, ensuring that they would be greater than 0. During our preliminary attempts to fit the nROUSE model, we found that the optimization routine had issues with local maxima. We resolved this issue via a brute force approach. For each subject, we carried out maximum likelihood estimation 100 times using different, randomized starting points. The multiple iterations with diverse starting point values greatly increased our chances of finding the global maximum. This approach, though inelegant, worked well due to the speed with which we could estimate the nROUSE parameters.

We hypothesized that we would need all three free parameters (representing inhibition, the noise multiplier, and temporal attention, respectively) in order to fit the nROUSE to the forced-choice data. In order to test this hypothesis, we carried out a set of model comparisons. For convenience, these comparisons were done at the group-level, fitting constrained variants of the nROUSE model to frequency correct across prime type and duration after aggregating over subjects. Hence, in addition to our unrestricted model with all three free parameters, we fit several restricted models (nested in the unrestricted model), where we constrained one to all of the free parameters. We then computed Akaike's Information Criterion (AIC; [Akaike, 1973](#)) and the Bayesian Information Criterion (BIC; [Schwarz, 1978](#)). Via this approach we could determine the model that best fit the data while controlling for the increased complexity induced by having more free parameters.

4.2. Identifying response time outliers

Outliers in response time data can negatively impact maximum likelihood methods, resulting in biased estimates (e.g., [Ratcliff & Tuerlinckx, 2002](#)). Even a small number of response time contaminants can have a large adverse effect on the estimates of the model parameters. Therefore, we pre-processed our data before fitting the diffusion race model (we did not pre-process our data when fitting the nROUSE model since we only applied it to the forced-choice accuracy data).

We used two approaches to identify outlying responses. First, we used a set of global cutoffs to remove obvious outliers. We removed any response time faster than 200 ms or slower than 2362 ms. We based our choice for the lower boundary on an event-related potential component known as the N170, which is believed to represent early semantic processing (e.g., [Serenó, Brewer, & O'Donnell, 2003](#); [Serenó, Rayner, & Posner, 1998](#)). Therefore, we assumed that non-contaminant responses would need to be slower than 170 ms. We settled on 200 ms in order to incorporate the additional time for a motor response. We determined our choice for the upper boundary by setting our cut-off equal to the mean response time plus four standard deviations, where the mean and standard deviation were calculated after collapsing the data across all 25 subjects, producing a cut-off of 2362 ms. Because response time distributions for subjects varied widely, this cut-off trimmed data primarily for the slowest subjects (another way to put this is that the slowest subjects were slow on average in part because they were more likely to have trials in which they engaged in off-task activities).

We determined our lower and upper cut-offs based on theoretical and statistical considerations, respectively. Again, these across-subjects cut-offs were used as a first pass to eliminate obviously excessively fast or slow response times. However, neither cut-off addresses individual differences. In the case of the lower cut-off, as pointed out by one reviewer, the actual timing of the N170 component varies across studies and subjects. Therefore, while relying on this component results in a less arbitrary cut-off compared to most studies, excessively fast response times could still be present, especially for subjects with on average slower response times. As for the upper cut-off, as noted earlier, the global cut-off is effective at identifying outliers from slower subjects, but likely fails to trim contaminant response times for faster subjects. We therefore also used another approach to identify outlying response times; we used a descriptive mixture model to identify outlier response times for each subject. However, by eliminating obvious outliers with the global cut-offs, we sought to improve the efficacy of the second model-based approach.

As a second pass to eliminate outliers, we relied on a mixture of a descriptive positively-skewed distribution and a uniform

distribution. Specifically, we posited that the data-generating process for an individual's response times (collapsing over all conditions) can be approximated via an inverse Gaussian distribution. However, with some small mixture probability, contaminant response times occur with uniform likelihood over the range of response times. We identified and trimmed outlier response times based on the ratio of the likelihood of a response time under the uniform density over the contrasting likelihood under the inverse Gaussian density. A ratio greater than one indicated that the response time was more likely under the uniform distribution than the inverse Gaussian; we labeled these response times as outliers and removed them from the analyses. Again, note that this approach for contaminant response is purely descriptive, not theory based. Details on the proportion of responses removed by our trimming approach are provided in [Appendix C](#) in our online supplementary materials.

4.3. Estimation for diffusion race model

This section details how we estimated the parameters for the diffusion race model. The diffusion race model has three types of parameters: drift rates, thresholds, and non-decision components. To reiterate, drift rates represent the average rate of evidence accumulation, thresholds represent the total amount of evidence needed to pick a particular choice, and the non-decision component is a shift parameter represent the time taken for encoding, motor processes, and other ancillary processes. The non-decision component cannot exceed the fastest response time. Therefore, we chose to reparameterize this value as the product of a proportion and the fastest response time. [Table 1](#) indicates how we mapped parameters over the different conditions for the same-different task. The symbols κ , ξ , and θ represent the threshold, drift rate, and proportion parameters, respectively. Furthermore, t_1 to t_{16} represent the observed minimum response time for a given condition and response type. Each condition has two associated accumulators, a racer for picking 'same' and a racer for picking 'different', respectively. For each racer, we list the mappings over all three types of parameters.

Given two prime durations, two prime types, and two types of correct response, there were a total of eight conditions. We specified a separate drift rate for each racer and condition. Hence, we estimated 16 drift rates per subject. We constrained thresholds based on the response type ('same' versus 'different') and based on whether the prime matched the test word. There were four possible combinations of prime and test word (foil prime and foil test, foil prime and target test, target prime and foil test, and target prime and target test). We therefore estimated eight different threshold values for subjects.

We specified the non-decision component in a novel way. We sought to estimate this component separately for each condition and racer. However, it is extremely difficult to estimate 16 separate parameters for the non-decision component, especially with only 80 trials per condition. Therefore, we instead initially fixed each non-decision component to be equal to the fastest response time observed for the associated condition and response type. Unfortunately, the non-decision component is typically of shorter duration than the observed fastest response time. Simply setting the non-decision components to equal these latencies will potentially overestimate the actual values. Therefore, we scaled down the fixed values for the non-decision times by weighting them with a single free parameter, a shared proportion θ . Consider [Table 1](#). As noted earlier, the values t_1 to t_{16} represent the observed minimum response time for a given condition and response type (i.e., 'same' versus 'different'). Then, as an example, the non-decision component for the short foil-primed condition when the correct answer was 'same' is $\theta \times t_1$ for picking 'same' and $\theta \times t_2$ for picking 'different'. So, to emphasize, subjects can have differing non-decision components for picking 'same' versus 'different' and for the different conditions, but these estimates will be scaled by a single parameter, θ .

We estimated the diffusion race parameters by first implementing the likelihood function given in [Appendix A](#) in the statistical software environment R. Once again, we used the C++ interface Rcpp for speed. We then carried out maximum likelihood estimation using the R package *optimx* (version 2013.8.7 [Nash & Varadhan, 2011](#)). We could not use the brute force method for finding the global maximum that we employed when estimating the nROUSE model because it took far longer to estimate the diffusion race model parameters. Instead, we estimated the model parameters using several different optimization routines, and selected the estimates with the highest sum of the log-likelihoods. We attempted to find the global maximum using R implementations of the Nelder-Mead simplex routine ([Nelder & Mead, 1965](#)), the Broyden-Fletcher-Goldfarb-Shanno algorithm ([Broyden, 1970](#); [Fletcher, 1970](#); [Goldfarb, 1970](#); [Shanno, 1970](#)), the non-linear minimization routine from [Schnabel, Koontz, and Weiss \(1985\)](#), the unconstrained PORT routines adapted from [Gay \(1990\)](#), and finally, a R implementation by Ravi Varadhan of a bounded Nelder-Mead optimization ([Kelley, 1999](#)). These optimization routines have differing trade-offs in terms of their degree of robustness versus converging upon the true global maximum. To further aid in the exploration of the parameter space, we used the log-transform of the drift rates and thresholds in order to restrict their values to be greater than zero. We restricted the proportion applied to the non-decision component to lie between zero and one via the logistic function.

The estimation of the nROUSE and diffusion race model allowed us to address our two core questions. First, we compared the hypothesis that the transition to negative priming with long duration primes reflects strategic discounting of the primed choice against the assumption of the nROUSE model that it reflects perceptual disfluency. To do so, we analyzed the threshold estimates using two repeated measures ANOVAs with planned contrasts. We then applied a Bayes factor test to compare a model that assumed only a same-different bias against a model that posited a bias towards unprimed choices and against primed choices at long durations. Second, we examined how the nROUSE model fared predicting drift rates from the same-different task based on forced-choice accuracy.

5. Empirical results

We now report the descriptive results. To the best of our knowledge, this is the first time that response times have been compared

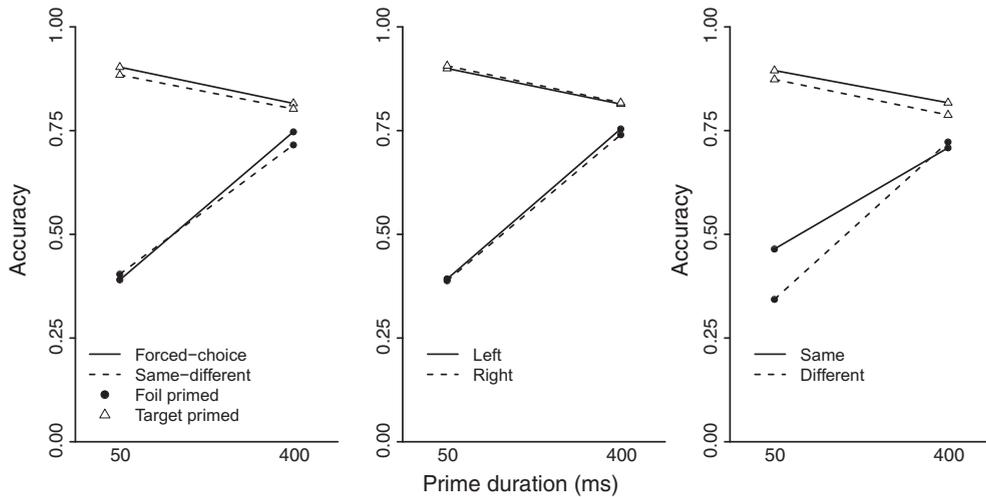


Fig. 6. Accuracy by task, prime type, and prime duration, averaged over subjects. The triangles indicate the target was primed, while the filled circles indicate the foil was primed. Solid and dashed lines denote the task or the type of correct response depending on the panel. The first panel collapses over the type of correct response (Left versus right for the forced-choice task, 'same' versus 'different' otherwise). The remaining two panels present performance based on the type of correct response for the forced-choice and same-different tasks respectively.

for forced-choice and same-different responses collected from the same subject. A comparison between the tasks can therefore provide useful insights. In our discussion of the following figures, for convenience and clarity of presentation we denote the type of prime as either 'Foil primed' or 'Target primed' across both tasks. However, as noted in the method, 'Foil primed' is somewhat of a misnomer in the case of the same-different task. Again, because this type of prime was matched with the conditions in which the foil was primed for the forced-choice task, we refer to both as 'Foil primed' to make clear as to how we are comparing conditions across tasks. We focus on aggregate measures of performance, averaging across subjects. For presentational clarity once again, we do not include measures of variability in these figures. However, because knowledge of such variability is critical when examining and evaluating empirical data, in Appendix B of our online supplementary material we present boxplots of the distributions over subjects and conditions for both our accuracy and response time results.

To support our preliminary interpretations of our empirical data, we used mixed effects Bayes factor tests (e.g., Rouder, Morey, Speckman, & Province, 2012). To compute these tests, we used the BayesFactor package in R (version 0.9.12-2; Morey & Rouder, 2015), relying on the default priors of the package. Note that these tests serve as a useful heuristic to lend credence to our claims, but we doubt our data meet the assumptions of the standard mixed effects model (i.e., the assumptions of a repeated-measures ANOVA). We partially address this issue by transforming our data to more closely meet the assumptions, but in the end, we base our primary conclusions on our preferred approach: the modeling results from applying the nROUSE and sequential sampling models to the data.

Fig. 6 presents the average accuracy across subjects for each task by prime type and duration. In the figure, triangles represent performance when the target was primed, while filled circles represent performance when the foil was primed. The first panel shows accuracy collapsed over the type of correct response (i.e. whether the correct answer was on the left or right or was 'same' or 'different'). The remaining two panels present accuracy separately for each task, allowing comparisons between performance based on the type of correct response.

There was a strong interaction between the type of priming and prime duration for accuracy: regardless of the type of testing, following a brief prime, accuracy was high in the target primed condition whereas it was below chance in the foil primed condition. In contrast to the pattern at short durations, accuracy was similar for the target primed and foil primed conditions following a longer duration prime. Remarkably, average accuracy was nearly identical in comparisons between forced-choice vs. same-different testing, providing a clear indication that the same perceptual information was used to guide responses for both test formats.

A Bayes factor test supports that the preferred model to account for accuracy performance includes only main effects and an interaction for prime duration and type. To carry out the test, we first applied a correction to address cases where subjects picked the correct choice 100% of the time (this occurred for only two subjects, in one condition out of the 16 in each case). Specifically, we adjusted frequency correct for each subject over each condition by adding 0.5, and we added 1 to the total number of trials (e.g., Hautus, 1995). We then computed the logit of proportion correct over subjects and conditions. Finally, we fit a mixed effects model with a random effect of subject (the null model), and we compared this model against models that included all possible combinations of main effects and interactions for the type of task, prime duration, and prime type. The model with main effects and an interaction of prime duration type, conditioned on our data and the default priors, was 9×10^{89} times more likely than the null model, and was 3.09 times more likely than the next most likely model (relative to the null), which included a main effect of task. According to Kass and Raftery (1995), the latter ratio indicates 'positive' evidence that there is no effect of task, though according to Jeffreys (1961), this ratio is still 'barely worth mentioning'.

The interaction between prime duration and type of priming replicates previous results, although the current study did not find a full crossover to negative priming but rather a large reduction in positive priming. Note that while the nROUSE model predicts a

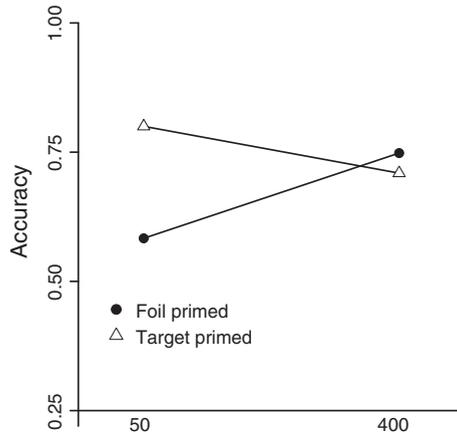


Fig. 7. Accuracy by prime type and duration for the first 20 trials only, averaged over subjects. The triangles indicate the target was primed, while the filled circles indicate the foil was primed.

cross-over in accuracy, for prime durations of 400 ms the predicted cross-over is relatively small (see Fig. 1). Furthermore, unlike previous studies, this study involved much more practice with the task (across 2 days of testing). Therefore it is possible that the lack of cross-over reflects learning on the part of the subjects over multiple sessions. We explored this possibility by looking only at accuracy for the first 20 trials (irrespective of task), aggregating over subjects. Fig. 7 presents the results. For the initial 20 trials, there was in fact a small degree of cross-over, with higher accuracy for the foil primed condition and lower accuracy for the target primed condition. Hence, in the initial stages of the experiment, the observed patterns matched the predictions of the nROUSE model that there would be a small degree of cross-over in accuracy. This finding supports the interpretation that the overall lack of cross-over at long durations reflects gradual learning on the part of the participants.

The mixed effects modeling indicates that the tasks differed little in terms of accuracy performance. Still, an examination of the middle and far-right panels of Fig. 7 suggest that the two test formats had slight differences in terms of response bias. More specifically, although there was no bias for ‘left’ or ‘right’ responses in forced-choice testing, there was a small degree of response bias for ‘same’ responses over ‘different’ responses for same-different testing, which was most pronounced in the short foil-primed condition. When applying a set of mixed effects models testing the main effects and interactions of prime duration and type as well of type of correct answer (‘same’ versus ‘different’), the most likely model had main effects and an interaction for prime type and duration, and included a main effect of type of correct answer (This model was 2×10^{40} times more likely than the null model with only a random effect of subject). However, the mixed effects model was only 2.72 times more likely than a similar model with no effect of type of correct answer; the presence of response bias in the same-different task is inconclusive based on the accuracy data alone.

Fig. 8 presents median response times, averaged over subjects, based on accuracy, prime type, and prime duration. The left-most panel presents performance for the forced-choice task, while the right-most panel presents performance for the same-different task.

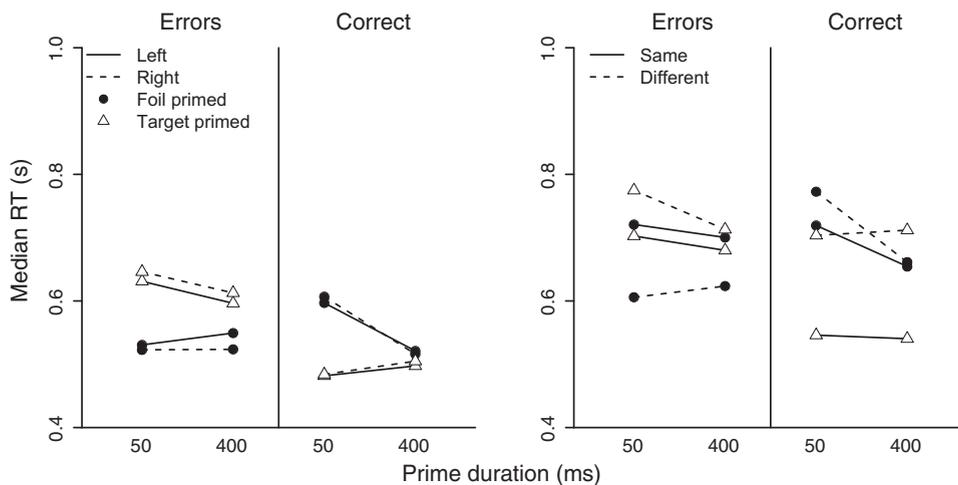


Fig. 8. Median response times (RT) averaged over subjects, based on accuracy, prime type, and prime duration. Triangles indicate the target was primed, while filled circles indicate the foil was primed. Solid and dash lines denote the type of correct response. The left-most panel shows performance for the forced-choice task, while the right-most panel shows performance for the same-different task. Within a panel, median response times for errors are shown on the left, while those for correct responses are shown on the right.

Within a panel, median response times for errors are shown on the left, while those for correct responses are shown on the right.

The figure reveals a complex set of effects for average median response time based on the different conditions and accuracy. As seen in the left panel, response times in forced choice testing were largely unaffected by whether the correct response appeared on the left versus the right. We conducted two Bayes factor tests to evaluate this hypothesis. First, we took the log of the median response times over subjects and conditions. We then split the data up separately for correct and incorrect responses. For each data set, we compared all possible combinations of fixed effects for choice type ('left' versus 'right') and prime duration and type. The most likely model (against a null model with only a random effect of subjects) for correct responses had main effects and an interaction for prime duration and type (with a Bayes factor of 5×10^8). This model was 5.836 times more likely than the next most likely model, which included a main effect of choice type. In contrast, the most likely model for error median response times included only a main effect of prime type (with a Bayes factor of 7×10^8). This model was 9.371 times more likely than the next most likely model that actually included an effect of choice type (in this case a main effect of prime and choice type). Hence, the response time data provides further support for the conclusion that response bias did not play an important role in forced choice testing.

For correct forced choice responses (shown on the right in the first panel), median RTs were the mirror image of accuracy, with faster target primed decisions than foil primed decisions following a short duration prime, but essentially no difference between these conditions following a long duration prime. In contrast, error responses (shown on the left in the first panel) exhibited a main effect of prime type, matching the effect of prime type found in the accuracy data. Errors were faster for foil-primed conditions, and slower for target-primed conditions.

As seen in the second panel, the effects for response times in the same-different task are harder to untangle, presumably because, unlike the situation with forced-choice testing, response bias played an important role in same-different testing. However, regardless of response bias, on average the median response times for same-different testing were slower than forced-choice task. We confirmed this hypothesis via a Bayes factor test, testing a set of mixed effects models for all possible combinations of fixed effects for the type of test and accuracy (collapsing over prime duration and type for simplicity). A model that only included a fixed effect of accuracy was in fact slightly less likely than the null model with only a random effect of subject (with a Bayes factor of 0.5). In contrast, the most likely model had both a main effect of task and accuracy (with a Bayes factor of 1×10^{14}). This result provides further support for the conclusion that forced-choice decisions reflect a race between the target and foil. More specifically, if forced-choice is a race to identify the target versus the foil, then response times are a minimum statistic, and are expected to be faster than single-item response times based on these same target and foil identification latencies (although this supposition depends on the decision process for same-different testing).

An initial clue regarding the same-different decision process comes from noting that the errors when the correct answer was different show the same priming pattern as the errors in forced-choice testing while, at the same time, the correct responses when the correct answer was same show the same priming pattern as the corrects in forced choice testing. This is to be expected if the time to identify the target underlies both correct responses in forced-choice testing and correct 'same' responses in same-different testing, while at the same time, the time to identify the foil underlies both the error responses in forced-choice testing and error 'same' responses in same-different testing. In other words, this follows if rapidly identified words in same-different testing are taken as evidence in favor of a 'same' response. Occasions where subjects gave a response of 'different' in same-different testing (i.e., the errors for same trials and corrects for different trials) may at first appear to defy explanation, although we will return to this later when considering different ways to model the evidence accumulation process that underlies 'different' responses.

For same-different tasks, two common findings for perceptual matching are more frequent errors for 'different' responses, and faster correct 'same' responses relative to correct 'different' responses (Van Zandt, Colonius, & Proctor, 2000). Both effects are present in the current data, but are modulated by prime type and duration. For instance, errors are more frequent for 'different' responses, but only in the short foil prime condition. Correct 'same' responses are faster than correct 'different' responses, but this is most pronounced only for the target primed condition.

The results shown here provide new insights into the nature of task differences for immediate word priming. Both the forced-choice and same-different task exhibit strong positive priming following short duration primes. Critically, across both tasks, this positive priming was greatly reduced or eliminated following long duration primes. There were slight task differences for accuracy measures, but the results with median response time indicate strong differences in performance between the forced-choice and same-different tasks. Therefore, our data provide strong constraints for response time and choice modeling.

6. Modeling results

In this section, we present our estimation results for the nROUSE and diffusion race model. We first discuss our findings on whether the nROUSE model best accounts for the forced-choice accuracy data with three free parameters. We then evaluate how well the diffusion race model fit the patterns of interest with the same-different choice and response time data. Finally, we present our findings on whether the nROUSE identification latencies generated from the forced-choice accuracy data correlate with either drift rates from the same-different task, and whether the threshold estimates support a strategic discounting strategy or not.

6.1. nROUSE results

We used the forced-choice data to obtain parameter estimates for the nROUSE model. We allowed three parameters for the nROUSE model (the noise multiplier, inhibition, and temporal attention parameters) to vary in order to account for individual differences across the behavioral data. We had hypothesized that all three of these parameters were necessary to explain individual

Table 2
nROUSE model comparisons.

Model	Number of parameters	AIC	AIC weight	BIC	BIC weight
NM	1	63.65	0.00	63.04	0.00
D	0	61.69	0.00	61.69	0.00
TA	1	37.11	0.01	36.50	0.00
NM-TA	2	33.75	0.03	32.52	0.03
I	1	32.54	0.06	31.92	0.04
NM-I	2	31.14	0.12	29.91	0.10
I-TA	2	29.88	0.23	28.65	0.19
NM-I-TA	3	28.11	0.55	26.27	0.64

differences in accuracy performance. We tested this hypothesis by comparing, at the group-level, the model with three free parameters against all possible sets of nested models with fewer free parameters. Table 2 presents the results, reporting the models and their associated AIC and BIC values, along with their Akaike and Schwarz weights. The Akaike and Schwarz weights show the relative probability that a model in the comparison set will fit new data from the same data generating process (e.g., Wagenmakers & Farrell, 2004). For conciseness, we abbreviate the parameter labels to NM (Noise multiplier), I (Inhibition), TA (Temporal attention), and D (all parameters fixed to the original default values). Entries in the table have been sorted from the highest values (worst fit) to the lowest values (best fit).

As seen in Table 2, the full model in which all three parameters can freely vary has the lowest AIC and BIC values, and the highest set of weights. Note that while the BIC typically has a higher penalty for complex models compared to the AIC, in this case due to the small number of observations (only a set 4 frequencies), the BIC actually ends up with a smaller penalty (i.e., roughly 1.3 times the number of parameters instead of 2). While the three parameter model has the best fit, two of the nested models still had fairly high Akaike and Schwarz weights. Therefore, we also evaluated how the three parameters capture individual differences by examining the aspects of the data with which the parameters correlated. Fig. 9 shows scatter plots for different aspects of the behavioral data as they relate to each of the free parameters.

Average threshold target durations varied greatly across the 25 subjects, ranging from 38 ms to 113 ms, with a mean of 65 ms and a standard deviation of 21 ms. Differences in the temporal attention parameter were hypothesized to cause the large variations in the threshold target durations. The mean value for this parameter was .96 (minimum of .29, maximum of 2.39, with a standard deviation of .52), and this parameter strongly correlated with individual differences in threshold target duration in the expected manner ($r = -.78$), with lower temporal attention requiring longer target durations to maintain accuracy at 70% (see panel A of Fig. 9).

The degree of positive priming at short prime durations was indexed by the accuracy difference between the target and foil primed conditions. All of the 25 subjects exhibited positive priming but the magnitude of this priming effect differed across subjects (mean priming effect of .51, with a minimum of .23, a maximum of .80 and a standard deviation of .13). As seen in panel B of Fig. 8, the noise multiplier parameter strongly correlated with the short duration priming effect ($r = -.75$), with greater decisional noise serving to reduce the difference between the target primed and foil primed conditions.

Finally, the magnitude of the priming effect at long durations was also indexed by the accuracy difference between the target and foil primed conditions. Subjects varied greatly, with a mean long duration priming effect of .07, with a minimum of $-.21$, a maximum of .39, and a standard deviation of .17). As seen in panel C of Fig. 9, the inhibition parameter strongly correlated with the long duration priming effect ($r = .93$), with greater inhibition producing positive priming rather than negative priming.

Based on the model comparison results at the group-level, and the correlations between the model parameters and empirical measures at the individual level, individual differences in the forced-choice accuracy data appear to reflect differences in the degree of attention subjects had for the target flash, the amount of inhibition they exhibited, and the values for the noise multiplier parameter. Notably, the model results help account for the failure to find a full cross-over from positive to negative priming. Subjects had a higher degree of inhibition compared to the results of Rieth and Huber (2017), perhaps reflecting learning over the multiple sessions (see Fig. 7). We now turn to the estimation results for the diffusion race model.

6.2. Diffusion race model results

After fitting the nROUSE model to the forced-choice accuracy data, we then fit the diffusion race model separately to the same-different choice and response time data. We ultimately wanted to assess whether the perceptual identification latencies (the internal measure of perceptual fluency for the nROUSE model) would predict drift rates, and whether the threshold estimates supported a strategic discounting approach or not. However, it would be pointless to carry out such a comparison if the diffusion race model failed to fit the same-different data. Therefore, in the following section we evaluate, at the group-level, the goodness-of-fit to the choice and response time data for the diffusion race model.

To examine the predictions of the diffusion race model at the group level, we generated model predictions for each subject, and then aggregated over these predictions by taking the mean. We focused on the model predictions for the joint cumulative density, as this provides information across multiple dimensions: model performance can be evaluated based on both the predicted proportion of responses and the predicted response time distributions for each response type. We therefore present figures with aggregate quantile estimates for the joint cumulative response time distribution across subjects for both the observed data and model predictions. These

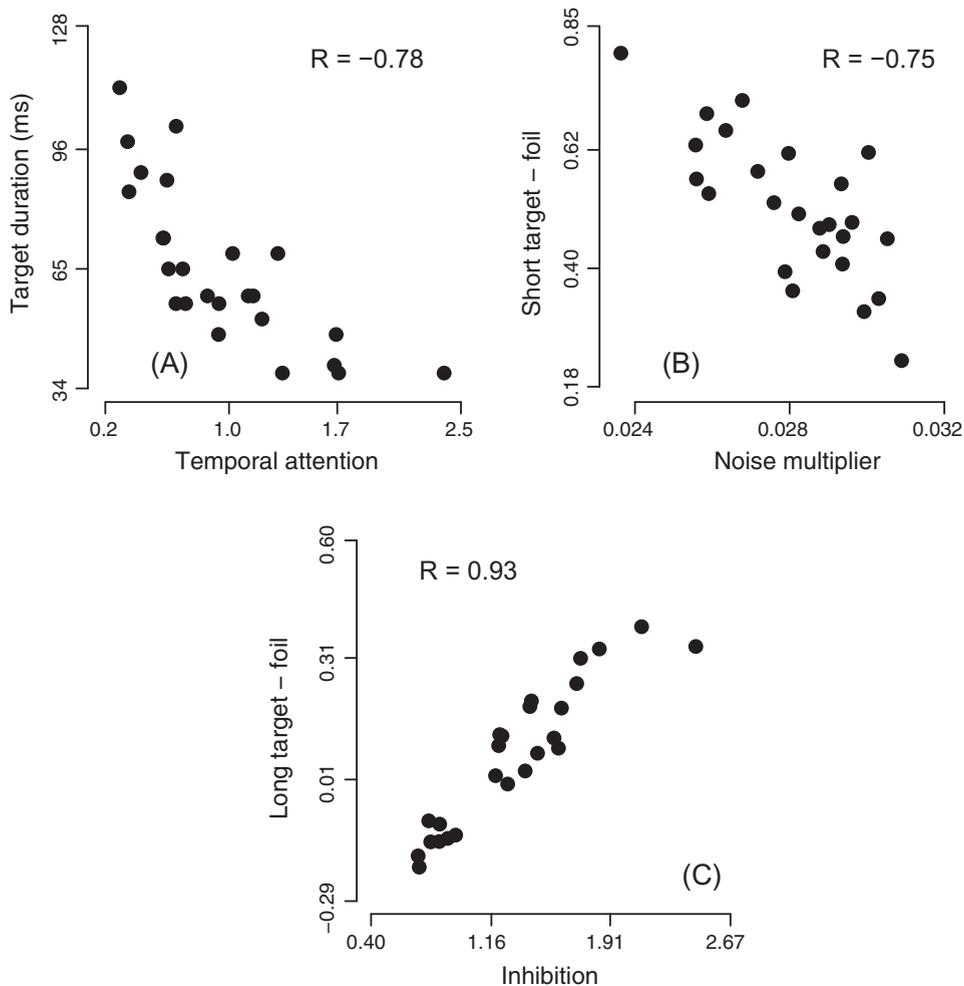


Fig. 9. Scatterplots of the relation between subject-level estimates for nROUSE parameters and aspects of the behavioral data. Panel A compares estimates for the temporal attention parameter to the duration of the target presentation that results in threshold performance. Panel B compares estimates of the noise multiplier parameter to the accuracy difference between the target and foil primed conditions following a short duration prime. Panel C compares estimates for the inhibition parameter to the accuracy difference between the target and foil primed conditions following a long duration prime.

plots were constructed by first determining the response times associated with the cumulative density values of 10, 30, 50, 70, and 90% for each conditional response time distribution (i.e., first for ‘left’ or ‘same’ responses, then for ‘right’ or ‘different’ responses). To get the joint response distributions, the cumulative density values are then weighted by the proportion of responses available for each response time distribution. Finally, on the far right, the average choice proportions are shown to allow interpretations regarding accuracy.

As an example, suppose the model predicts a subject will pick ‘same’ for a proportion of 0.7 out of the total number of trials. Also suppose that the 90% quantile for the response time distribution conditioned on the subject picking the left response is 0.8 s. The value of this point on the x-axis will therefore also be 0.8, and the value on the y-axis would be 0.63, which is the cumulative density value 0.9 weighted by the proportion 0.7. In the following figures, distributions are presented for the type of choice data instead of accuracy in order to properly reflect how we fit the data.

Figs. 10 and 11 present the joint cumulative distribution functions over prime type and duration for when ‘different’ or ‘same’ responses were correct, respectively. As to be expected given the large number of drift rate and threshold parameters, the diffusion race model did an excellent job of capturing the trends in the data. Critically, when the foil was primed for 50 ms, the model captures the dip in performance to below chance levels for when ‘different’ responses were correct, and the much smaller dip in performance leading to at chance performance for when ‘same’ responses were correct. Indeed, the model nicely captures both choice and response time data for when ‘same’ responses were correct. The model exhibits some misfit with response times for the infrequent errors in the target-primed conditions when the ‘different’ response was correct. However, the optimization routines were unlikely to select parameters to best-fit these errors, as their infrequent occurrence downweights their impact on the sum of the log-likelihoods.

To more closely evaluate the goodness-of-fit of the diffusion race model to the same-different data, we present boxplots showing the distribution of residuals for key statistics. We focus on the difference between the observed and predicted median response times

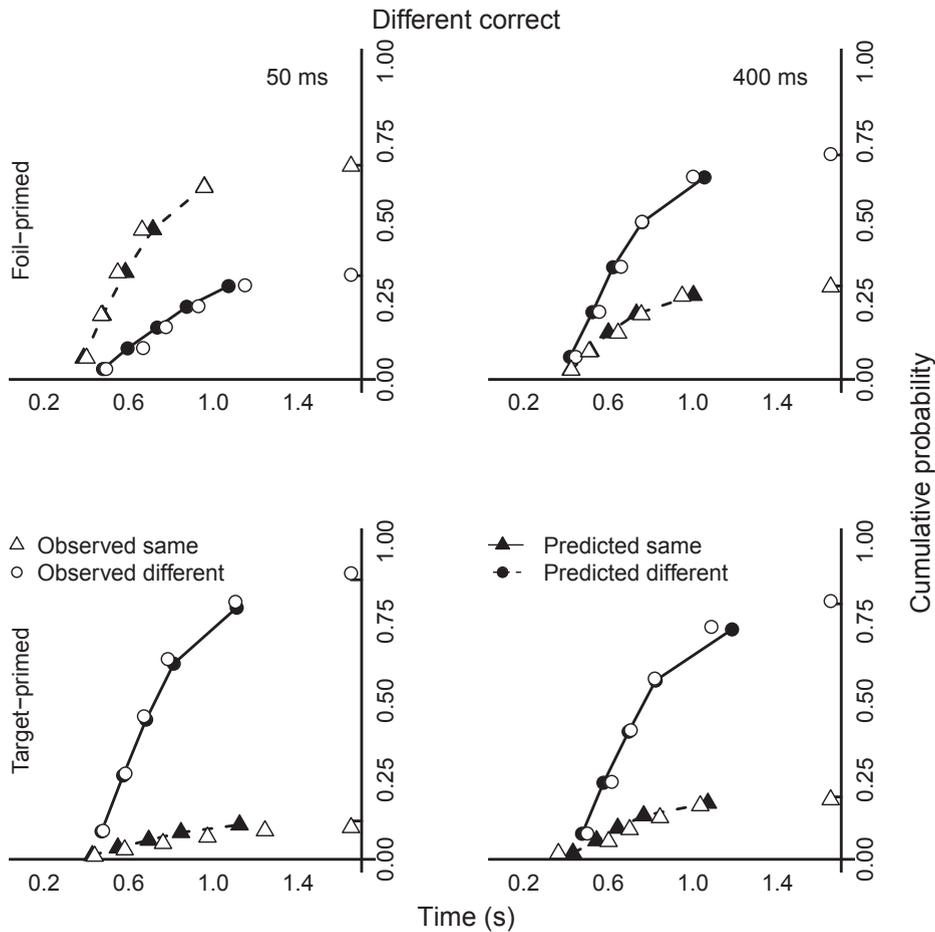


Fig. 10. Joint cumulative distribution functions for choice and response time over prime type and duration when picking ‘different’ was correct. The y-axis shows the cumulative probability for picking ‘different’, while the x-axis shows time in seconds. Triangles represent distributions for ‘same’ responses, while circles represent distributions for ‘different’ responses. Filled points connected by lines represent the diffusion race model predictions. Points on the right-hand axis denote the average choice proportions. Distributions are represented by the 10, 30, 50, 70, and 90% response time quantiles. The top panels are for foil-primed conditions, the bottom panels are for target-primed conditions. The left panels are for short prime durations, the right panels are for long prime durations.

for both correct and error responses, as well as the difference between the observed and predicted proportion correct. Fig. 12 shows the distributions for these residuals over prime type and duration, as well as the type of correct response. The endpoints for the boxplots represent the inter-quartile range (the difference between the third and first quartiles), while the box shows a bootstrapped 95% confidence interval around the average residual, denoted by the innermost bar. We estimated the 95% confidence interval via a resampling technique. For a given type of correct response, as well as prime type and duration, we sampled with replacement the set of residuals for our 25 subjects. We then recomputed the average over these residuals, and repeated this procedure 10,000 times.

To represent differences in uncertainty based on the average number of trials underlying a given set of residuals, we plot a 95% confidence interval denoting the range of residual values predicted from sampling error alone (the gray dashed lines in the figure). To generate these intervals, we first averaged the best-fitting parameters over subjects. We then computed a parametric bootstrap, simulating observations from the diffusion race model. We simulated observations equal to the number of trials per condition for a single subject (i.e., 80 trials per condition). We then computed the difference between the observed test statistic (e.g., the median response time for correct responses) and the predicted value based on the generating parameters (found by integrating over the density function). We repeated this procedure 1000 times, and then obtained estimates of the 2.5% and 97.5% quantiles for the distribution of residuals in order to compute the 95% confidence interval. This intervals provide useful information on the uncertainty we would expect from sampling error alone. For instance, in Fig. 11 the intervals widen for median error response times when the target was primed, as errors were much less frequent.

Fig. 12 allows us to quantify where the diffusion race model best fits the observed data, and for which aspects of the data it struggles to fit. The model fits the proportion correct quite well, with only some very slight underestimation of proportion correct in the target primed conditions when ‘different’ responses were correct. The model captures response time data fairly well, an impressive feat considering the small number of trials per condition. It tends to underestimate median response times for correct responses with the short duration prime when picking ‘different’ is correct, and overestimate response times for the short target prime

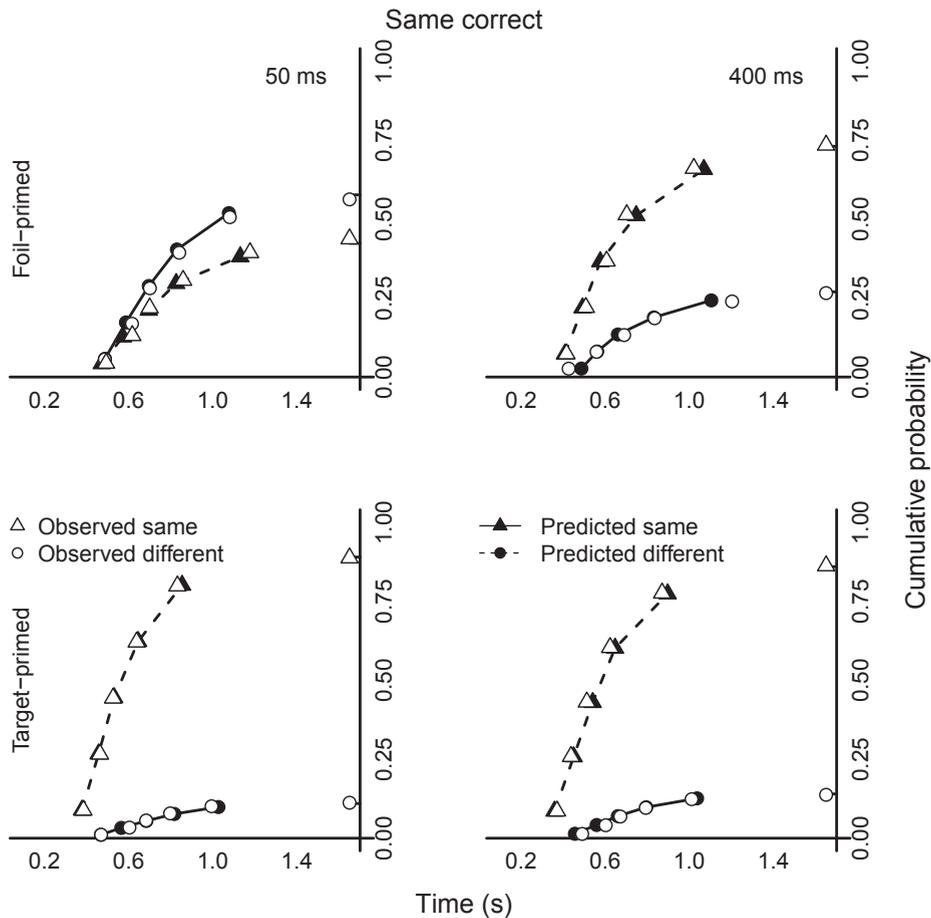


Fig. 11. Joint cumulative distribution functions for choice and response time over prime type and duration when picking ‘same’ was correct. The y-axis shows the cumulative probability for picking ‘different’, while the x-axis shows time in seconds. Triangles represent distributions for ‘same’ responses, while circles represent distributions for ‘different’ responses. Filled points connected by lines represent the diffusion race model predictions. Points on the right-hand axis denote the average choice proportions. Distributions are represented by the 10, 30, 50, 70, and 90% response time quantiles. The top panels are for foil-primed conditions, the bottom panels are for target-primed conditions. The left panels are for short prime durations, the right panels are for long prime durations.

condition when picking ‘same’ is correct. The model also tended to underestimate median response times for infrequent errors, particularly when the target was primed and the ‘different’ response was correct. Overall, the diffusion race model had a high level of goodness-of-fit to the data, allowing us to examine the relationship between the parameters and the perceptual identification latencies with confidence.

6.3. Relation between models

We successfully fit the nROUSE model to the forced-choice accuracy data and the diffusion race model to the same-different choice and response time data. We were also able to generate predicted perceptual identification latencies for each subject by simulating the nROUSE model with its associated best-fitting parameters. With the predicted latencies from the nROUSE model and parameter estimates from the diffusion race model, we could then address our two main questions of interest: (1) can the nROUSE model capture individual differences in both accuracy and response time for the same-different task via a relationship with drift rates, and (2) do the long prime duration conditions from the same-different task support a strategic discounting hypothesis or the nROUSE model assumption that perceptual disfluency drives priming effects?

We assessed the relationship between the perceptual identification speeds (the inverse of the latencies) and the drift rates for picking ‘same’ via three statistics: Pearson’s R, Kendall’s tau, and Spearman’s rho. Pearson’s R is the most well-known correlation coefficient, expressing the degree of linear relationship between two variables (where 1 represents a perfect positive linear relation, and -1 represents a perfect negative linear relation). Kendall’s tau measures the ordinal association between two variables (where 1 represents when all pairs of observations have equivalent rank orders, and -1 represents when all pairs of observations have completely opposite rank orders). Spearman’s rho is similar, measuring how well the relation between two variables can be described by a monotonic function. Kendall’s tau and Spearman’s rho are more robust than Pearson’s R, but if there is indeed a linear relation between the variables, they can be less powerful.

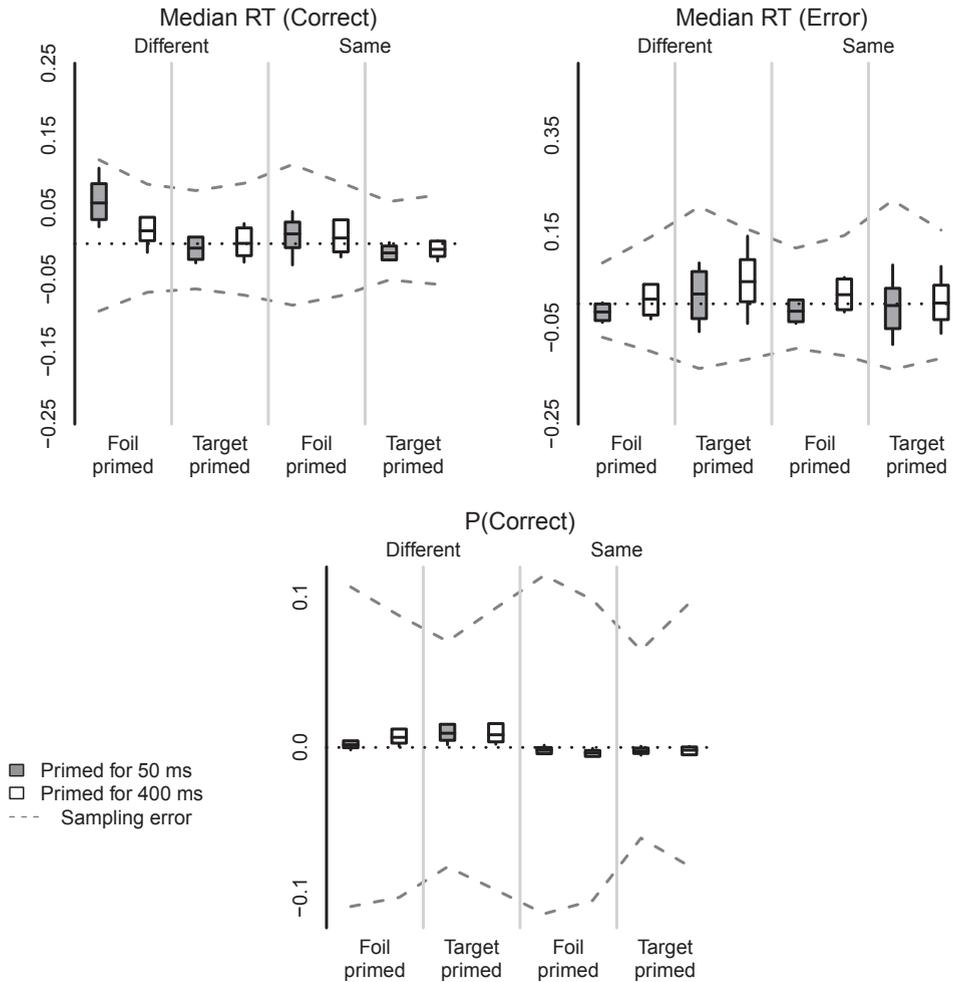


Fig. 12. Boxplot distributions of the residuals for the median correct and error response times, as well as the proportion correct, over prime type and duration as well as the type of correct response. Dashed lines denote the expected degree of sampling error for each condition. Boxplots denote, from outward to inward, the inter-quartile range, \pm one standard error, and the average residual value over subjects.

For each subject, we first took the inverse of 8 perceptual identification latencies from the nROUSE model. We did so for ease of interpretation, as fluency is typically conceptualized as speed instead of latency. Under this transformation, if perceptual fluency underlies the decision process, we would expect a positive relation between drift rates and the inverse identification latencies. In contrast, with strategic discounting, we would not predict perceptual disfluency at longer prime durations. Instead, negative priming would occur because subjects would purposefully discount the prime by picking the less fluent choice. Therefore, we would expect no correlation between drift rates and identification speeds from the nROUSE model, as the mismatch between predictions for the short and long prime durations would cancel out any relationship. After standardizing all variables (which simplifies estimation and plotting of the results), we computed the three correlation coefficients, assessing the relationship between the 8 perceptual identification latencies from the nROUSE model and the associated 8 drift rates for ‘same’ responses. For completeness, in [Appendix B](#) of the online supplementary materials, we present the scatterplots for the relationship between the nROUSE latencies and drift rates for each of the 25 subjects.

Fig. 13 presents the distributions for each of the three correlation coefficients (Pearson’s R, Kendall’s tau, and Spearman’s rho) for drift rates. The y-axis reports the magnitude of the correlation coefficients, and the x-axis indicates which set of distributions is associated with which correlation coefficient. We used Gaussian kernel density estimation via a base function in R to determine the distribution of coefficients. As predicted, regardless of the type of correlation coefficient, all subjects exhibited moderate to strong positive correlations between the inverted identification latencies and drift rates. This finding supports the nROUSE model assumption that priming effects reflect perceptual fluency at short prime durations, and perceptual disfluency at longer prime durations.

Fig. 14 presents the median values over subjects for the drift rates and threshold estimates, along with their inter-quartile ranges. The left panel has the median values for drift rates, with drift rates for picking ‘same’ marked with white squares and drift rates for picking ‘different’ marked with black circles. The x-axis gives the condition for each pair of drift rates, with each line indicating the

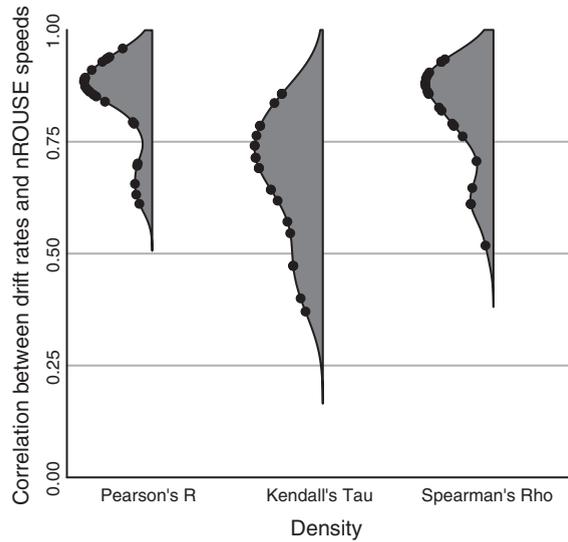


Fig. 13. Distributions of correlation coefficients between inverted perceptual identification latencies and diffusion race model parameters for the ‘same’ response. The y-axis indicates the magnitude of the correlation coefficient. Dots represent the specific correlation coefficient and its associated density for a given subject.

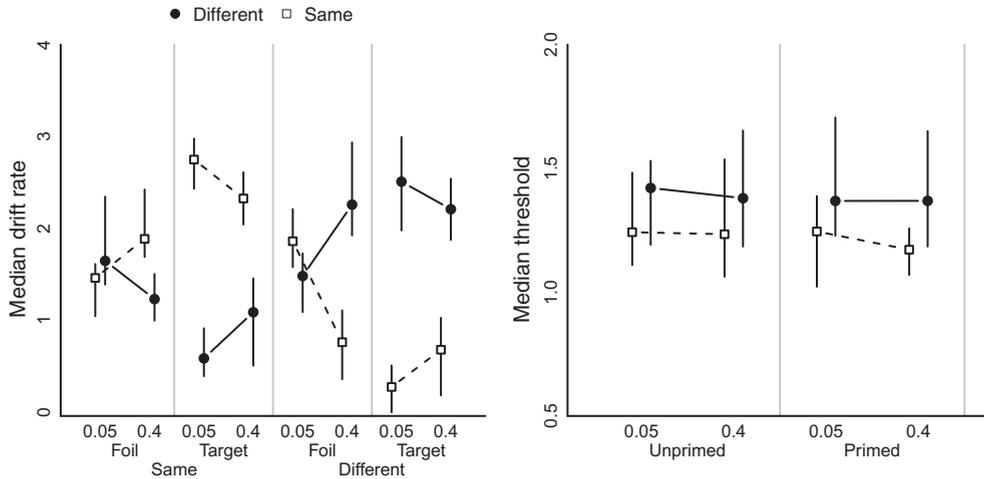


Fig. 14. Median values over subjects for the drift rate and threshold parameter estimates obtained from fitting the diffusion race model to the same-different task. Error bars represent the inter-quartile range. White squares represent parameter estimates corresponding to picking the ‘same’ response, while black circles represent parameter estimates corresponding to picking the ‘different’ response. The y-axis represents the value of the parameter estimates, while the x-axis shows the condition for a given pair of estimates. The labels of the x-axis for drift rates present the prime duration, prime type, and type of correct response for each condition, respectively. The labels of the x-axis for thresholds present the prime duration and whether the onscreen choice was primed or not. Solid lines once again denote thresholds for picking ‘different’ while dashed lines denote thresholds for picking ‘same’.

prime duration, type, and type of correct response respectively. The right panel structures the threshold estimates based on prime duration and whether the onscreen choice was primed or not. Solid lines once again denote thresholds for picking ‘different’ while dashed lines denote thresholds for picking ‘same’.

Critically, the drift rates exhibit the cross-over pattern originally predicted by the nROUSE model. Drift rates are highest for the correct response when the target is primed for 50 ms, and lowest when the foil is primed for 50 ms. In contrast, drift rates increase when the foil is primed for 400 ms, but decrease when the target is primed for 400 ms. Again, as evidenced earlier by the distribution of correlations, this pattern of drift rates reflects a switch from perceptual fluency to perceptual disfluency, as predicted by the nROUSE model. Drift rates for the ‘different’ response also exhibit an interesting pattern in which they provide a ‘mirror-image’ of the drift rates for ‘same’ responses. If the drift rate for the ‘same’ response is high, the drift rate for the ‘different’ response is low, and vice versa.

The panel on the right indicates that, for threshold estimates, subjects on average had a clear bias against picking ‘different’, replicating previous findings for the same-different task (e.g., Van Zandt et al., 2000). However, no obvious effect of bias towards or against the primed choice is present. To confirm this, as noted earlier, we applied two mixed effects linear regressions with categorical predictors to the threshold estimates. The first model had an intercept term and a response bias effect, capturing the bias against picking ‘different’. The other regression model incorporated a second categorical predictor, coded to reflect a planned contrast

capturing strategic discounting. Specifically, the predictor was fixed to 0 for short prime durations, but set to + 1 for the thresholds reflecting unprimed choices, and –1 for thresholds reflecting primed choices. Negative values for this regression coefficient, then, indicate strategic discounting, with higher threshold values for the primed choice. We also allowed there to be an interaction between the strategic discounting and the response bias predictors. Both models included a random effect for subjects. These regression models allowed us to address our second question on whether subjects exhibited strategic discounting for long-duration primed choices. First, we assessed whether the single predictor for response bias was sufficient for capturing threshold estimates using a Bayes factor test. However, evidence was equivocal: the ratio of the strategic discounting model over the response bias model was 0.92. Next, we examined the credible intervals for the regression coefficient posteriors of the strategic discounting model. Critically, we found that the coefficient for strategic discounting overlapped with zero, and went in the wrong direction, with a 95% credible interval of –0.002 to 0.056.

7. Discussion

A wide variety of theories rely on the concept of fluency (the meta-cognitive evaluation of speed of processing). With immediate repetition priming at short durations, several theories posit that it is perceptual fluency that drives priming effects. The primed choice is processed faster, thereby appearing more fluent. With perceptual identification, these theories posit that subjects use a heuristic when picking between alternatives: a more fluent choice is assumed to be so because it must have been recently viewed (i.e., the briefly flashed target stimulus). However, at long prime durations, subjects on average pick the unprimed choice more often. Most theories explicitly or implicitly assume that this form of negative priming reflects strategic discounting on the part of the subject. Subjects become aware of long duration primes, and strategically discount the influence of the prime, purposely choosing the less fluent choice. The nROUSE model instead argues that subjects experience perceptual disfluency: due to neural habituation, the primed choice ends up appearing less fluent, and the unprimed choice more fluent. Hence, even with long duration primes subjects continue to rely on the heuristic of fluency to drive their judgments, and do not engage in strategic discounting.

The nROUSE model and its assumptions of perceptual fluency and disfluency with increasing prime duration fill an important niche in the literature on priming and fluency. The basis of the nROUSE model, the Bayesian ROUSE model, is a theory of temporal parsing in which the perceptual system infers whether the currently available perceptual information arises from the most recent visual object (i.e., the target) or whether it instead reflects lingering information from a prior visual object (i.e., the prime). As compared to the Bayesian ROUSE model, the neural nROUSE model specifies how the brain might approximate this Bayesian inference process through neural habituation. Because the nROUSE model is a specific neural implementation of the more abstract Bayesian model, it necessarily makes auxiliary assumptions; assumptions that if untrue would call into question previous applications of the nROUSE model. The assumption that subjects rely on perceptual fluency at both short and long durations is a critical component of the nROUSE model, as without it the model would make incorrect predictions for forced-choice accuracy at long durations.

We therefore compared the strategic discounting and perceptual fluency/disfluency theories, using a novel approach. We first had subjects complete both forced-choice and same-different testing during perceptual identification with immediate repetition word priming. Priming effects for accuracy remain unchanged when one changes the mode of testing from forced-choice to same-different (see Fig. 6). Hence, according to the nROUSE model, perceptual fluency should also underly the decision process in same-different testing. While the model is agnostic about processing for ‘different’ responses, it posits that processing for ‘same’ responses should match the processing in the forced-choice task across both short and long prime durations. In contrast, the strategic discounting approach posits that while subjects rely on fluency during short duration primes, they strategically become biased against primed choices at long prime durations.

To untangle the rate of perceptual processing from decisional biases, we cannot analyze response times and accuracy separately; we must consider their joint distribution. To do so, we applied a sequential sampling model, the diffusion race model, to our data. We converted choice and response time into measures of drift rate (i.e., the rate at which perceptual evidence accumulates) and threshold values (the amount of evidence needed for a decision, capturing biases and response caution). This allowed us to address our main question: is the transition from positive to negative priming due to a switch from perceptual fluency to disfluency (as predicted by the nROUSE model), or due to strategic discounting on the part of the subject? We tested these contrasting accounts by (1) assessing the correlation between predicted identification latencies (generated from fitting the nROUSE model to forced-choice accuracy) and drift rates from the same-different task, and (2) examining if there was any evidence of bias against the primed choice based on threshold estimates from the same-different task.

We successfully fit the nROUSE model to each individual’s forced-choice accuracy data. We therefore could generate predictions for the speed of identification for choice alternatives in the visual system, a direct measure of perceptual fluency. We also successfully fit the diffusion race model to the choice and response time data from the same-different task. We were able to obtain drift rate estimates for picking ‘same’, and threshold values for the primed and unprimed choices. In other words, we extracted the identification speeds and diffusion race model parameters from separate tasks, with differing response time results. Despite this, across subjects the identification speeds strongly predicted drift rates, as predicted under the nROUSE model’s assumption of perceptual fluency/disfluency driving the effects. Furthermore, a mixed effects regression model applied to threshold estimates revealed that additional bias beyond a bias against responding ‘different’, rather than strategic discounting, in fact reflected a bias towards the prime choice at long prime durations. Moreover, the pattern of drift rates across short and long prime durations reflected higher rates for primed choices at short durations, and lower rates for primed choices at long durations, as predicted by the nROUSE model. In other words, subjects did not strategically discount the influence of the primed choice at long durations. Therefore, we found strong support for the assumptions of the nROUSE model. More generally, these results also support the claim that brief prime presentations

increase perceptual fluency for primed items whereas longer prime presentations eliminate, or even reverse this fluency, producing a perceptual disfluency for primed items.

7.1. Fluency

As noted in the introduction, the nROUSE model is not the first to use fluency to explain the immediate effect of a short duration masked prime (for a recent study demonstrating the role of fluency in masked priming, see [Bodner, Johnson, & Masson, 2015](#)). Perhaps the earliest such account comes from spreading activation theory (e.g., [Anderson, 1983](#)), which remains a dominant explanation of word priming. On this account, presentation of the prime causes a spread of activation that pre-activates a related target word in advance of the target word's appearance. In other words, target identification is easier (more fluent) because it is already partially active. However, this account does not specify why priming should reverse for a long duration prime.

Other fluency-based priming accounts can explain this reversal. For instance, in the context of priming memory responses, fluency plays a key role in fluency-attribution theory (e.g., [Jacoby et al., 1989](#); [Jacoby & Whitehouse, 1989](#)). The core idea of this theory is that the feeling of familiarity is based on an attributional process. This process may include other influences other than memory retrieval, for instance if processing fluency is used as a heuristic to guide recognition judgments. Fluency is a useful heuristic here because prior learning for an item should speed up subsequent processing of that item. However, if other factors boost fluency, such as the presentation of a prime, people may misattribute the increased fluency as an indication of prior experience. As it applies to briefly presented primes, this account is remarkably similar to the nROUSE model. However, there is a critical difference between the nROUSE model and fluency-attribution theory: Although both accounts assume that brief primes increase perceptual fluency, leading to performance gains for primed items, fluency-attribution theory assumes that perceptual fluency still exists in response to longer duration primes. Thus, only decisional (i.e., attributional) factors counteract the influence of fluency. In contrast, the nROUSE model assumes that perceptual fluency is directly reduced by perceptual habituation, rather than counteracted in the decision process.

A related theory of fluency and familiarity is discrepancy-attribution theory (e.g., [Whittlesea & Williams, 1998, 2000, 2001a, 2001b](#)). Whittlesea and Williams posit that it is not fluency per se that determines familiarity, but instead the discrepancy between the expected and experienced degree of fluency. In other words, if a test stimulus is surprisingly fluent, the subject assumes that this unexpected fluency reflects recent experience with the stimulus. In the context of immediate word priming, short duration primes produce perceptual fluency, but provided that the prime duration is brief, the subject may not know to adjust their expectations regarding fluency, leading to a bias for the primed choice. Long duration primes also increase fluency, but this increase in fluency is expected, so there is no discrepancy, which eliminates or reverses the bias. Thus, as with fluency-attribution theory, this account assumes that the elimination of priming with a longer prime duration is a decisional factor rather than perceptual habituation.

[Jacoby and Whitehouse \(1989\)](#) applied fluency-attribution theory to immediate word priming in the context of recognition memory. In their study, subliminal prime presentations biased subjects to judge a primed test word as previously studied, regardless of whether the test word was or was not studied. In contrast, when the prime was easily seen, this bias effect was reversed. To determine whether the change from positive to negative priming for the immediate priming of episodic recognition responses reflects strategic discounting, [Huber, Clark, et al. \(2008\)](#) examined this effect across a series of five experiments, reporting evidence against a decisional explanation of this priming change (see also [Higham & Vokey \(2000\)](#)). For instance, Experiment 2 compared prime words presented for 100 ms to ones that were presented for several seconds. Similar to the current perceptual identification results, short duration primes produced a positive priming effect (but in this case it was a tendency to claim that the primed test word was previously studied, rather than presented as a briefly flashed target). Also similar to current results, this effect was eliminated following longer duration primes. Critically, in terms of differentiating between the different theories, the short duration primes were either followed by the episodic recognition task or by a prime identification task (i.e., forced-choice to determine which word was the prime). This manipulation occurred randomly across trials such that at the time when the prime appeared, the subject could not know whether they would be asked to identify the prime or whether the prime was simply a precursor to the recognition test. Even though average accuracy to identify the 100 ms prime was 96%, there was a large positive priming effect for the recognition task. Thus, in this case, the elimination of priming when comparing a brief versus a long duration prime could not be attributed to a change in conscious awareness for the prime's identity, considering that subjects were clearly aware of the prime's identity in the short prime duration condition.

The study of [Huber, Clark, et al. \(2008\)](#) provided evidence against the proposal that the change from short to long duration primes reflects a change in prime awareness, supporting the claim that the change with increasing prime duration is an automatic process. Nevertheless, it is conceivable that subjects were aware of both brief and long duration primes and yet elected to use this knowledge differently in each condition. Another way to address this question is with manipulations that are predicted to change the nature of priming for long duration primes. More specifically, both the Bayesian and neural ROUSE models predicted that the negative priming often seen for repetition priming following a long duration prime can be made to flip to positive priming with particular manipulations. These manipulations are always made in a within-subject manner, randomly occurring across trials, making it unlikely that the results reflected a change in the decision strategy. The models made three key predictions for long duration primes, all of which were confirmed: (1) for repetition priming, negative priming with dissimilar choice words but positive priming with similar choice words ([Huber et al., 2001](#); [Huber, Shiffrin, Lyle, & Quach, 2002](#)) (2) negative priming for repetition priming, but positive priming for moderate orthographic priming ([Huber et al., 2002](#)) or for semantic priming ([Rieth & Huber, 2017](#)) and (3) for repetition priming, negative priming with targets presented at threshold, but positive priming when no target was flashed at all ([Huber et al., 2002](#); [Weidemann, Huber, & Shiffrin, 2008](#)). All of these effects occurred for long duration primes and yet in all experiments short duration primes unflinchingly produced positive priming. Thus, if negative priming following a long duration prime is a decision strategy, one

would need to assume that subjects not only choose to apply this strategy only to long duration primes, but more specifically choose to apply this strategy only to the case of repetition priming in situations of sufficiently strong target evidence and when choosing between dissimilar choices. There is no obvious reason why subjects should adopt such a complicated decision strategy, particularly when placed in a situation where the primes were truly non-diagnostic (i.e., just as likely to indicate the wrong answer as the correct answer), with trial by trial feedback reinforcing the non-diagnosticity of the primes in all conditions. Finally, it should be noted that these truly were a priori predictions (the Bayesian ROUSE model simulations were run in advance of data collection).

Aside from response time modeling and the above mentioned predictions of the ROUSE model, another way to address the role of decision strategies is by explicitly changing the decision strategy. Using the same priming paradigm as in the current study, Weidemann et al. (2008) did this by comparing the typical non-diagnostic situation to conditions where the prime was made diagnostic such that knowing which word was primed could be used to indicate the correct answer even in the absence of any perceptual information from the briefly flashed target word. This was achieved by varying the proportion of target primed versus foil primed trials (rather than having equal occurrences of these two conditions). Across four experiments, whenever prime diagnosticity manipulations affected priming, there was nevertheless a change in priming when comparing the short and long duration priming conditions such that short duration priming was positive and long duration priming reduced, eliminated, or reversed the priming effect. In other words, diagnosticity shifted the overall bias to choose or not choose primed words, but it did not affect the relative difference between the short and long duration priming conditions.

This prior work suggests that the change with increasing prime duration is not a change in decision strategy. However, the prior work never determined whether the positive priming effect with a short duration prime reflects increased fluency, as assumed by spreading activation theory, fluency-attribution theory, discrepancy attribution theory, and the nROUSE model. Putting aside the debate regarding decision strategies for easily seen long duration primes, the current study supports this assumption held by all of these models as regards short duration primes.

Of course, there are accounts of immediate word priming that do not involve perceptual fluency. For instance, compound-cue theory (Ratcliff & McKoon, 1988) posits that the prime and the target form a compound cue that is used to probe memory. A non-linear familiarity value can be calculated for this cue (e.g. the strength of the relation between internal concepts for the prime and target can be multiplied and these products can then be summed together) and used to render a judgment. Therefore, if a prime is related to the target or to common concepts, familiarity is boosted without needing to assume that the prime directly affects target processing. In this way, the theory captures positive priming at short durations. It is less clear how compound-cue theory could explain the change from positive to negative priming with increasing prime durations. In related work, Ratcliff and McKoon (2001) proposed a multinomial model of these effects in which the subject first chooses to either reject the primed alternative or not (discounting) on each trial. If subjects do not discount the primed choice, they then may or may not mistake the prime for the target (source confusion). This simple model adequately describes the change in accuracy with increasing prime durations as reflecting the relative balance between source confusion and discounting. However, because this model does not make any representational assumptions, it does not make a priori predictions regarding the roles of choice word similarity, target duration, or type of priming, all of which were found to turn negative priming into positive priming for long duration primes. Furthermore, because it does not make any processing assumptions, the multinomial model does not make any predictions regarding perceptual fluency or response times.

Norris's (2006) Bayesian reader model also accounts for priming without resorting to perceptual fluency. In this model, subjects behave like rational Bayesian decision makers. They combine prior information (i.e. word frequencies) with current perceptual evidence to update hypotheses (e.g. whether a target is a word or non-word). Norris and Kinoshita (2008) present a version of the model for masked priming in a lexical decision task. Perceptual evidence accumulates in a noisy fashion according to a random walk, and each evidence sample is used to update the priors for the hypotheses. A likelihood ratio is then computed, indicating the strength of evidence for a particular hypothesis. While the prime and target provide independent sources of evidence, Norris and Kinoshita note that, like the original ROUSE model, subjects fail to treat the prime and target as distinct events. In simulations, evidence is sampled from the prime for a set duration, and then from the target. As noted by Norris (2006) the Bayesian reader model is highly flexible and applicable to a wide array of tasks. Furthermore, the model may be simulated to provide predictions on both choice probabilities and response time distributions. For immediate word priming in a perceptual identification task, the two hypotheses being evaluated would be which of the choice words matched the briefly flashed target. However, the variant of the Bayesian reader model proposed by Norris and Kinoshita (2008) was not intended for visible primes, as the authors argue the mechanisms for short-duration (masked) priming are distinct from long-duration (visible) priming. The mechanism they used for masked priming would not work if simply applied as is to long-duration primes: as prime duration increases, the evidence from the prime would overwhelm the evidence from the target, producing even stronger positive priming rather than negative priming. Still, the Bayesian reader model could potentially account for both masked and visible primes after specification of appropriate mechanisms for the decision process and long-duration priming.

The alternative theories just discussed account for priming effects based on a rational analysis of word reading or on abstract cognitive processes. In other words, these accounts explain priming at the computational or algorithmic levels, using Marr's (1982) terminology. Furthermore, while several of these accounts appeal to perceptual fluency, said fluency remains a latent, unobservable process within these models. The nROUSE model therefore fills an important gap in the literature. The model proposes explicit neural mechanisms that account for the priming effects via the rate-coded synaptic depression equation and the layers for visual, orthographic and lexical-semantic representations. Furthermore, the model allows for a direct measure of fluency via the assumption that perceptual fluency is based on the time to peak activation. The current results support the fluency assumption made by the nROUSE and other models regarding positive priming with short duration primes. The current results also support the unique prediction of the nROUSE model that the change to a long duration prime reduces or reverses this fluency.

Aside from confirming predictions of the nROUSE model, the current study suggests that the drift rates derived from application of a sequential processing decision model could be interpreted as a measure of fluency, or at the very least a useful proxy for fluency. Thus, the methodology we developed in the current study may be useful for measuring fluency in a variety of situations and paradigms, including same-different tasks, as considered next.

7.2. Linking forced-choice and same-different tasks

Same-different tasks can be challenging to fit with decision models because they often produce two types of effects: first, average response times for correct ‘same’ responses tend to be faster than those for ‘different’ responses, and second, incorrect ‘different’ responses tend to be more frequent than incorrect ‘same’ responses (Van Zandt et al., 2000). Note that both these effects are present in our data, though modulated by prime duration. Our study adds an interesting new layer as we tested whether performance could be linked between forced-choice and same-different testing when we assessed whether the perceptual identification latencies from the nROUSE model (obtained from fits to forced-choice accuracy data) could predict drift rates and thresholds (obtained by fitting the diffusion race model to same-different choice and response times). Therefore, our results have implications how forced-choice and same-different tasks may be linked, and indeed on how decisional processing for same-different testing works.

As a useful contrast, consider alternative accounts of same-different judgments, particularly in regards to the assumed perceptual match process. Krueger’s (1978) noisy operator theory posits that noise perturbs stimulus features in the system, so differences can be perceived even with identical stimuli. However, there will be fewer differences on average for matches compared to mismatches. It is assumed that a comparison process is carried out with each glance or pass over the onscreen stimulus, counting the number of mismatches between features. If there are enough mismatches, a ‘different’ response is made, but if there are too few mismatches, a ‘same’ response is made. If there is an inconclusive number of mismatches, another pass is made and the mismatches are recounted. In cases in which Ratcliff’s diffusion model has been applied to perceptual matching (e.g., Ratcliff, 1981; Van Zandt et al., 2000), it is proposed that a measure of overlap between the target percept and the test stimulus is computed and compared against a critical value. If the overlap is less than the critical value, evidence accumulates for a ‘different’ response, whereas if the overlap is greater, evidence accumulates for a ‘same’ response. Alternatively, consideration of response times in same-different tasks can also be eschewed completely. For instance, Hochhaus and Johnston (1996) used a signal detection theory approach to analyze priming data in a same-different task.

Our approach instead posits a race between processes assessing the match or mismatch, respectively, between the target percept and the onscreen test stimulus, similar to the assumptions Van Zandt et al. (2000) used in fitting the Poisson race model to same-different data. This assumption is distinctly different compared to the noisy operator model, which assumes evidence is based on the mismatch between features only. The inclusion of a match process was necessary in the current situation to make the link with forced-choice responding because we assumed that the forced-choice decision is based on a race between the match processes alone (as per the nROUSE model). When we fit the diffusion race model to the same-different data, we made no assumptions regarding the mismatch process; processing of evidence for perceptual differences was represented by separate drift rates and thresholds for each condition. Fig. 13, however, reveals an interesting pattern: the drift rates for the mismatch process appear to mirror the drift rates for matches, such as might occur if these processes are in direct competition with each other, as seen with dynamic inhibition between the operations of each process in the leaky competing accumulator model (Usher & McClelland, 2001).

Our models and data imply that subjects can rely on fluency, the speed with which an onscreen alternative is identified, to render both ‘same’ and ‘different’ judgments, rather than direct evaluations of the degree of match or mismatch. Subjects may assume that a quickly identified choice reflects a match, and a slowly identified choice reflects a mismatch. Future work could explore this hypothesis via further manipulations of the degree of fluency for stimuli in same-different tasks (e.g., manipulating the size or visual quality of the stimuli). Furthermore, our results indicate that the same measure of fluency that a subject uses to render a ‘same’ judgment can also be used to carry out forced-choice comparison, as evidenced by the strong correlation between the nROUSE latencies derived from the forced-choice task and the drift rates derived from the same-different task. In other words, our results are in accord with meta-cognitive literature (e.g., Alter & Oppenheimer, 2009) arguing that fluency, the internal measure of speed of processing, can have a larger impact on a variety of judgments than the actual content being processed.

7.3. The non-decision component

When applying sequential sampling models to priming data, the prime can impact not only drift rates and threshold parameters, but also the non-decision component. As indicated by its name, the non-decision component is an estimate of the proportion of the response time that does not reflect the decisional process. This component is a summation of multiple ancillary processes, such as the time for the motor response and the time it takes the perceptual system to process the on-screen stimuli. Past work has implicated the non-decision component with identity primes (e.g., Gomez, Perea, & Ratcliff, 2013).² Assumptions that priming impacts the non-decision component are in fact in accord with the nROUSE model, given that it models perceptual encoding. In theory, the nROUSE model could potentially be used to separate perceptual encoding from the motor process. Unfortunately, our data are not well suited to testing any assumptions regarding the non-decision component as we have too few observations (only 80 trials per condition) to obtain any accurate measurements of the fastest response times needed to accurately estimate this parameter. For this reason, we

² We’d like to thank an anonymous reviewer for raising this point regarding the non-decision component and reminding us of the related literature.

assumed only a single parameter governing the non-decision component. However, we still attempted to capture the predictions of the nROUSE model that the encoding times differ across priming conditions by having the single parameter reflect a proportion weighting the fastest response time for each condition of interest.

7.4. Conclusions

The nROUSE model is a general account of perceptual dynamics, explaining how neural habituation for recently viewed objects serves to temporally parse one object from the next, thus avoiding visual blending when objects are viewed in rapid succession. However, this mechanism for parsing the visual stream comes at a cost, producing negative priming when the previously object was viewed for an extended duration. As applied to the identification of targets presented at the perceptual threshold, this model made the auxiliary assumption that perceptual decisions are based on perceptual fluency. Thus, brief primes boost fluency for primed test items but long duration primes eliminate or reverse this fluency owing to habituation. In contrast, past works such as [Jacoby and White's \(1998\)](#) fluency-attribution theory or Whittlesea and colleague's discrepancy-attribution theory posit that instead subjects strategically discount the influence of the long-duration primes, purposefully choosing the less fluent choice. To properly compare these contrasting accounts, one must examine the joint distribution of accuracy and response time. Applying the nROUSE model and diffusion race model separately to forced-choice and same-different task, we found support for the fluency assumption and we confirmed predictions of the nROUSE model regarding the magnitude of fluency in different priming conditions. However, we found little support that subjects strategically discounted the influence of long-duration primes.

Acknowledgments

We'd like to thank Joseph Houpt, Sachiko Kinoshita, and an anonymous reviewer for extremely helpful comments on earlier versions of this manuscript.

Funding

This work was supported by the National Science Foundation grant BCS-1431147.

Supplementary data

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.cogpsych.2017.11.004>.

References

- Akaike, H. (1973). Information theory and an extension of the maximum likelihood principle. In B. N. Petrov, & F. Caski (Eds.). *Proceedings of the second international symposium on information theory* (pp. 267–281). Budapest: Akademiai Kiado.
- Alter, A. L., & Oppenheimer, D. M. (2009). Uniting the tribes of fluency to form a metacognitive nation. *Personality and Social Psychology Review*, 13, 219–235. <http://dx.doi.org/10.1177/1088868309341564>.
- Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.
- Bodner, G. E., Johnson, J., & Masson, M. E. (2015). Fluency can bias masked priming of binary judgments: Evidence from an all-nonword task. *Canadian Journal of Experimental Psychology/Revue Canadienne de Psychologie Expérimentale*, 69, 200–205. <http://dx.doi.org/10.1037/cep0000047>.
- Broyden, C. G. (1970). The convergence of a class of double-rank minimization algorithms. *Journal of the Institute of Mathematics and its Applications*, 6, 76–90.
- Davelaar, E. J., Tian, X., Weidemann, C. T., & Huber, D. E. (2011). A habituation account of change detection in same/different judgments. *Cognitive, Affective, & Behavioral Neuroscience*, 11, 608–626. <http://dx.doi.org/10.3758/s13415-011-0056-8>.
- Eddelbuettel, D., & François, R. (2011). Rcpp: Seamless R and C++ integration. *Journal of Statistical Software*, 40, 1–18. <<http://www.jstatsoft.org/v40/i08/>> .
- Eddelbuettel, D., & Sanderson, C. (2014). Rcpparmadillo: Accelerating R with high-performance C++ linear algebra. *Computational Statistics and Data Analysis*, 71, 1054–1063. <http://dx.doi.org/10.1016/j.csda.2013.02.005>.
- Fletcher, R. (1970). A new approach to variable metric algorithms. *Computer Journal*, 13, 317–322.
- Gay, D. M. (1990). *Usage summary for selected optimization routines. Technical report 153*. AT & T Bell Laboratories Murray Hill.
- Goldfarb, L. (1970). A family of variable metric updates derived by variational means. *Mathematics of Computation*, 24, 23–26.
- Gomez, P., Perea, M., & Ratcliff, R. (2013). A diffusion model account of masked versus unmasked priming: Are they qualitatively different? *Journal of Experimental Psychology: Human Perception and Performance*, 39, 1731–1740.
- Hautus, M. J. (1995). Corrections for extreme proportions and their biasing effects on estimated values of d'. *Behavior Research Methods Instruments, & Computers*, 27, 46–51.
- Hertwig, R., Herzog, S. M., Schooler, L. J., & Reimer, T. (2008). Fluency heuristic: A model of how the mind exploits a by-product of information retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34, 1191–1206. <http://dx.doi.org/10.1037/a0013025>.
- Higham, P. A., & Vokey, J. R. (2000). Judgment heuristics and recognition memory: Prime identification and target-processing fluency. *Memory & Cognition*, 28, 574–584. <http://dx.doi.org/10.3758/BF03201248>.
- Hochhaus, L., & Johnstone, J. C. (1996). Perceptual repetition blindness effects. *Journal of Experimental Psychology: Human Perception and Performance*, 22, 355–366. <http://dx.doi.org/10.1037/0096-1523.22.2.355>.
- Huber, D. E. (2008). Immediate priming and cognitive aftereffects. *Journal of Experimental Psychology: General*, 137, 324–347. <http://dx.doi.org/10.1037/0096-3445.137.2.324>.
- Huber, D. E. (2015). Using continual flash suppression to investigate cognitive aftereffects. *Consciousness and Cognition*, 35, 30–32. <http://dx.doi.org/10.1016/j.concog.2015.04.017>.
- Huber, D. E., Clark, T. F., Curran, T., & Winkielman, P. (2008). Effects of repetition priming on recognition memory: Testing a perceptual fluency-disfluency model. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34, 1305–1324. <http://dx.doi.org/10.1037/a0013370>.
- Huber, D. E., & O'Reilly, R. C. (2003). Persistence and accommodation in short-term priming and other perceptual paradigms: Temporal segregation through synaptic

- depression. *Cognitive Science*, 27, 403–430.
- Huber, D. E., Shiffrin, R. M., Lyle, K. B., & Quach, R. (2002). Mechanisms of source confusion and discounting in short-term priming 2: Effects of prime similarity and target duration. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28, 1120–1136. <http://dx.doi.org/10.1037/0278-7393.28.6.1120>.
- Huber, D. E., Shiffrin, R. M., Lyle, K. B., & Ruys, K. I. (2001). Perception and preference in short-term word priming. *Psychological Review*, 108, 149–182. <http://dx.doi.org/10.1037/0033-295X.108.1.149>.
- Huber, D. E., Tian, X., Curran, T., O'Reilly, R. C., & Woroch, B. (2008). The dynamics of integration and separation: ERP, MEG, and neural network studies of immediate repetition effects. *Journal of Experimental Psychology: Human Perception and Performance*, 34, 1389–1416. <http://dx.doi.org/10.1037/a0013625>.
- Jacoby, L. L., Kelley, C. M., & Dwyan, J. (1989). Memory attributions. In H. I. Roediger, & F. M. Craik (Eds.). *Varieties of memory and consciousness: Essays in honour of Endel Tulving* (pp. 391–422). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- Jacoby, L. L., & Whitehouse, K. (1989). An illusion of memory: False recognition influenced by unconscious perception. *Journal of Experimental Psychology: General*, 118, 126–135. <http://dx.doi.org/10.1037/0096-3445.118.2.126>.
- Jeffreys, H. (1961). *Theory of probability*. Oxford, UK: Oxford University Press.
- Kanwisher, N. G. (1987). Repetition blindness: Type recognition without token individuation. *Cognition*, 27, 117–143. [http://dx.doi.org/10.1016/0010-0277\(87\)90016-3](http://dx.doi.org/10.1016/0010-0277(87)90016-3).
- Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90, 773–795.
- Kelley, C. T. (1999). *Iterative methods for optimization*. Raleigh, North Carolina: Society for Industrial and Applied Mathematics.
- Krueger, L. E. (1978). A theory of perceptual matching. *Psychological Review*, 85, 278–304. <http://dx.doi.org/10.1037/0033-295X.85.4.278>.
- Kucera, H., & Francis, W. (1967). *Computational analysis of present-day American English*. Providence, Rhode Island: Brown University Press.
- Logan, G. D., Van Zandt, T., Verbruggen, F., & Wagenmakers, E.-J. (2014). On the ability to inhibit thought and action: General and special theories of an act of control. *Psychological Review*, 121, 66–95. <http://dx.doi.org/10.1037/a0035230>.
- Luce, R. D. (1986). *Response times: Their role in inferring elementary mental organization*. Oxford, New York: Oxford University Press.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco, California: W.H. Freeman.
- McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: I. An account of basic findings. *Psychological Review*, 88, 375–407. <http://dx.doi.org/10.1037/0033-295X.88.5.375>.
- Morey, R. D., & Rouder, J. N. (2015). BayesFactor: Computation of Bayes factors for common designs. r package version 0.9.12-2. < <https://CRAN.R-project.org/package=BayesFactor> > .
- Nash, J. C., & Varadhan, R. (2011). Unifying optimization algorithms to aid software system users: Optimx for R. *Journal of Statistical Software*, 43, 1–14. <<http://www.jstatsoft.org/v43/i09/>> .
- Nelder, J. A., & Mead, R. (1965). A simplex method for function minimization. *Computer Journal*, 7, 308–313. <http://dx.doi.org/10.1093/comjnl/7.4.308>.
- Norris, D. (2006). The Bayesian reader: Explaining word recognition as an optimal Bayesian decision process. *Psychological Review*, 113, 327–357. <http://dx.doi.org/10.1037/0033-295X.113.2.327>.
- Norris, D., & Kinoshita, S. (2008). Perception as evidence accumulation and Bayesian inference: Insights from masked priming. *Journal of Experimental Psychology: General*, 137, 434–455. <http://dx.doi.org/10.1037/a0012799>.
- R Core Team (2016). R: A language and environment for statistical computing. Austria: R Foundation for Statistical Computing Vienna. < <http://www.R-project.org/> > .
- Ratcliff, R. (1981). A theory of order relations in perceptual matching. *Psychological Review*, 88, 552–572. <http://dx.doi.org/10.1037/0033-295X.88.6.552>.
- Ratcliff, R., Chierian, A., & Segraves, M. (2003). A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of simple two-choice decisions. *Journal of Neurophysiology*, 90, 1392–1407. <http://dx.doi.org/10.1152/jn.01049.2002>.
- Ratcliff, R., Hasegawa, Y. T., Hasegawa, R. P., Childers, R., Smith, P. L., & Segraves, M. A. (2011). Inhibition in superior colliculus neurons in a brightness discrimination task? *Neural Computation*, 23, 1790–1820. http://dx.doi.org/10.1162/NECO_a.00135.
- Ratcliff, R., Hasegawa, Y. T., Hasegawa, R. P., Smith, P. L., & Segraves, M. A. (2007). Dual diffusion model for single-cell recording data from the superior colliculus in a brightness-discrimination task. *Journal of Neurophysiology*, 97, 1756–1774. <http://dx.doi.org/10.1152/jn.00393.2006>.
- Ratcliff, R., & McKoon, G. (1988). A retrieval theory of priming in memory. *Psychological Review*, 95, 385–408. <http://dx.doi.org/10.1037/0033-295X.95.3.385>.
- Ratcliff, R., & McKoon, G. (2001). A multinomial model for short-term priming in word identification. *Psychological Review*, 108, 835–846. <http://dx.doi.org/10.1037/0033-295X.108.4.835>.
- Ratcliff, R., & Tuerlinckx, F. (2002). Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychonomic Bulletin & Review*, 9, 438–481. <http://dx.doi.org/10.3758/BF03196302>.
- Rieth, C. A., & Huber, D. E. (2010). Priming and habituation for faces: Individual differences and inversion effects. *Journal of Experimental Psychology: Human Perception and Performance*, 36, 596–618. <http://dx.doi.org/10.1037/a0018737>.
- Rieth, C. A., & Huber, D. E. (2013). Implicit learning of spatiotemporal contingencies in spatial cueing. *Journal of Experimental Psychology: Human Perception and Performance*, 39, 1165–1180. <http://dx.doi.org/10.1037/a0030870>.
- Rieth, C. A., & Huber, D. E. (2017). Comparing different kinds of words and word-word relations to test an habituation model of priming. *Cognitive Psychology*, 95, 79–104. <http://dx.doi.org/10.1016/j.cogpsych.2017.04.002>.
- Rouder, J. N., Morey, R. D., Speckman, P. L., & Province, J. M. (2012). Default Bayes factors for ANOVA designs. *Journal of Mathematical Psychology*, 56, 356–374.
- Rusconi, P., & Huber, D. E. (2017). The perceptual wink model of non-switching attentional blink tasks. *Psychonomics Bulletin & Review*, 1–23. <http://dx.doi.org/10.3758/s13423-017-1385-6>.
- Schnabel, R. B., Koontz, J. E., & Weiss, B. E. (1985). A modular system of algorithms for unconstrained minimization. *ACM Transactions on Mathematical Software*, 11, 419–440.
- Schwarz, G. (1978). Estimating the dimension of a model. *Annals of Statistics*, 6, 461–464. <http://dx.doi.org/10.1214/aos/1176344136>.
- Schwarz, N., Bless, H., Strack, F., Klumpp, G., Rittenauer-Schatka, H., & Simons, A. (1991). Ease of retrieval as information: Another look at the availability heuristic. *Journal of Personality and Social Psychology*, 61, 195–202. <http://dx.doi.org/10.1037/0022-3514.61.2.195>.
- Sereno, S. C., Brewer, C. C., & O'Donnell, P. J. (2003). Context effects in word recognition evidence for early interactive processing. *Psychological Science*, 14, 328–333. <http://dx.doi.org/10.1111/1467-9280.14471>.
- Sereno, S. C., Rayner, K., & Posner, M. I. (1998). Establishing a time-line of word recognition: Evidence from eye movements and event-related potentials. *Neuroreport*, 9, 2195–2200.
- Shanno, D. F. (1970). Conditioning of quasi-newton methods for function minimization. *Mathematics of Computation*, 24, 647–656.
- Smith, P. L. (2000). Stochastic dynamic models of response time and accuracy: A foundational primer. *Journal of Mathematical Psychology*, 44, 408–463. <http://dx.doi.org/10.1006/jmps.1999.1260>.
- Tian, X., & Huber, D. E. (2010). Testing an associative account of semantic satiation. *Cognitive Psychology*, 60, 267–290. <http://dx.doi.org/10.1016/j.cogpsych.2010.01.003>.
- Tian, X., & Huber, D. E. (2013). Playing “duck duck goose” with neurons change detection through connectivity reduction. *Psychological Science*, 24, 819–827. <http://dx.doi.org/10.1177/0956797612459765>.
- Tsodyks, M. V., & Markram, H. (1997). The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proceedings of the National Academy of Sciences*, 94, 719–723. <http://dx.doi.org/10.1073/pnas.94.2.719>.
- Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky competing accumulator model. *Psychological Review*, 108, 550–592. <http://dx.doi.org/10.1037/0033-295X.108.3.550>.
- Van Zandt, T., Colonius, H., & Proctor, R. W. (2000). A comparison of two response time models applied to perceptual matching. *Psychonomic Bulletin & Review*, 7, 208–256. <http://dx.doi.org/10.3758/BF03212980>.
- Wagenmakers, E.-J., & Farrell, S. (2004). AIC model selection using Akaike weights. *Psychonomic Bulletin & Review*, 11, 192–196.
- Weidemann, C. T., Huber, D. E., & Shiffrin, R. M. (2008). Prime diagnosticity in short-term repetition priming: Is primed evidence discounted, even when it reliably

- indicates the correct answer? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34, 257–281. <http://dx.doi.org/10.1037/0278-7393.34.2.257>.
- Whittlesea, B. W., & Williams, L. D. (1998). Why do strangers feel familiar, but friends don't? A discrepancy-attribution account of feelings of familiarity. *Acta Psychologica*, 98, 141–165. [http://dx.doi.org/10.1016/S0001-6918\(97\)00040-1](http://dx.doi.org/10.1016/S0001-6918(97)00040-1).
- Whittlesea, B. W., & Williams, L. D. (2000). The source of feelings of familiarity: The discrepancy-attribution hypothesis. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26, 547–565. <http://dx.doi.org/10.1037/0278-7393.26.3.547>.
- Whittlesea, B. W., & Williams, L. D. (2001a). The discrepancy-attribution hypothesis: I. The heuristic basis of feelings and familiarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27, 3–13. <http://dx.doi.org/10.1037/0278-7393.27.1.3>.
- Whittlesea, B. W., & Williams, L. D. (2001b). The discrepancy-attribution hypothesis: II. Expectation, uncertainty, surprise, and feelings of familiarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27, 14–33. <http://dx.doi.org/10.1037/0278-7393.27.1.14>.