

A pragmatic account of the processing of negative sentences

Ann E. Nordmeyer

anordmey@stanford.edu
Department of Psychology
Stanford University

Michael C. Frank

mcfrank@stanford.edu
Department of Psychology
Stanford University

Abstract

Previous work suggests that negative sentences are more difficult to process than positive sentences. A supportive context, however, can mitigate this effect. We investigate the role of context on negation by measuring the processing cost of negation with and without a visual context (Study 1) and then systematically varying the strength of the context (Study 2). We find that a supportive visual context has a graded effect on negation processing. We then create a model to compute the informativeness of an utterance in context, and find that a model that considers both the surprisal of an utterance and the surprisal of seeing a referent is highly correlated with reaction times. Our data suggest that pragmatic factors likely explain the processing costs of negation.

Keywords: Negation; sentence processing; pragmatics; Bayesian modeling.

Introduction

Language is a powerful tool that allows us to describe not only the state of the world as we see it, but also the world as it is not. If I am a regular at a coffee shop and always order chai, but the shop has run out today, the barista might say “We don’t have any chai today” when I enter. Negative sentences are very informative when expectations are violated.

Although negation is critical for communicating many meanings, processing negative sentences can be slow and effortful. In sentence verification tasks, participants who are asked to evaluate the truth of a sentence describing a picture take significantly longer to evaluate negative sentences compared to positive ones (Clark & Chase, 1972; Carpenter & Just, 1975; Just & Carpenter, 1971, 1976). In EEG experiments, sentences in which the final noun is semantically unexpected elicit an N400 response, and this response is found even when a negative makes the sentence logically true (e.g. “A robin [is/is not] a truck”)—suggesting that negation is slow to integrate with the rest of the sentence (Fischler, Bloom, Childers, Roucos, & Perry, 1983; Lüdtke, Friedrich, De Filippis, & Kaup, 2008). Similar results have been found in probe-recognition tasks (Kaup & Zwaan, 2003; Kaup, Lüdtke, & Zwaan, 2006; Hasson & Glucksberg, 2006). Collectively, this work suggests that processing negative sentences is often difficult.

There is a critical difference, however, between evaluating a sentence in the lab and comprehending speech in the real world. According to Grice’s Cooperative Principle (Grice, 1975), speakers should produce utterances that are truthful, relevant, and informative. Negative sentences presented without context violate this principle. If the barista says “we don’t have chai today” to a customer who always orders coffee, this utterance would be neither relevant nor informative. In gen-

eral, negations are produced when there is some expectation that the speaker wishes to reverse.

Congruent with this Gricean account, a number of studies have shown that a supportive context mitigates the processing cost of negation (Wason, 1965; Glenberg, Robertson, Jansen, & Johnson-Glenberg, 1999; Lüdtke & Kaup, 2006; Nieuwland & Kuperberg, 2008; Dale & Duran, 2011). Some contexts are more effective than others at reducing processing demands. For example, contexts that explicitly mention a negated characteristic (Lüdtke & Kaup, 2006) or that present the negation within a dialogue (Dale & Duran, 2011) elicit faster reaction times, perhaps because the negation is more informative. But although these findings are congruent with the idea that pragmatic expectations are the source of negation’s processing cost, they do not directly test that hypothesis. The goal of our current work is to make such a test.

We propose that negative sentences are more informative in contexts that set up a strong expectation that is violated. If the processing cost of negation is pragmatic, then more informative negative sentences should elicit smaller reaction times. How should we quantify informativeness in context? Recent modeling work quantifies pragmatic reasoning in simple experimental contexts (Frank & Goodman, 2012; Goodman & Stuhlmüller, 2013). The assumption underlying this work is that speakers are informative—they will produce utterances that will pick out smaller subsets of the context, leaving as little ambiguity as possible for the listener. We use this definition of informativeness to provide a quantitative interpretation of our hypothesis.

To link informativeness—as computed in a probabilistic model—to reaction time, we assume that reaction time is proportional to *surprisal*. Surprisal is an information-theoretic measure of the amount of information carried by an event (in this case, an utterance in some context) based on its probability. Surprisal has been used effectively to predict reaction times from probabilistic models (Levy, 2008); this work provides inspiration for our current model.

We test the hypothesis that pragmatic surprisal explains the processing cost of negative sentences. Study 1 measures this processing cost, replicating previous findings that context facilitates the processing of negation. Study 2 investigates the effect of the strength of the context by parametrically varying the base rate of a negated feature. We compute the surprisal of sentences in these contexts, and find that a model of pragmatic informativeness predicts the relationship between context and reaction time. These results support the idea that context affects negative sentence processing by modulating listeners’ expectations.




Figure 1: An example trial, consisting of two separate slides (shown sequentially): a context slide and a trial slide for a true negative trial.

Study 1: Context vs. No Context

To test whether non-linguistic contextual expectations alleviate the processing cost of negative sentences, we constructed a simple sentence verification task based on Clark and Chase (1972). Previous studies of the relationship between context and negation have required participants to actively engage with the context, either by describing pictures (Wason, 1965) or reading sentences (Glenberg et al., 1999). Here, participants passively viewed a visual context, eliminating linguistic confounds in previous work.

Method

Participants We recruited 100 participants to participate in an online experiment through the Amazon’s Mechanical Turk (mTurk) website.¹ Participants ranged in age from 18-65; 63 were male and 37 female. We restricted participation to individuals in the United States. We paid participants 30 cents to participate, which took approximately 5 minutes to complete.

Stimuli Twenty-eight trial items were created in which a character was shown holding either two of the same common, recognizable objects (e.g. two apples), or holding nothing. On each trial a sentence of the form “[NAME] [has/has no] [ITEM]” was written. Half of the sentences were positive and half were negative, and they were paired with pictures such that half were true and half were false. The experiment was fully crossed, with participants receiving seven true positive, seven false positive, seven true negative and seven false negative sentences in a randomized order over the course of the study.

Participants were randomly assigned to the “no context” condition or the “context” condition. Participants in the no context condition saw a blank screen with a fixation cross before each trial, while participants in the context condition viewed a context slide. The context slide showed three characters, each holding the same two identical items. The characters all differed from the trial character and from each other in hair and shirt color. A sentence instructed participants to “Look at these [boys/girls]!” (Fig. 1).

¹Previous work has shown that mTurk is an effective tool for collecting RT data (Crump, McDonnell, & Gureckis, 2013).




Figure 2: Reaction times for each trial type across different conditions. Responses to true sentences are shown on the left, and false sentences are shown on the right. Negative sentences are shown in grey, and positive sentences in black. Error bars show 95% confidence intervals.

Procedure Participants were first presented with an instructions screen which described the task and informed them that they could stop at any time. Once they accepted the task, they were given eight positive sentence practice trials with feedback about incorrect responses.

In each trial, participants saw a context (3s) and then a picture and a sentence. They were asked to read the sentence and respond as quickly and accurately as possible with a judgment of whether it was true or false when applied to the picture. We recorded reaction times for each trial, measured as the time from when the picture and sentence were presented to the moment when the response was made.

Data Processing We excluded from analysis 6 participants who did not list English as their native language, 7 participants for having participated in a previous pilot study, and 4 participants for having an overall accuracy of below 80%. Thus, data from a total of 83 participants were analyzed. We also excluded trials with RTs greater than 3 standard deviations from the log-transformed mean.

Results & Discussion

Negative sentences were difficult to process when presented without context; in context, this effect disappeared (Fig. 2). This result is congruent with previous work on sentence verification, which has also found a main effect of negation (e.g. Clark & Chase, 1972) and with work examining the role of context in negation (e.g. Wason, 1965; Nieuwland & Kuperberg, 2008; Dale & Duran, 2011).

To examine the reliability of these findings, we fit a linear mixed-effects model to participants’ reaction times. We examined the interaction between sentence type, truth value, and context on reaction times.² Results of this model show

²All mixed-effects models were fit using the lme4 package in R version 2.15.3. The model specification was as follows: $RT \sim sentence \times truth \times context + (sentence \times truth | subject) +$

a main effect of truth value, with significant faster reaction times for true sentences compared to false sentences ($\beta = -196$, $p < .001$).³ Although there was no main effect of negation across both conditions, there was an interaction between sentence type and truth value ($\beta = 260$, $p < .001$), replicating the finding that participants respond fastest to true positive sentences but slowest to true negative sentences (Clark & Chase, 1972). Critically, there was a significant 3-way interaction between context condition, sentence type, and truth value ($\beta = -227$, $p < .01$), suggesting that this interaction was primarily driven by the slow RTs for true negative sentences in the no context condition.

To understand why context had the strongest effect on true negative sentences, consider what a true negative trial looks like in the no context condition. These are trials in which the participant has no expectation about what the character might be holding, because no context was provided to set up such an expectation. The participant would then see a picture of an empty-handed boy with the sentence “Bob has no apples.” These types of trials likely cause participants to falter because there is no reason for “apples” to be mentioned at all. However, when a participant first views a context such as the one in Fig. 1, they can form an expectation that boys typically have apples. Now, when participants see a boy with no apples, a sentence such as “Bob has no apples” makes sense.

Study 1 contributes to a body of evidence suggesting that negative sentences are more felicitous when they negate an expectation, and that such expectations can be set up by an appropriate context. In Study 2, we examine how systematically manipulating the strength of the context might produce changes in reaction times by altering the expectations created by the context.

Study 2: Varying strength of context

Should all contexts be equally helpful in processing negation? In Study 2, we parametrically manipulated the strength of the context. Participants saw contexts consisting of either three (Study 2a) or four (Study 2b) characters in which some subset of the characters were holding the target item. If the context gives participants a glimpse into the “world” that each trial exists in, this represents a small sample of the base rate of what the characters in this world look like. By manipulating this base rate, we can change peoples’ expectations about the trial character. If the differences in reaction times between the no context and the context condition in Study 1 are due to the relative informativeness of the negative utterance based on the context, we should expect to see a relationship between the strength of the context and reaction time.

(sentence \times truth | item). Significance was calculated using the standard normal approximation to the t distribution (Barr, Levy, Scheepers, & Tily, 2013). Data and analysis code can be found at http://github.com/anordmey/cogsci14_negatron

³Coefficient weights are interpretable in milliseconds.

Method

Participants We again recruited participants from mTurk, 200 in 2a (129 male, 71 female) and 400 in 2b (205 male, 195 female), ages 18 – >65. We again restricted participation to individuals in the US and paid 30 cents for this 5 minute study.

Stimuli Study 2a used the same 28 trial items and sentence types as those used in Study 1. A between-subjects factor determined what type of context participants saw. Context conditions showed $\frac{0}{3}$, $\frac{1}{3}$, $\frac{2}{3}$, or $\frac{3}{3}$ of the characters holding objects. Trial stimuli were identical to those in Study 1.

Study 2b used 48 items. The contexts were the same as in Study 2a, except that each context contained 4 characters and there were therefore 5 context conditions ($\frac{0}{4}$, $\frac{1}{4}$, $\frac{2}{4}$, $\frac{3}{4}$ or $\frac{4}{4}$).

Procedure The procedure for Study 2a was identical to that of Study 1, with participants randomly assigned to condition. In Study 2b, participants were given 4s (instead of 3s) to view the context before the experiment advanced. This latency was changed to give participants more time to look at the slightly larger contexts; the procedure was otherwise identical.

Data Processing We excluded 35 participants who did not list English as their native language (9 in 2a and 16 in 2b), 24 participants for participating in a previous iteration of the experiment (3 in 2a and 21 in 2b), and 35 participants for having an overall accuracy below 80% (11 in 2a and 24 in 2b). Thus, we analyzed data from a total of 177 participants in Study 2a and 339 participants in Study 2b. As in Study 1, we only analyzed correct trials and excluded trials with RTs greater than 3 SDs from the log-transformed mean.

Because we were interested in the effect of context, results from these two studies were combined and analyzed together with context condition re-coded as a continuous variable by calculating the proportion of people in each context condition who had a target item (e.g. the $\frac{1}{3}$ condition in Study 2a was recoded as .33).

Results and Discussion

As the proportion of target items in the context increased, reaction times tended to decrease, particularly for negative and false sentences, supporting our hypothesis (Fig. 3). Unexpectedly, reaction times increased slightly when all characters in the context had target items, resulting in a U-shaped relationship between context and RT.

We fit a linear mixed-effects model to reaction times in response to sentences. We examined the interaction between sentence type, truth value, and context on reaction times.⁴ As in Study 1, we found a significant effect of truth value, with significantly faster reaction times for true sentences compared to false sentences ($\beta = -154$, $p < .001$). Although there was

⁴The model specification was as follows:
 $RT \sim sentence \times truth \times context$
 $+ (sentence \times truth | subject) + (sentence \times truth | item)$.




Figure 3: Reaction times for each trial type across different conditions. Responses to true sentences are shown on the left, and false sentences are shown on the right. Negative sentences are shown in grey, and positive sentences in black. Data for Study 2a (3-person contexts) are shown above, and data for Study 2b (4-person contexts) are shown below. The context condition is noted by a fraction representing the number of characters in the context who held target items. Error bars show 95% confidence intervals.

not a significant main effect of negation, there was a significant interaction between sentence type and truth value, such that the difference between true positive and true negative was greater than the difference between the two types of false sentences ($\beta = 159, p < .001$). There was also a linear effect of context, such that as the proportion of people with the target item in the context increased, reaction times decreased ($\beta = -197, p < .001$). As before, there was a significant 3-way interaction between context, sentence type, and truth value, such that the linear effect of context was most striking in true negative sentences ($\beta = -141, p < .001$).

Responses in the $\frac{3}{3}$ and $\frac{4}{4}$ conditions suggest that the relationship between context and RT is not linear (Fig. 3). We added a quadratic term to our model to test for this nonlinear effect of context ($\beta = 610, p < .001$). The quadratic model fit our data significantly better in a likelihood comparison test ($\chi^2(1) = 80.59, p < .001$).

Quantitative manipulation of the strength of the context resulted in systematic changes in the processing cost of negation, particularly for true negative sentences. This finding is consistent with our initial hypothesis: As the proportion of people in the context with the target item increases, describ-

ing the trial picture as *not* having that target item becomes more informative. That is, the more people in the context who have apples, the more we expect a person with nothing to be described as “a boy with no apples.”

Model

Studies 1 and 2 show that a simple visual context can facilitate the processing of negation, with contexts that set up a strong expectation leading to faster RTs for negative sentences. We hypothesized that this effect was driven by the expectation that speakers are informative (Grice, 1975; Frank & Goodman, 2012): If everyone in a context has a specific feature, and the target character is lacking that feature, it is highly informative to describe the target character in terms of the negation of the expected feature. In contrast, if no one has a feature, it’s pragmatically odd to negate it. In this section, we formalize these intuitions. Due to the Gricean nature of the intuition—which lead us to consider a truthful speaker as well—we focus here on predicting the processing of true sentences.

Model 1: Utterance Surprisal

We modeled the behavior of participants in our experiments by assuming that reaction time is proportional to the surprisal of the utterance w , given the context C and the speaker’s intended referent r_s (following Levy, 2008):

$$RT \sim -\log(P(w|r_s, C)). \quad (1)$$

We then define the probability of the utterance as proportional to its utility (following Frank & Goodman, 2012):

$$P(w|r_s, C) \propto e^{U(w;r_s,C)}, \quad (2)$$

This utility is defined as the informativeness of w minus its cost $D(w)$:

$$U(w;r_s,C) = I(w;r_s,C) - D(w). \quad (3)$$

Informativeness in context is calculated as the number of bits of information conveyed by the word. We assume that w has a uniform probability distribution over its extension in context (e.g. “boy with apples” applies to any boy who has apples, leading to a probability of $1/|w|$ of picking out each individual boy with apples) :

$$I(w;r_s,C) = -(-\log(|w|^{-1})). \quad (4)$$

The cost term $D(w)$ can then be defined in any number of ways; in this model we define it as the number of words in the utterance multiplied by a cost-per-word parameter. Note that in our experiment, the negative sentences always have exactly one word more than the positive sentences.

We created a sparse vocabulary which represented possible words to describe the characters. This included the target utterance (e.g. “apples” and “no apples”), as well as words

that were uniformly true or false of all characters. Combining Equations 2–4, and normalizing Eq. 2 over all possible words in the vocabulary V , we have:

$$P(w|r_s, C) = \frac{e^{\log(|w|^{-1}) - D(w)}}{\sum_{w' \in V} e^{\log(|w'|^{-1}) - D(w')}}. \quad (5)$$

Combining Eq. 1 with Eq. 5, this model predicts that as the number of e.g. boys with apples in the context increases, the informativeness of the negative sentence “Bob has no apples” increases, because it selects an increasingly smaller subset of the context. Highly informative sentences will have high probability, hence lower surprisal and faster RTs.

We fit this model to data from Study 2a, with cost = 3 (Table 1). When the model was fit to the combined data from Studies 2a and 2b, the cost-per-word parameter remained the same (Fig. 4). This model accounted for a substantial amount of variance in participant reaction times from Study 2 ($r = .76, p < .001$). Nevertheless, the model fails to capture the U-shaped relationship seen in Study 2; specifically, it underestimates the surprisal of $\frac{0}{3}$ and $\frac{0}{4}$ contexts for positive sentences, and $\frac{3}{3}$ and $\frac{4}{4}$ contexts for negative sentences. In these trials, participants may have found the target picture surprising regardless of the sentence that they read. For example, in $\frac{0}{3}$ and $\frac{0}{4}$ contexts followed by a true positive trial, participants saw several boys with nothing, and then saw a boy holding something.

Model 2: Utterance and Referent Surprisal

To account for reaction time related to seeing the target picture, we included the surprisal of the referent r_s as well as the surprisal of the utterance w given the referent. We estimated the probability of seeing the referent via the count of the target property in the context, smoothed with a parameter λ :

$$P(r_s|C) = \frac{\#MatchingPeople + \lambda}{\#TotalPeople + 2\lambda} \quad (6)$$

We then added $-\log(p(r|C))$ (Eq. 6) to $-\log(p(w|r_s, C))$ (Eq. 5), resulting in:

$$RT \sim -\log(P(w|r_s, C)) - \beta \log(P(r_s|C)). \quad (7)$$

Note that this formulation is quite similar to a model which accounts for the prior probability of the referent $p(r_s)$; the only difference is our use of a weight β to adjust the different effects of these two probabilities.

Consider the example in Fig. 1, in which you see three boys with apples and then a boy with no apples. The sentence “Bob has no apples” is highly probable—and thus low surprisal—in this context, because it uniquely identifies the target character (Eq. 5). In the current model, however, we must also calculate the surprisal of seeing the target character (i.e. the referent). In this example, the referent surprisal is high, because the probability of seeing a boy with no apples in this context is low (Eq. 6).

Table 1: Model parameters and correlations between model predictions and reaction times, for both Model 1 (Utterance surprisal only) and Model 2 (Utterance and referent surprisal). Parameters are either fit to Study 2a only or to both studies, as indicated.

Model	Data	cost	λ	β	r
1 (Utt only)	Study 2a	3			.84
	Study 2b (2a params)	3			.71
	Both (2a params)	3			.76
	Both	3			.76
2 (Utt + Ref)	Study 2a	.5	.1	.3	.95
	Study 2b (2a params)	.5	.1	.3	.86
	Both (2a params)	.5	.1	.3	.89
	Both	.4	.2	.4	.90

We again fit this model to data from Study 2a and compared model predictions to data from Study 2b as well as the combined data from 2a and 2b (Table 1). Using the parameters fit to Study 2a, the model accounted for a substantial amount of variance in participant reaction times from Study 2 ($r = .89, p < .001$). We also fit the model to the combined data from Study 2, which resulted in similar parameter values (Table 1), and continued to account for a substantial amount of the variance in RTs ($r = .90, p < .001$; Fig. 4).

General Discussion

What makes negation so hard? It takes longer to evaluate negative sentences than positive sentences when presented without context, but these effects are mitigated in context. We suggested a Gricean account: the processing cost of negation is related to the degree to which it violates expectations about communication in context. In our studies, by changing the proportion of people in the context who held a target item, we systematically manipulated participants’ contextual expectations. We found a parametric relationship between the strength of the context and reaction times, and this relationship was well fit by a model of the surprisal of a sentence and its referent given the context.

Previous work on sentence processing has suggested that processing negation is fundamentally difficult, perhaps due to the processing cost of negating a proposition (e.g. Clark & Chase, 1972) or the cost of suppressing an affirmative representation (e.g. Kaup & Zwaan, 2003). Our work here suggests that the difficulty of negation may not be unique to negation at all; instead, general pragmatic mechanisms could be driving this effect. Due to the specific pragmatics of negation, negative sentences presented without context are uninformative and are thus unlikely to be produced, leading to increased surprisal and slower processing times. In conversation, however, negative sentences are often produced when some expectation has been violated, decreasing surprisal and processing time.

Although our specific focus was to understand the processing of negative sentences, this work has implications for sentence processing more generally. Debates about the ef-




Figure 4: Best-fitting model predictions for a model of utterance surprisal (above) and a model of total surprisal, Eq. 7 (below). Positive sentences are represented in purple and negative sentences in blue. Context conditions are identified as fractions, written next to the relevant data point. Arrows indicate data points that are not well captured by our initial model of utterance surprisal.

fects of pragmatics on linguistic processing exist in other domains (e.g. the processing of scalar implicatures, Huang & Snedeker, 2009, 2011; Grodner, Klein, Carbury, & Tanenhaus, 2010). We believe that formal models of pragmatics can provide insight into these debates and, more generally, into the role that context plays in linguistic processing.

Acknowledgments

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship.

References

Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68, 255–278.

Carpenter, P., & Just, M. (1975). Sentence comprehension: A psycholinguistic processing model of verification. *Psychological Review*, 82, 45–73.

Clark, H., & Chase, W. (1972). On the process of comparing sentences against pictures. *Cognitive Psychology*, 3, 472–517.

Crump, M. J., McDonnell, J. V., & Gureckis, T. M. (2013). Evaluating Amazon’s Mechanical Turk as a tool for experimental behavioral research. *PloS One*, 8, e57410.

Dale, R., & Duran, N. (2011). The cognitive dynamics of negated sentence verification. *Cognitive Science*, 35, 983–996.

Fischler, I., Bloom, P., Childers, D., Roucos, S., & Perry, N. (1983). Brain potentials related to stages of sentence verification. *Psychophysiology*, 20, 400–409.

Frank, M., & Goodman, N. (2012). Predicting pragmatic reasoning in language games. *Science*, 336, 998.

Glenberg, A., Robertson, D., Jansen, J., & Johnson-Glenberg, M. (1999). Not propositions. *Journal of Cognitive Systems Research*, 1, 19–33.

Goodman, N. D., & Stuhlmüller, A. (2013). Knowledge and implicature: Modeling language understanding as social cognition. *Topics in Cognitive Science*, 5, 173–184.

Grice, H. (1975). Logic and conversation. 1975, 41–58.

Grodner, D. J., Klein, N. M., Carbury, K. M., & Tanenhaus, M. K. (2010). Some, and possibly all, scalar inferences are not delayed: Evidence for immediate pragmatic enrichment. *Cognition*, 116, 42–55.

Hasson, U., & Glucksberg, S. (2006). Does understanding negation entail affirmation? An examination of negated metaphors. *Journal of Pragmatics*, 38, 1015–1032.

Huang, Y. T., & Snedeker, J. (2009). Online interpretation of scalar quantifiers: Insight into the semantics–pragmatics interface. *Cognitive psychology*, 58, 376–415.

Huang, Y. T., & Snedeker, J. (2011). Logic and conversation revisited: Evidence for a division between semantic and pragmatic content in real-time language comprehension. *Language and Cognitive Processes*, 26, 1161–1172.

Just, M., & Carpenter, P. (1971). Comprehension of negation with quantification. *Journal of Verbal Learning and Verbal Behavior*, 10, 244–253.

Just, M., & Carpenter, P. (1976). Eye fixations and cognitive processes. *Cognitive Psychology*, 8, 441–480.

Kaup, B., Lüdtke, J., & Zwaan, R. (2006). Processing negated sentences with contradictory predicates: Is a door that is not open mentally closed? *Journal of Pragmatics*, 38, 1033–1050.

Kaup, B., & Zwaan, R. (2003). Effects of negation and situational presence on the accessibility of text information. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29, 439–446.

Levy, R. (2008). Expectation-based syntactic comprehension. *Cognition*, 106, 1126–1177.

Lüdtke, J., Friedrich, C., De Filippis, M., & Kaup, B. (2008). Event-related potential correlates of negation in a sentence-picture verification paradigm. *The Journal of Cognitive Neuroscience*, 20, 1355–1370.

Lüdtke, J., & Kaup, B. (2006). Context effects when reading negative and affirmative sentences. In *Proceedings of the 28th Annual Conference of the Cognitive Science Society* (pp. 1735–1740).

Nieuwland, M., & Kuperberg, G. (2008). When the truth is not too hard to handle. *Psychological Science*, 19, 1213.

Wason, P. (1965). The contexts of plausible denial. *Journal of Verbal Learning and Verbal Behavior*, 4, 7–11.