Contrast Effects in Priming Paradigms: Implications for Theory and Research on Implicit Attitudes

Laura D. Scherer and Alan J. Lambert
Washington University in St. Louis

Contrast effects have been studied in dozens of experimental paradigms, including the measurement of attitudes in the social psychological literature. However, nearly all of this work has been conducted using explicit reports. In the present research the authors employed a variety of different types of priming tasks in order to gain insight into the nature of contrast effects and the role that automatic processes might play in their emergence. They report 6 experiments. In Experiments 1 and 2 the replicability and robustness of automatized contrast effects across 2 types of implicit tasks are established. Experiments 3–6 were conducted in order to further understand the nature of these effects and whether they are best understood in terms of spreading activation vs. response-based models of priming. In the course of accounting for their findings, the authors propose and validate a response-mapping framework, which provides insight into some longstanding ambiguities in the priming literature. Implications for theories of contrast and models of evaluative priming are discussed.

Keywords: priming, implicit attitudes, contrast, automatic evaluation, response mapping

Supplemental materials: http://dx.doi.org/10.1037/a0015844.supp

When are mice *large* and elephants *small*? This strange state of affairs could easily occur depending on the context in which you made these judgments. In the context of rating amoebas or other single-celled organisms, mice might be considered rather large. By the same token, elephants might seem small if you had just been estimating the size of planetary bodies. What do these considerations have to do with implicit measures of attitudes? In this article we show that these considerations are very much related indeed. In particular, we show how a fundamental understanding of implicit attitudes and the measures used to assess them can both inform, and be informed by, a consideration of the kinds of contrast effects that have long been the focus of experimental psychologists (cf. Helson, 1947; Parducci, 1965; Volkmann, 1951).

A Brief Review of Contrast Effects

One of the important principles to arise out of decades of research on psychological judgment is that context matters. This principle was first demonstrated in the field of psychophysics across a variety of sensory modalities, including judgments of the heaviness of lifted weights, temperature of water, or the brightness of lights (for reviews, see Eiser, 1990; Petty & Cacioppo, 1981; Wedell, Hicklin, & Smarandescu, 2007). The typical finding obtained in this literature is contrast. For example, participants will rate a given weight as heavier if they had recently lifted a series of relatively lightweight objects than if they had not. Such effects have been demonstrated in a variety of other judgmental domains, including matters of longstanding interest to social psychologists, including inferences about attitudes as held by others (Sherif & Hovland, 1961), impressions of single group members (Simpson & Ostrom, 1976), and even judgments of one's own happiness and mental well-being (Brickman, Coates, & Janoff-Bulman, 1978).

Ambiguities Arising From Reliance on Explicit Measures

The majority of contrast effects have been obtained using explicit measures, in which participants are posed with a probe about the target (e.g., *how heavy is this weight?*; *how attractive is this person?*) and are then asked to overtly report their internal experiences regarding the stimulus in question. Researchers in this area have long been aware of the limitations of using explicit measures. As Eiser (1986) has noted, "[t]he problem is that sensation and subjective experience can only be measured indirectly, by getting the subject to make a response (e.g., by marking a category on a rating scale)" (p. 126). To clarify, the issue is not whether perception is considered to be "subjective" or not. Indeed, the subjective nature of perception is taken as a given in this area. Rather, the point is that there can often be a serious disconnect between people's subjective *perceptions* of a stimulus (how the stimulus feels) and overt *reports* about that stimulus (how the stimulus is rated). The possible disconnect between one's perception of the

Laura D. Scherer and Alan J. Lambert, Department of Psychology, Washington University in St. Louis.

Correspondence concerning this article should be addressed to Laura D. Scherer, Department of Psychology, Washington University, 1 Brookings Drive, St. Louis, MO 63130. E-mail: scherer@wustl.edu

 0022-3514/09/$12.00 DOI: 10.1037/a0015844

stimulus and how that stimulus is overtly rated has resulted in a longstanding debate in the contrast literature, as noted below.

Perceptual Versus Scale-Based Conceptualizations of Contrast Effects

One view of contrast is that such effects involve bona-fide changes in perceptual experience of the target stimulus (Helson, 1947; Sherif & Hovland, 1961). According to this *change in perception* view, for example, a given "midrange" face could actually elicit stronger feelings of spontaneous attraction than it ordinarily would, if one had earlier been exposed to a series of extremely unattractive faces. An alternative, *scale-based* explanation of contrast effects distinguishes between the *language* one uses to describe the stimulus, as opposed to one's phenomenological experience of the stimulus itself (Ostrom, 1970; Steele & Ostrom, 1974; Upshaw, 1969).

To illustrate, suppose that you were asked to rate the size of a moderately sized stimulus (say, a border collie) with respect to a scale ranging from 0 (*extremely small*) to 10 (*extremely large*). Under normal circumstances, the anchor labels "extremely small" and "extremely large" are likely to call to mind reasonably familiar exemplars of small and large stimuli (e.g., pebbles or skyscrapers, respectively). However, the meaning of "extremely small" is likely to be different if you had recently judged the size of some unusually tiny things (e.g., quarks, protons). In this particular context, therefore, the lower end of the rating scale would now be reserved for exceptionally small objects (e.g., subatomic particles). By extension, this means that ratings of moderately sized things, like collies, might be shifted upward on the rating scale. Although such effects are sometimes referred to as judgmental "distortions," a better descriptor might be judgmental *displacements.* This emphasizes the idea that the introduction of relatively extreme stimuli can lead people to reserve the lower or upper ends of response scale for these targets, leading more moderate stimuli to be displaced ("reassigned") to other response categories.

Summary. Nearly everything psychologists know about contrast effects has been derived from the use of explicit measures. This state of affairs can be seen as responsible, at least in part, for a long-running ambiguity in this literature that has never been completely resolved (Biernat, 2005): Do contrast effects involve actual changes in the cognitive representation of the stimulus? According to the perceptual change view, such change has in fact occurred, whereas the scale-based perspective holds that such effects represent "only a change in response language: the weight feels the same, but is described as being heavier; the girl looks the same, but is described as prettier" (Petty & Cacioppo, 1981, p. 117).

There is an ironic aspect to this reliance on explicit measures in the contrast literature. Despite (or, perhaps, because of) the widespread use of explicit measures in this area, one of the chief contributions of theory and research on contrast is that it has led to a better understanding of the surprisingly complex nature of explicit measures (e.g., Biernat & Manis, 1994; Schwarz, Groves, & Schuman, 1998). Perhaps most relevant to present concerns, the current consensus among researchers is that explicit measures may never be able to completely resolve the debate between perceptual versus response-based models of contrast (Biernat, 2005; Eiser, 1990). In this sense, while theory and research on contrast have been informed by the use of explicit measures, our understanding of these measures themselves has also been informed through this research. This point becomes relevant in the discussion to follow.

Contrast Effects in Implicit Measures

Over the last 10 years, social psychologists have become increasingly interested in implicit, as opposed to explicit, measures of attitudes (Bargh & Chartrand, 1999; Fazio, Jackson, Dunton, & Williams, 1995; Greenwald, McGhee, & Schwartz, 1998). Although there is some lively debate as to the exact meaning of the term "implicit" (De Houwer, 2006), for purposes of this article the main differences between explicit and implicit measures is that the latter (a) do not directly involve direct queries about the attitude object; (b) offer much less opportunity for people to exercise control over their responses; and (c) are typically more sensitive in detecting the presence of, as well as changes in, evaluative associations with the attitude object.

Given the longstanding interest in contrast among attitude theorists, and in light of the aforementioned features of implicit attitude tasks (e.g., sensitivity, relative immunity to conscious control, ability to tap spontaneous attitudinal reactions), one might imagine that researchers would have already used such measures in order to understand the nature of contrast effects. In fact, surprisingly little research in the implicit area has been devoted to addressing the aforementioned ambiguities that have arisen in the classic research on contrast.

One clarification regarding our terminology is worth noting here. Throughout this article, we use the term "contrast" in a manner similar to the way this term has been used in the literature we have reviewed thus far: scenarios in which a relatively "midrange" (i.e., neutral) stimulus appears to take on value that is inversely proportional to the implications of the more extreme, contextual stimuli with which it appears (Sherman, Ahlm, Berman, & Lynn, 1978). The term "contrast" has been used in two other lines of work in the priming area, including research on reverse priming effects (e.g., Glaser & Banaji, 1999; Klauer, Teige-Mocigembad, & Spruyt, 2009) and also double priming paradigms (e.g., Deutsch & Gawronski, 2009; Gawronski, Deutsch, & Seidel, 2005). Upon closer inspection, however, the processes involved in those other paradigms are considerably different from those under consideration in this article. We compare and contrast the implications of our own research with these and other lines of work in the General Discussion section.

Another aspect of the research on implicit attitudes is worth mentioning here. A line of research conducted over the last decade or so has observed that implicit attitudes toward social groups can change as a result of participants' previous experiences with group exemplars (Blair, Ma, & Lenton, 2001; Dasgupta & Greenwald, 2001). For example, making an admirable African American (e.g., Bill Cosby) more accessible prior to an implicit measure of attitudes toward African Americans as a whole can lead to more favorable responses than if a less likeable exemplar (e.g., Mike Tyson) has been made more accessible. However, careful scrutiny of this literature calls into question the idea that automatized associations have changed at all (cf. Devine, 2001, for a related point). In this literature, implicit attitudes toward the relevant exemplars (e.g., Bill Cosby, Mike Tyson) remains unchanged, and the experiment can simply be seen as altering which of these fixed representations comes to mind when participants are thinking about African Americans as a whole.

The aforementioned research is quite different from the present research. In particular, our work is concerned with the possibility that the actual representation of the stimulus has changed as a result of the immediate context in which it is presented. In order to explore this idea, we attempt to systematically change automatized responses to the same, specific, identifiable stimulus by altering the nature of the context in which the stimulus is presented. In so doing, we believe we address more precisely the nature of automatized associations and whether these associations can depend on the nature of the situational context.

Summary of Research Goals

The overarching objective of our research was to draw upon research and theory on implicit attitudes in order to better understand fundamental issues relevant to the effects of context on social attitudes, most specifically contrast effects (here, we use "contrast effects" in the sense originally formulated by the older, explicit literature). Our first two experiments were designed to show direct evidence for *automatized contrast effects*. By "automatized contrast effects," we refer to contrast effects that occur independent of cognitive control (Payne, Jacoby, & Lambert, 2005) and, hence, occur spontaneously, are relatively effortless, and occur independent of conscious intention.

Having obtained evidence for automatized contrast, this led us to consider some larger issues regarding priming tasks in general. As it turns out, our "contrast paradigm" fortuitously offered surprisingly good leverage in addressing some contentious issues that have vexed priming paradigms, such as whether spreading activation versus response-based processes represent the most viable theoretical models for understanding priming effects. As we show, these considerations have great theoretical and practical importance for a number of domains, including the kinds of conclusions that social psychologists typically draw from implicit measures of outgroup prejudice. We explored these considerations in Experiments 3–6.

In a sense, then, the trajectory of our research tended to parallel research and theory on contrast using explicit measures. As discussed earlier, demonstration of contrast using explicit measures ultimately led to greater understanding of the complexities and difficulties in interpreting such measures. An analogous state of affairs arose for us as well. In the course of using implicit measures to study contrast effects, these efforts ended up yielding several new insights into the nature of implicit measures themselves.

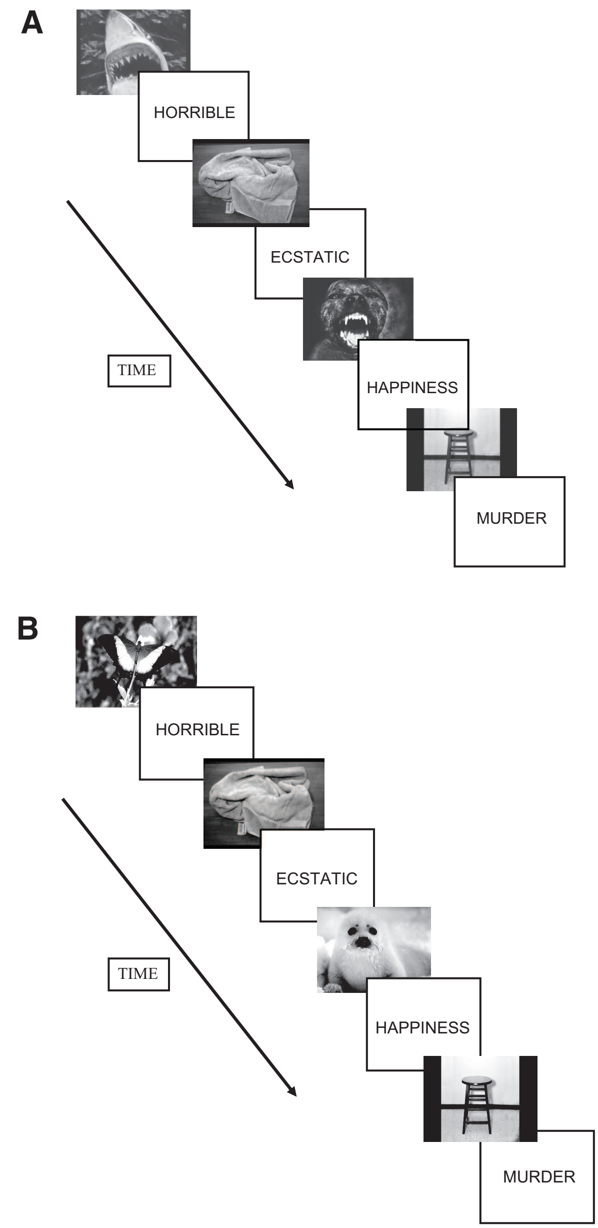

Figure 1. A: Depiction of a "negative" priming context, in which neutral primes are embedded in a task amongst negative primes. B: Depiction of a "positive" priming context, in which neutral primes are embedded in a task amongst positive primes. A color version of this figure is available online at http://dx.doi.org/10.1037/a0015844.supp.

Experiment 1

A schematic representation of the design used in this experiment is provided in Figures 1A and 1B. In the task represented here, participants are asked to categorize target words as quickly as possible as either evaluatively good or bad, while ignoring the primes that immediately precede them. Note that some of the primes are either unambiguously negative (Figure 1A) or positive (Figure 1B). For these unambiguous stimuli, one would naturally expect them to elicit strong priming effects due to their strong preexisting associations, and this is what we also find in our research. That is, unambiguously negative primes should facilitate reactions toward negative words but inhibit responses toward positive words, whereas unambiguously positive primes should elicit the opposite pattern of responses.

What of the neutral primes, however? Note that in each of our two examples, half of the primes are neutral and are embedded amongst the other unambiguous primes (that are either positive [Figure 1A] or negative [Figure 1B]). The question at hand is what kind of spontaneous, unintentional effect these neutral primes might have on participants' responses to the target stimuli that immediately follow them.

At first blush, one might expect that the neutral primes would simply have no effect at all. In this paradigm, we purposely chose neutral stimuli that would have the fewest possible preexisting evaluative associations. Indeed, this being the case, one might suppose that these stimuli would be excellent choices as “control” or “baseline” primes. In fact, we are not aware of any extant theory that would specifically predict that these neutral primes would elicit automatic facilitation or inhibition on the target stimuli. To this extent, one might argue that the most likely outcome would be a null effect.

As should be obvious by now, we were interested in a second and far more interesting possibility: that the neutral stimuli would be subject to contrast effects; that is, they could acquire the capacity to facilitate or inhibit responses to the target stimuli in a way analogous to effects typically observed with unambiguous stimuli. When embedded amongst negative primes, neutral primes should elicit priming effects analogous to those seen with positive primes. Conversely, when embedded amongst positive primes, neutral primes should show effects that are normally associated with negative primes.

Method

Participants and Design

In this experiment, 108 undergraduates participated in return for partial course credit. In the primary task, our design included one between-subjects factor, the nature of the primes presented in the task. Half of the participants received a fully randomized set of positive and neutral primes (positive context), whereas the other half of participants received a randomized set of negative and neutral primes (negative context). A different randomized order of primes and words was generated for each participant.

In addition to the task described above, all participants completed an entirely separate task block, which consisted of positive versus negative primes only, with no neutral primes. This block allowed us to address a number of important issues, the nature of which we discuss presently. Half of the participants completed this block after the primary block of trials (i.e., those involving neutral primes), whereas the other half of participants completed this block first. A schematic display of the design used in this study is presented in Table 1.

Procedure

Participants were instructed via computer to complete a task that was meant to measure “how various tasks influence our evaluations of things in our environment.” They were informed that a picture would flash very briefly and would be replaced by a word that was either objectively good or bad. Participants were told to ignore the picture and identify the word as either good or bad, using the keys labeled “good” and “bad” on the keyboard. The location of the good and bad keys on the left or right side of the keyboard was counterbalanced between subjects. Participants were told to try to respond within the time limit, even if it meant that they would make some errors.

One group of participants received a priming block containing neutral and positive primes (*positive context condition*); the other group received a priming block containing neutral and negative primes (*negative context condition*). As noted above, all participants also completed a block of trials containing positive and negative primes only. Inclusion of this block served two purposes. First, it allowed us to test whether our task would replicate the usual findings obtained in the evaluative priming literature, which most often uses tasks that present participants with both positive and negative primes. Second, this aspect of our design also allowed us to show that the extreme stimuli themselves are *not* subject to context effects. For example, we expected that the unambiguous negative primes should elicit much the same reactions, regardless of whether they appear in the context of neutral or positive primes. Similarly, reactions elicited by the unambiguous positives primes should not vary across condition. These assumptions are somewhat self-evident, given that we deliberately selected these stimuli to have exceptionally clear and unambiguous implications in the first place. Nonetheless, the integrity of our statistical analyses rests on the assumption that the meaning of the stimuli designed to instantiate a particular context did not, themselves, change across experimental condition.

Table 1
Schematic Overview of Design for Experiment 1

Condition	Block 1	Block 2
	Positive priming context	
Counterbalancing condition 1	**Neutral primes and positive primes**	Positive and negative primes
Counterbalancing condition 2	Positive and negative primes	**Neutral primes and positive primes**
	Negative priming context	
Counterbalancing condition 1	**Neutral primes and negative primes**	Positive and negative primes
Counterbalancing condition 2	Positive and negative primes	**Neutral primes and negative primes**

Note. Critical blocks are printed in bold type; priming context refers to the valence of primes into which the neutral primes were embedded in the relevant critical block.

As noted earlier, our task employed an extremely short response window. On each trial, the rectangular prime appeared in the center of the screen for 200 ms and was immediately followed by the target word, which remained on the screen for 100 ms. The target word was then replaced by a mask that was approximately the same size and shape as the prime picture and which remained on the screen for 500 ms. This allowed participants a total of 600 ms (prime duration plus mask duration) in which to respond to the target. If participants did not respond within this window, the words “Faster Please!” appeared in red. If participants responded within the window, the next trial began immediately.

Participants were given 10 practice trials at the beginning of each of the key experimental blocks. Within each block, the presentation of the primes was randomized, and each type of prime was paired once each with five good and five bad target words. This design resulted in a total of 160 trials per block, and 320 trials total (160 trials $\times$ 2 blocks). After completing the priming task, participants were fully debriefed and dismissed.

Stimuli

The priming tasks consisted of 24 priming pictures (all approximately 300 × 230 pixels), all of which were selected from the International Affective Picture System (IAPS; Lang, Bradley, & Cuthbert, 2008) on the basis of the normative data collected by Lang and his colleagues. The primes consisted of eight pictures that had been rated as very positive (e.g., puppies, kittens, rabbits; average valence rating = 7.67, average SD = 1.43), eight that had been rated as very negative (e.g., a snarling dog, cockroaches, spiders; average valence rating = 3.73, average SD = 1.88), and eight pictures of objects that had received ratings around the midpoint of the scale (e.g., a towel, mug, lamp; average valence rating = 4.87, average SD = 0.99). The IAPS scale ranged from 1 (*negative*) to 9 (*positive*).

The target words were selected from a set of eight positive words (*happiness*, *laughter*, *sunny*, *ecstatic*, *honest*, *enjoyable*, *jewel*, *luck*) and eight negative words (*execution*, *offended*, *fever*, *horrific*, *murder*, *depressed*, *agony*, *kill*). These words were selected from a preexisting database (Larsen, Mercer, & Balota, 2006), which allowed us to select words on the basis of their sharply divergent affective implications while controlling for their length, frequency in spoken language, and orthographic neighborhood. The words were presented in Times New Roman bold 40-point font.

Results and Discussion

Data Analytic Technique

Separate indices of errors for the good and bad words were averaged across the valid trials for each of the three types of primes (positive, negative, neutral).[1] Valid trials included those in which responses were made within the 600 ms response window (which included 84% of responses). Next, a proportion of errors was calculated by dividing the average number of errors by the total number of valid responses on trials that contained the prime type. This resulted in two error indices for each prime: proportion of errors on good words and proportion of errors on bad words. In this paradigm, fewer errors on good words relative to bad words following a particular prime indicated positive automatic bias (and vice versa: fewer errors on bad words indicates negative bias). Preliminary analyses included counterbalancing order of task blocks as a factor, but counterbalance order did not show any significant effects (all Fs < 1). As a result, counterbalancing order is not considered further.

Preliminary Analyses on the Positive Versus Negative Primes Only

Recall that participants completed a block of trials consisting of positive and negative primes only (see Table 1). Analyses of this block in isolation revealed a main effect of word type, $F(1, 107) = 6.01$, $p < .05$, which was qualified by the expected two-way Prime × Word Type interaction, $F(1, 107) = 65.63$, $p < .001$. This two-way interaction reflected the fact that on negative prime trials, the error rate on bad words was reliably lower than on good words (Ms = 5.32 vs. 8.56), whereas this difference was reversed on positive prime trials (Ms = 7.31 vs. 5.95). The former difference was reliable, $F(1, 107) = 38.95$, $p < .001$. The latter difference was smaller, but still reliable, $F(1, 107) = 10.26$, $p < .01$. A distinct issue concerns whether the effects elicited by these positive and negative primes were any different when they appeared in the context of neutral primes. They were not: The positive primes elicited approximately the same size of priming effects, regardless of the task block in which they were presented. The same was true for the negative primes (effects of task block for positive and negative primes, both ps $> .10$).

Main Analyses

One of the primary questions at hand is whether spontaneous evaluations of the neutral primes were influenced by task context. In these analyses, therefore, we focus only on the block that involved neutral primes. A 2 (prime: valenced vs. neutral) × 2 (word type: good vs. bad) × 2 (context: whether unambiguous prime is positive or negative) revealed a significant main effect of valence, $F(1, 106) = 4.64$, $p < .05$, indicating that participants made more errors on good words than bad words overall. The only other significant effect was a three-way interaction between prime, word type, and task context, $F(1, 106) = 45.74$, $p < .001$. The data relevant to this analysis are shown in Figure 2, which displays the pattern of errors on good and bad words as a function of whether the priming task consisted of neutral and negative primes (left side) or neutral and positive primes (right side).

Further planned comparisons were performed to reveal the nature of the three-way interaction. Among participants assigned to the positive prime context, the predicted two-way Prime Type (positive vs. neutral) × Word Type (good vs. bad) interaction emerged, $F(1, 54) = 26.38$, $p < .001$. The positive prime facilitated positive word judgments, whereas the neutral prime facilitated negative word judgments. An analogous interaction emerged among participants assigned to the negative priming context, $F(1, 52) = 20.19$, $p < .001$. In this condition, the negative prime facilitated negative word judgments, whereas the neutral prime facilitated positive word judgments. As seen in Figure 2, the effects elicited by the neutral primes were strongly contingent on the context in which they appeared. When the neutral primes appeared amongst negative primes, the neutral primes showed a pattern of facilitation and inhibition normally associated with positive primes (i.e., relatively low error rate on positive words and high error rates on negative words). Conversely, when neutral primes appeared amongst positive primes, the neutral primes elicited effects typically associated with negative primes. This asymmetry was confirmed by statistical analyses on the neutral primes only, which revealed a significant two-way interaction involving word type and priming context, $F(1, 107) = 12.03$, $p = .001$. Simple effects tests revealed that the difference in error rate on the good versus bad words was reliable when neutral primes appeared in the context of positive primes (and, hence, exhibited properties normally seen with negative primes; $p < .05$). For neutral primes

[1] Analyses on reaction times lead to roughly the same conclusions as error rates. In cases where participants were more accurate, they were also faster (e.g., they were faster to respond to positive words that were preceded by positive primes). However, we forced participants to respond very quickly, which restricted the variability in reaction time scores. As a result, the patterns of effects that were observed for reaction times often did not reach significance.

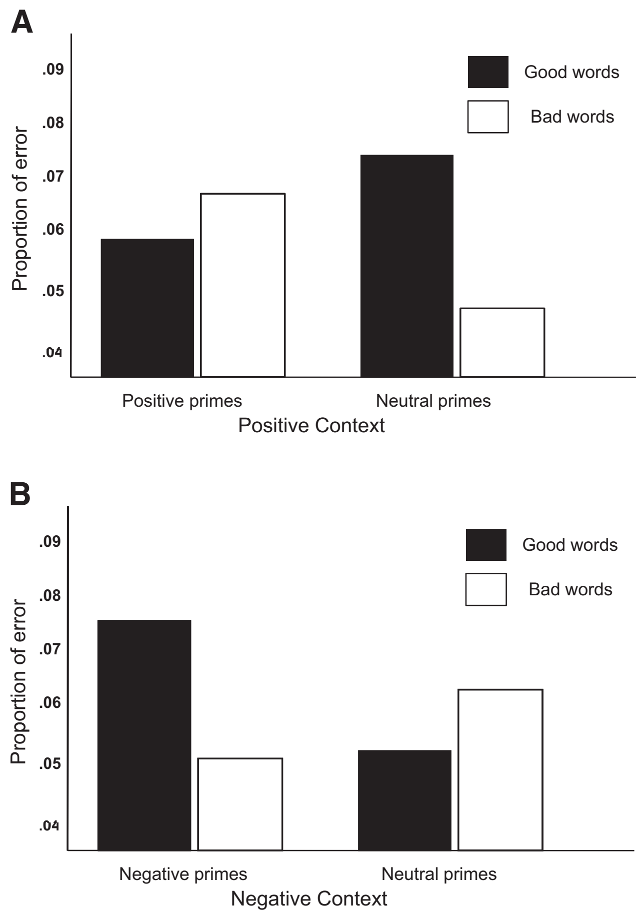

Figure 2. Experiment 1 results. This figure displays the proportion of errors following each prime type, by word type and task context. Greater errors on positive relative to negative words following a particular prime type indicate that the prime elicited negative bias, and the opposite pattern of errors indicates positive bias.

in the negative condition, this pattern was reversed, although not reliably ($p > .10$).

Additional Analyses

One of the interesting implications of Figure 2 is that the priming effects elicited by the neutral primes were roughly equivalent in strength to the valenced primes. These observations were confirmed more formally by theory-driven comparisons involving the neutral primes, which showed no differences whatsoever (all Fs < 1.0) between (a) neutral primes and negative primes, when the former appeared in the context of positive primes, and (b) neutral primes and positive primes, when the former appeared in the context of negative primes.

Experiment 2

A second experiment was conducted in order to show that the effects demonstrated in the previous study could also be generated in a different type of evaluative priming task. Along with showing the replicability and generalizability of our findings, this demonstration is important because it shows that our earlier findings are not an artifact of the particular parameters of the kind of task used in Experiment 1. To this end, Experiment 2 used the affect misattribution paradigm (AMP; Payne, Cheng, Govorum, & Stewart, 2005).

The AMP shares a few basic features with the task used in Experiment 1. Like the task used earlier, the AMP is a sequential priming task, in which participants are told to ignore the prime and respond only to the subsequent target. The AMP also requires participants to choose between two possible evaluative response options. Aside from these two similarities, however, the two tasks are quite different. In the evaluative priming task used earlier, the prime is followed by target *words* (e.g., *happy*, *horrible*), and participants are asked to identify the words as either good or bad. In contrast, in the AMP the prime is followed by Chinese characters, none of which have any a priori meaning to our participants. For each of these characters, participants are asked to rate the target in terms of its aesthetic pleasantness or unpleasantness. Effects in this task are thought to be caused by participants' tendency to misattribute their evaluations of the prime as like or dislike of the subsequent Chinese character (Payne, Cheng, et al., 2005). If contrast were to occur, one should find (a) a tendency for participants to label Chinese characters as pleasant when followed by neutral primes that have been embedded along with negative primes and (b) a tendency for participants to label Chinese characters as unpleasant when followed by neutral primes that have been embedded amongst positive primes.

Method

Participants and Design

A total of 70 participants participated in return for partial fulfillment of course credit. This experiment included one between-subjects factor, which was whether the neutral primes appeared in the task along with positive or negative primes.

AMP Procedure and Stimuli

The priming stimuli were identical to those used in Experiment 1. The targets included 80 Chinese characters that were originally used in Payne, Cheng, et al. (2005). The timing of the trials closely followed Payne, Cheng, et al. The prime was presented in the center of the computer screen for 75 ms and was immediately replaced by one of the Chinese characters. After 100 ms the character was replaced by the same mask used in Experiment 1 and which remained on the screen until participants made their response. Once a response was made, the next trial began immediately. Each individual prime was presented six times per block, in random order, for a total of 48 trials per block (8 primes $\times$ 6 presentations). The Chinese characters that appeared immediately after each prime were selected at random from the pool of 80 characters.

Participants were instructed to ignore the prime picture and to indicate whether the Chinese character was more or less pleasant than average by pressing one of two keyboard keys, labeled "pleasant" and "unpleasant," respectively. The location of the pleasant and unpleasant keys on the left or right side of the keyboard was counterbalanced. Participants were additionally instructed to work through the task quickly and to not think too hard about their responses. The entire experiment took less than 5 min to complete.

Results

For each prime, a single index was calculated by dividing the number of "pleasant" responses by the total number of responses. For this index, higher numbers indicate that the prime elicited more pleasant responses than unpleasant responses; a score of .50 indicates equal numbers of pleasant and unpleasant responses.

Figure 3 displays the pattern of data as a function of priming context and prime type. A 2 (prime: unambiguous [either positive or negative] vs. neutral) $\times$ 2 (context: whether the unambiguous primes were positive or negative) mixed-model analysis of variance (ANOVA) was performed on these data, revealing a significant main effect of context, $F(1, 68) = 6.49$, $p < .05$, indicating that participants made more "pleasant" responses in the positive context than the negative context. This main effect was qualified by a two-way Prime $\times$ Context interaction, $F(1, 68) = 31.14$, $p < .001$. Further planned comparisons indicated that, as expected, participants made more "pleasant" responses following the positive primes than the negative primes, $F(1, 68) = 31.51$, $p < .001$. Of greater importance, participants made significantly more "pleasant" responses following neutral primes in the negative condition compared to the positive condition, $F(1, 68) = 5.21$, $p < .05$ (see Figure 3). This represents a contrast effect, in which the neutral primes elicited more favorable responses when the preceding primes were negative than when they were positive. Paralleling similar analyses in Experiment 1, one can also test whether the neutral primes elicited effects comparable in magnitude to the unambiguous primes. When the neutral primes were embedded amongst the negative primes, these neutral primes elicited positive priming effects that were not significantly different from the effects elicited by the unambiguously positive primes ($p > .20$). When embedded amongst the positive primes, the neutral primes elicited negative priming effects that were nonetheless not quite as strong as those elicited by the unambiguously negative primes ($p < .05$).

Discussion of Experiments 1 and 2

Experiments 1 and 2 showed that neutral primes can elicit automatized evaluations that are about as strong as those elicited by unambiguously positive and negative stimuli. Specifically, we obtained strong contrast effects. The neutral primes activated favorable responses when presented along with negative primes and activated unfavorable responses when embedded amongst the positive primes. At this point an important question emerges, however. What psychological processes are driving these contrast effects? The answer to this question depends, in large part, on one's theoretical perspective regarding evaluative priming effects, an issue we explore in more depth ahead.

A Spreading Activation View of Priming

One early (and, for many years, the dominant) view of evaluative priming effects drew heavily from research and theory on semantic priming effects (Neely, 1991). According to this view, mere presentation of the prime was

> assumed to activate its representing node in a lexical or semantic network . . . and the activation then spreads to nodes of evaluatively consistent targets, but not of inconsistent targets, thereby facilitating processing of the target whenever prime and target are evaluatively consistent. (Klauer & Musch, 2003, p. 22)

The relevance of this model to evaluative priming tasks that use words as targets (cf. Experiment 1) is self-evident. However, it also provides a fairly straightforward account of the AMP (cf. Experiment 2) insofar as spreading activation processes triggered

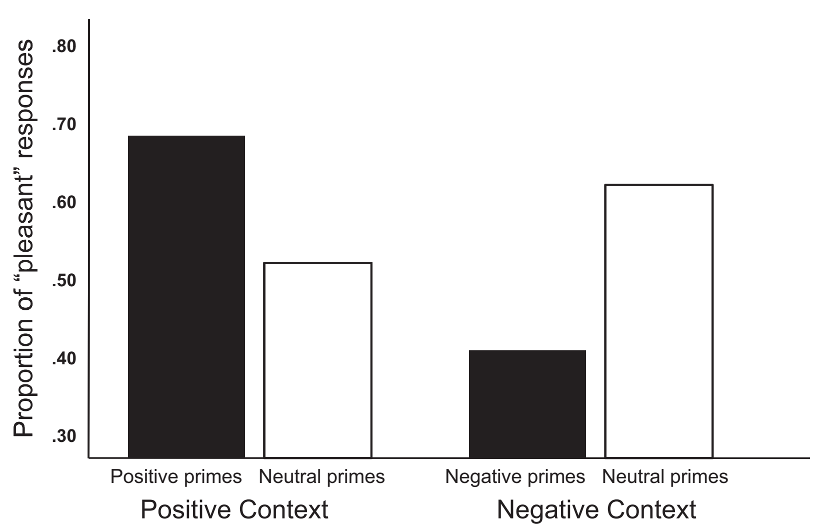

Figure 3. Experiment 2 results. This figure displays the proportion of "pleasant" responses following each of the primes, by task context. Greater number of "pleasant" responses following a prime indicates that the prime elicited a relatively positive response bias.

by the prime lead to a misattribution of positive or negative affect to the Chinese characters.

Potential implications of spreading activation for explaining contrast effects. Spreading activation models are not models of contrast effects per se. However, one could use such an account to explain the kinds of effects we obtained in Experiments 1 and 2. For example, take the neutral primes that were presented in the context of negative primes. In this case, the neutral primes facilitated responses toward positive targets and inhibited responses toward negative targets. This being the case, one might draw the conclusion that there has been genuine change in the underlying evaluation of these primes; that is, these primes are now eliciting automatic activation of positive or negative associations, depending on the context in which they are presented. Note that this could account for the pattern of findings for the neutral primes in both the evaluative priming task (Experiment 1) as well as the AMP (Experiment 2).

There is some theoretical precedence for anticipating this sort of change of meaning process. As we noted earlier, some of the earliest models of contrast (e.g., Helson, 1947) assume that contrast effects involve an underlying change in how stimuli are actually experienced. This suggests the possibility that the neutral primes evoke phenomenologically different reactions depending on the context in which they occur. Proponents of this view have not always been clear in terms of the exact mechanisms by which such perceptual changes occur (Eiser, 1986). However, one potential mechanism might involve changes in attentional focus, that is, the kinds of features that people attend to and/or remember about the target stimulus. For example, scissors are useful tools, but they can also be used as lethal weapons. Hence, the presence of contextual stimuli might highlight certain associations that could be somewhat positive or negative, even for apparently mundane objects.

Shortcomings. However, there are at least two problems with this view. First, it seems just as plausible, if not more so, to suggest that neutral primes would acquire evaluative meaning *consistent* with the meaning of the primes with which they appear (e.g., wooden stools might take on a more positive meaning when embedded in the context of extremely positive primes).[2] Hence, this selective attention account would need to explain why participants would attend to features diametrically opposite to, rather than similar to, those contained in the contextual primes.

A second problem concerns the size of the effects obtained. Not only did we obtain contrast with the neutral primes, but in some cases these effects were very much comparable to the effects obtained with the unambiguous primes! It is conceivable that the neutral primes might have undergone *some* change of meaning, but it is difficult to imagine that the change of meaning would have been so profound that (for example) a stool could became just as negative as a snarling dog or a fearsome-looking shark. These considerations led us to consider the possibility of an alternative account, which we describe below.

A Response-Based Perspective of Priming

Response-based models (also known as "response compatibility" or "response interference" models) were initially conceptualized to explain sequential evaluative priming tasks in which both the prime and target have preexisting evaluative meaning (e.g., *kitten* → *wonderful*). The basic idea behind response-based models is that the "activation of the attitude associated with the [prime] suggests a response that either facilitates or inhibits the response to the following target word" (Ferguson & Bargh, 2003, p. 173). For example, presentation of a prime associated with a positive response (e.g., kitten) should lead to relatively fast responses toward targets that are also associated with positive responses but should interfere with responses to negative targets (e.g., *kitten* followed by the word *disgusting*). This is different from spreading activation, which states that facilitation should be caused by relatedness between the prime and target, irrespective of the nature of the required response. For example, spreading activation models predict that a picture of a snarling dog should facilitate responses to the word *horrible* regardless of whether the response requires an evaluative (e.g., good vs. bad) or semantic (e.g., adjective vs. noun) classification of the target.

A response-based view could also account for findings typically observed in the case of the AMP in which the target stimuli (Chinese characters) have no intrinsic meaning in their own right. In this case, presentation of unambiguously positive or negative primes could be seen as making positive or negative responses more accessible, respectively. Hence, in the very brief span of time in which participants must "decide" whether a Chinese character is pleasant or unpleasant, prior activation of a pleasant response (via presentation of a positive prime) may increase the probability that participants exercise a pleasant response toward the target on that particular trial.

Some experiments have rendered support for response-based models of priming. For example, evaluative priming effects are typically stronger when the targets are classified by their evaluative connotations, whereas semantic priming effects are stronger when the targets are classified by their semantic connotations (see Klauer & Musch, 2003). However, some researchers have argued that spreading activation models could account for these results, by assuming that the classifications required by the task make the evaluative or semantic connotations of the primes more salient (cf. Fazio, 2001). As Ferguson and Bargh (2003) have noted, the general consensus among researchers is that few, if any, studies have conducted definitive tests of the viability of spreading activation versus response-based models of priming. At various points throughout the rest of this article we discuss the implications of our research for this issue.

Potential implications for explaining contrast effects. As was the case with spreading activation models, response-based models cannot account for implicit contrast effects without assuming that the actual evaluative connotations of the neutral primes are altered by the task context. This is because these models assume that primes facilitate positive responses (for example) to the extent that those primes actually are, in fact, evaluatively positive. Primes that are truly neutral should not produce evaluative priming.

Nonetheless, consideration of such models suggests an alternative way of accounting for our contrast effects. For example, take

[2] The effect described here is equivalent to an assimilation effect, in which judgments of the target stimulus are biased in the direction of, rather than away from, the contextual stimuli (Eiser, 1990). In the General Discussion section we offer a brief discussion of these types of effects with regard to our paradigm.

the neutral primes that were paired with the positive primes in Experiment 1. If one assumes that the neutral primes are somehow tagged with a negative response (in spite of the fact that they are evaluatively neutral), this could explain why the neutral primes exhibited properties normally observed with negative primes. That is, it could explain how a neutral prime could facilitate negative responses, without having to assume that the neutral prime was actually perceived as being negative. Hence, the response-based view opens the door, at least, to the possibility that the effects observed for the neutral primes could have come about with no real change in their evaluative implications.

However, this leaves a great deal of the story unanswered. Most important, it does not specify how or why the neutral primes would come to be tagged with any responses in the first place. This is relevant to a more general weakness of response-based models of priming (for a related discussion, see Ferguson & Bargh, 2003). Response-based models offer a reasonably clear account of priming effects once it is established that a particular response has been linked to the prime and target. Yet these models are vague when it comes to explaining how these responses become linked to these stimuli in the first place.

Below, we propose a theoretical framework that builds on some of the core assumptions of response-based models of priming, especially the idea that priming effects ultimately involve compatibility of responses between primes and targets (Klauer & Musch, 2003). As we discuss, however, our framework is not simply a restatement of these earlier models. On the contrary, our framework goes beyond these previous formulations in several ways.

A Response-Mapping Framework

As we noted above, response-based models cannot, in their present form, offer a novel theoretical account of the contrast effects observed in Experiments 1 and 2. Our framework was specifically designed to address this issue. We propose that both explicit and implicit measures involve a fundamental process of *response mapping*, about which we say more presently. We use this framework retrospectively as well as prospectively, insofar as we (a) account for the findings obtained in Experiments 1 and 2 and (b) use this framework as a basis for proposing and testing several novel predictions that would not otherwise be forthcoming with extant models of priming and social judgment.

Response Mapping in Explicit Judgment Domains

At the beginning of this article we briefly alluded to scale-based interpretations of contrast effects (e.g., Parducci, 1965; Upshaw, 1969). Although this is not always stated explicitly, the logic of these models rests on a general assumption about the way that people use response scales. In particular, one must assume that (a) people are motivated to make meaningful perceptual discriminations among the stimuli presented and that (b) people are attempting to do the best they can given the types of response options that are available to them at the time (cf. Wedell et al., 2007).

These assumptions may seem rather self-evident. However, they have some surprisingly important implications. Returning to a variant on the simple example used earlier, suppose that you have been asked to judge the size of various stimuli and that you have been given only two response categories to express those judgments, *small* and *large*. Suppose further that in the context of these ratings, you have already used the response option "small" to refer to very small stimuli, such as single cell organisms. Now suppose that you are asked to judge the size of a mouse. How should you rate this stimulus? Mice are normally rated as small, but use of that response option in this particular rating context would violate participants' motivation to successfully convey that they have, in fact, recognized that mice are bigger than amoebas. Hence, one could say that the mouse is assigned the "large" rating by default, because the other response option—small—has already been reserved for the other stimuli being considered in that task. When pressed on the matter, we suspect that most respondents would admit that they do not really think that mice are "large." However, respondents might respond (with perhaps some indignation) that they were essentially doing the best they could with the response options that were provided for them by the experimenter.

We have deliberately chosen a simple example to illustrate what turns out to be a fairly complex set of issues, which have been explored in a number of different judgmental models, most notably range–frequency theory (Parducci, 1965). One of the features of this model is that it assumes that overt ratings can reflect a compromise between a frequency principle (a bias toward using the available response alternatives equally often) and a range principle (a motivation to accurately match the range of responses to the range of underlying stimuli). (See Wedell et al., 2007, for a related discussion.) Range–frequency theory represents a powerful, but rather complex, model that is capable of addressing judgmental settings far more complex than the simple example used here, and full appreciation of it requires understanding of some rather complex mathematical modeling assumptions.

Nevertheless, for the present purposes, the critical assumption of this and other "response-based" models of judgmental contrast (Biernat, 2005) is that human judgment often involves a basic process of *response mapping*. We define response mapping as the process by which people come to associate a particular response (e.g., "good," "2," "pretty") onto a given stimulus. As part of this process, the mapping of response options onto extreme stimuli can displace ratings of less extreme stimuli onto other sorts of responses, and this displacement essentially describes a contrast effect. In other words, mice can be *large*, and elephants can be *small*, provided that the category labels that are normally used to refer to these stimuli have already been assigned to some other even more extreme stimuli in that particular setting.

Response Mapping in Implicit Judgment Domains

The basic principles described above can, in principle, be applied to implicit judgment domains, including priming paradigms. Just as in the case of explicit judgments, priming paradigms can involve mapping of response options onto primes. Note, however, that there are a few important differences that characterize response mapping here as it might occur in priming paradigms.

In explicit rating paradigms, response mapping arises as a part of participants' intentional goal, which is to respond to and rate the target stimuli. The situation in priming paradigms is somewhat more complex. Participants' responses to the *targets* are part of their primary, intentional goal (e.g., to classify words as positive and negative). However, participants are also responding to the *primes*, even though they are not supposed to be attending to these stimuli at all. This creates an interesting state of affairs in evaluative priming paradigms: While participants intentionally sort the targets into "good" and "bad" categories, they are, at the same time, unintentionally sorting the primes into "good" and "bad" categories as well. That is, the response-mapping framework proposes that participants unintentionally impose the response categorization scheme onto the primes. The primes become associated, or "mapped," with the responses as a result of this process.

Under normal circumstances, the response-mapping assumption is not particularly noteworthy, in the sense that researchers almost always present an exactly balanced ratio of positive and negative primes. In fact, we are aware of very few studies in the published literature that do not follow this general rule. In other words, if participants end up mapping the positive primes with positive responses and the negative primes with negative responses, this could be seen merely as a validation of the initial assumptions of the experimenter.

However, note what happens when participants are presented with an array of priming stimuli consisting of extreme stimuli (e.g., unambiguously negative primes) along with an array of relatively neutral stimuli, such as was the case in Experiments 1 and 2. According to the response-mapping framework, participants will attempt to impose the good/bad classification scheme onto *all* of the primes, even though half of the primes are actually neutral. Clearly, the neutral primes do not fit either of the response labels perfectly. Nevertheless, there is a clear evaluative distinction between the neutral primes and the other extreme primes that are presented in the task. As a result, participants will end up mapping the neutral primes with the "good" or "bad" category labels, using the label that is most appropriate given the context. Just as a mouse may be mapped onto the "large" response in order to distinguish it from amoebas, a neutral prime may be mapped onto the "negative" response to distinguish it from unambiguously positive primes. Hence, when the neutral primes are embedded amongst the positive primes, the favorable response option ("good") is mapped onto to the positive primes, leaving the remaining ("bad") response option for the neutral stimuli. Conversely, when the neutral primes are presented along with the negative primes, the unfavorable response option ("bad") is mapped onto the negative primes, leaving the remaining ("good") response option for the neutral stimuli, by default.

In order to understand the essential elements of our framework, it can be heuristically useful to regard these mappings as involving two different stages. (The term "stage" is used in a general sense here; the temporal ordering of these stages is not crucial for present purposes.) One stage of the response-mapping process involves mapping the "good" and "bad" responses onto the target stimuli. This mapping is a direct consequence of the kinds of instructions that participants are typically given on priming tasks (e.g., "whenever you see a positive word, hit the 'good' key"). The other stage of the response-mapping process involves mapping responses onto the primes. For the reasons noted above this can, in some cases, involve mapping a positive or negative response onto a priori neutral primes.

This implicit response-mapping process is the most novel aspect of our framework and, hence, it is useful to emphasize a few important elements of it here. For one thing, decades of work in the priming literature have shown that priming effects reflect an excellent example of unintentional processing (Bargh, 1994). By the same token, we assume that the process of mapping responses onto the primes is similarly unintentional. The unintentional nature of this process as it applies to the primes is important to keep in mind, because certain aspects of this process may not seem particularly "rational." In particular, why would participants do something so illogical as map stimuli onto a response that has no relation to the stimuli at all? There are two answers to that objection. First, researchers have long recognized that automatic processing can involve the use of categories in ways that do not always follow strict rules of logic and rationality. Second, even if participants were fully aware of how they are mapping responses onto the primes, this involves considerations not altogether different from those pertaining to explicit ratings. In the same way that people may not "really" think of elephants as particularly small, they may not "really" think of wooden stools as positive or negative.

Before summarizing the critical elements of our framework, it is also worth noting that our framework assumes that participants essentially treat the primes as a distinct stimulus group, apart from the targets. In other words, the process of mentally classifying the primes into "good versus bad" categories occurs within the prime stimulus set, apart from the mental classification of the targets. These considerations are discussed in more detail in the General Discussion section.

A Summary of the Response-Mapping Framework

Although our framework may seem complex, the essential details are actually fairly straightforward and can be summarized as follows.

1. Explicit and implicit attitude tasks involve a basic process of response mapping, which we simply define as the process of selecting any available response and using it to refer to a given stimulus.

2. In the case of explicit rating tasks, this process is part of an intentional effort to judge the target stimuli. In the case of implicit tasks, the response-mapping process encompasses both the *intentional* goal to respond to the targets and well as the *unintentional* imposition of those response categories onto the primes.

3. The response-mapping process is not necessarily "rational," and participants may not even be aware that that such a process is even occurring, especially as it concerns the way that participants process information about the primes. Rather, response mapping can be better understood as an unintentional process of imposing the response categories onto the primes.

4. Relatively extreme and/or unambiguous stimuli will tend, by their very nature, to "dominate" the response-mapping process. For example, if participants are presented with a randomized array of unambiguously negative and neutral stimuli, the former stimuli will tend to stand out and will be assigned the most negative response option available.

5. Contrast can be understood as a process of judgmental displacement as a result of the preceding step. In other words, after respondents have "used up" the available response options to refer to the unambiguous stimuli, less extreme or midrange stimuli will be assigned to whatever

> remaining response options are available. This explains both why recycling (a "midrange" social issue) can be rated as either important or trivial depending on whether the available response options have been taken up by extremely important or extremely trivial issues (nuclear war or sidewalk spitting, respectively). This also explains why a towel could elicit automatic priming effects analogous to those associated with positive or negative stimuli, depending on the nature of the other primes presented in that context. Note, however, that such contrast effects are contingent on the response options that are available in the task, a point to which we return in Experiment 3.

Although the findings obtained in Experiments 1 and 2 are consistent with this framework, some critical, additional tests are needed to validate it, rule out some alternative explanations, and also to clarify some important aspects of our theoretical assumptions. The experiments to be reported in the rest of this article were designed with these objectives in mind.

Experiment 3

All aspects of the priming stimuli, including the creation of two priming contexts (cf. Figures 1A and 1B) were identical to those used in Experiment 1. However, we modified the target words and response options. In the earlier study, participants were presented with just two kinds of words (positive or negative) and two kinds of response options (positive or negative). Here, participants were presented with three types of target words, which included the same positive and negative words used in the earlier studies, along with neutral words (*initially*, *heels*, *inclined*, *intercom*, *stands*, *ounce*, *stretched*, and *hold*). Accordingly, participants were also presented with three corresponding response options (*positive*, *negative*, *neutral*) with which to categorize these targets. These modifications may seem modest, but in fact they provide a strong test of the response-mapping framework.

Predictions

The explicit goal of the participants in this task may be seen as an intentional effort to map the three available responses (*positive*, *negative*, *neutral*) onto the three types of target word types. More important, the response-mapping framework additionally assumes that participants should unintentionally map these response options onto the primes. This leads to a counterintuitive prediction that is entirely specific to the response-mapping perspective and cannot be easily accounted for by a spreading activation view.

Consider the case in which participants are presented with positive primes and neutral primes. Because participants now have three response options with which to respond to the words, this means that they can map positive responses onto the positive primes and, more importantly, they can also map a neutral response onto the neutral primes. A similar state of affairs pertains to cases in which participants are presented with negative and neutral primes. Once again, participants have three response options, and hence participants are likely to map the negative response onto the negative primes and a neutral response onto the neutral primes. The upshot of these considerations is that whenever a neutral prime (e.g., a picture of a towel) appears on a given trial, such primes should significantly facilitate responses toward neutral target words, because they have been mapped onto the same response. For example, on trials containing a neutral prime (e.g., a picture of a towel), the presentation of this prime should significantly facilitate responses toward an utterly neutral word (e.g., *initially*).

A response-mapping perspective provides a fairly parsimonious account of this type of finding, as it requires only the assumption that participants have mapped the same type of response onto the prime and target. On the other hand, a spreading activation account would have a great deal of difficulty accounting for these findings. Most problematic, perhaps, is the challenge in defining what sort of "spreading activation" process, exactly, might exist that could account (for example) for why a picture of a towel would facilitate responses toward the word *initially*.

Finally, contrast effects were not expected to arise from the neutral-object primes, even though the priming context in which they appeared was the same as in Experiment 1. This is because a neutral response option was provided in both conditions, and so the neutral primes should not be displaced onto the negative or positive response. Instead, neutral primes should facilitate responses only toward neutral words, and this should be true regardless of the context in which the neutral primes appear.

Method

Participants

In this study, 129 students participated in return for partial course credit. One participant was removed from analyses due to making an unusually high number of errors (greater than 2 *SD*s above the mean), leaving a total of 128 participants.

Procedure

Experiment 3 was identical to Experiment 1 (see Table 1), save for three details. First and most important, participants were asked to respond to good, bad, and neutral words (rather than just good and bad words). On the computer keyboard, three adjacent keys were labeled "good," "bad," and "neutral," and participants were instructed to use the index, middle, and ring fingers of their dominant hand to make their responses (instructions were communicated via experimenter and computer, to ensure maximum compliance). Also, the ordering of the relevant response on the keyboard was completely counterbalanced between subjects, to include all possible combinations of order. Another minor difference between Experiments 1 and 3 was that participants in Experiment 3 received feedback for both fast and slow responses (rather than just slow responses): In addition to seeing the words "Faster Please!" for responses made outside the 500 ms window, participants additionally saw the word "OK!" appear in green for responses that were within the window (both messages remained on the screen for 400 ms). Finally, every prime was presented with every word one time in each block, for a total of 192 trials per block (8 primes $\times$ 24 words) and 384 trials total (192 $\times$ 2 blocks).

Stimuli

Stimuli consisted of 12 of the pictures used in Experiment 1: four positive animals, four negative animals, and four neutral objects. Target words included the same 16 good and bad words used in Experiment 1, in addition to eight neutral words (e.g., *initially*) that were drawn from the same word pool that was used in Experiment 1. All the words were matched on length in letters, frequency in spoken language, and orthographic neighborhood. Additionally all words were pretested for emotional intensity and valence (Larsen et al., 2006).

Analogous to Experiment 1, we formed overall indices of error rates, except that in this case we formed three (rather than two) separate error indices corresponding to the three types of words in this study (positive, neutral, and negative).

Results

The analytic approach used in this study was similar to that of Experiment 1. However, it should be noted that the types of comparisons of main interest here were somewhat different. In Experiment 1, the comparisons of main interest focused on a particular type of prime and then compared error rates across different types of words. In Experiment 3, however, the theoretical issues were such that the more important comparisons involved focusing on a particular type of word (especially, neutral words) and investigating what types of primes did or did not facilitate reactions toward it, depending on the context in which those primes appeared.

Preliminary Analyses on the Positive Versus Negative Primes Only

In some ways our three-response design deviated significantly from the usual procedure used in evaluative priming studies. For this reason, it is useful at the outset to demonstrate that our task replicated the usual priming effects observed when positive or negative primes are paired with positive or negative words. These initial analyses included only the block in which positive and negative primes were presented. Analyses on this block revealed the expected Prime Type × Word Type interaction, $F(2, 126) = 17.41$, $p < .001$, demonstrating that the standard evaluative priming effect was replicated within this study. As expected, on trials involving positive words, participants revealed far fewer errors when these words were preceded by positive, relative to negative, primes (Ms = 11.67 vs. 13.82), $F(1, 127) = 34.82$, $p < .001$, whereas this pattern was reversed on trials involving negative words (Ms = 11.78 vs. 10.50), $F(1, 127) = 12.74$, $p < .01$. The primes elicited no difference in errors on neutral words (Ms = 13.60 vs. 14.04 for positive vs. negative primes, respectively; $p > .20$).

Main Analyses

Of greater interest was the pattern of data involving the blocks that included neutral primes. These data are most easily conveyed in tabular form, which we present in Table 2. This table displays error rates as a function of target word type and prime, both for the negative priming context (top) as well as the positive priming context (bottom).

Independent of all other considerations, error rates for negative words tended to be lower compared to those for positive or neutral words. This was true in both the negative priming context (Ms = 9.65 vs. 12.35 vs. 12.11 for negative, neutral, and positive words, respectively), $F(2, 61) = 8.53$, $p < .01$, as well as the positive priming context (Ms = 10.82 vs. 12.02 vs. 13.07), $F(2, 63) = 3.71$, $p < .05$. As seen in Table 2, this main effect of word type was apparent across all four rows of data. In other words, the relatively low rates of errors for negative words (vs. neutral or positive words) held regardless of the specific type of prime and the context in which those primes appeared.

Table 2
Mean Error Rates as a Function of Prime and Target Word Type for Experiment 3

Context and prime type	Target word type: Neutral	Negative	Positive
Negative priming context			
Neutral primes	11.59	10.09	11.79
Negative primes	13.10	9.21	12.42
Difference	-1.51^{*}	$0.88^{\dagger}$	-0.63
Positive priming context			
Neutral primes	11.44	10.74	13.56
Positive primes	12.60	10.90	12.58
Difference	-1.16^{*}	-0.16	$0.98^{\dagger}$

Note. Table does not present data from the priming block that contained only positive and negative primes (see text).
$^{\dagger}p \leq .10$. $^{*}p < .05$.

Of greater interest, however, were possible differences in the relative rates of errors across primes and the contingency of such effects as a function of priming context. Hence, we conducted a series of theory-driven contrasts to test the specific predictions of our framework. As we shall show, the finding of greatest interest was the fact that responses to neutral words were reliably facilitated only by neutral primes.

Negative priming context (neutral vs. negative primes). As seen at the top left column of Table 2, error rates on neutral words were significantly lower when participants were presented with neutral (vs. negative) primes, $F(1, 62) = 9.85$, $p < .01$. In contrast, the reverse pattern occurred for negative words (middle column), which revealed a tendency for error rates to be lower when participants were presented with negative (vs. neutral) primes ($p < .09$). As for the positive words (right column), there was no reliable effect of prime ($p > .15$). This asymmetry was confirmed by a significant Prime Type × Word type interaction, $F(2, 124) = 5.81$, $p < .01$.

Positive priming context (neutral vs. positive primes). Analogous to the previous set of analyses, error rates on neutral words were lower when participants were presented with neutral (vs. positive) primes, $F(1, 64) = 4.10$, $p < .05$. There was also a tendency for positive primes to produce lower error rates on positive words, although this effect did not reach significance ($p = .10$). The effect for negative words, on the other hand, was negligible ($F < 1.0$). This asymmetry was again confirmed by a significant Prime Type × Word Type interaction, $F(2, 128) = 3.49$, $p < .05$.

Additional Analyses

If contrast effects had occurred in this type of paradigm with the neutral primes, they would have manifested themselves by (a) neutral primes facilitating negative responses in the positive priming context and (b) neutral primes facilitating positive responses in the negative priming context. In this experiment, however, we did not expect such effects to occur because neutral primes were expected to be mapped onto the neutral responses, and this should be true regardless of the context in which those neutral primes occurred. The data showed that this was in fact the case. An analysis of the neutral primes across the two priming contexts revealed a main effect of word type, $F(2, 125) = 6.97$, $p < .01$, but the priming context had absolutely no effect on the reactions elicited by the neutral primes ($F < 1.0$ for the two-way Word Type × Context interaction).

Finally, recall that the negative and positive primes were presented to each participant in two separate priming blocks, as in Experiment 1 (see Table 1). For the sake of completeness, we also conducted analyses to determine whether either the positive or negative primes elicited different reactions across the two priming blocks in which they appeared. Neither two-way Word Type × Block interaction was significant ($p > .20$).

Discussion

One of the predictions borne out of our response-mapping model is that neutral primes could facilitate responses to neutral words. This is exactly what we observed in Experiment 3. It is difficult, if not impossible, to explain this finding using spreading activation models. What preexisting cognitive associations would connect, for example, a barstool with the words *initially*, *ounce*, or *hold*? Response mapping, on the other hand, can explain this finding in a fairly parsimonious way: By providing participants with a neutral response option, we allowed them to use the neutral response option to refer to both neutral primes as well as neutral words.

Experiment 4

One of the more provocative implications of our framework is that automatic priming effects can be driven by mechanisms that are independent of the actual valence of the prime. That is, neutral primes may, under some circumstances, facilitate responses toward good or bad words, independent of the fundamental evaluative meaning of the priming stimulus itself (which remains unchanged). We have already provided fairly strong evidence for this assumption. In particular, Experiment 3 suggested that contrast effects have nothing to do with any fundamental change in the evaluation of towels but rather, to changes in the types of responses that are mapped onto these stimuli.

The overriding goal of Experiment 4 was to submit our model to an even stronger test. To understand the rationale for this study, it is important to recognize that there is nothing about our framework that limits its applicability to evaluative tasks. Semantic responses can be mapped onto stimuli as well. Just as evaluative priming effects can be driven by mechanisms that are independent of the actual valence of the prime, it is also true that semantic priming effects can be driven by mechanisms that are independent of the actual semantic implications of the prime. This aspect of our model leads to some very counterintuitive predictions.

Before getting to a detailed description of the design, a concrete example may be useful. Imagine that you are presented with a picture of a towel. Like most towels you have seen, there is nothing whatsoever in this picture that is semantically related in any meaningful way to the category of either food or flowers. This being the case, it seems unlikely or even absurd to suggest that a picture of a towel could reliably facilitate responses toward food-related words (e.g., *hungry*, *eat*) or flower-related words (e.g., *tulip*, *bouquet*) in a semantic priming task. Yet, these are precisely the sorts of facilitation effects that could be accounted for by our response-mapping framework. That is, a towel could remain resolutely towel-like but at the same time reliably facilitate reactions toward either food or flower targets. Although this type of effect does not represent a "contrast effect" in the sense that that term is usually employed, the mechanisms responsible for it are directly related to those driving the effects observed in Experiments 1–3, as we shall show.

Overview of Design

In order to facilitate understanding of some nuances of our design, Table 3 provides a guide to the general design of the experiment. In this experiment, participants were always presented with a randomized list of two types of targets words: flower-related words (e.g., *bloom*, *grow*) and food-related words (e.g., *eat*, *hungry*). As each word appeared on the screen, participants were simply asked to categorize these targets by pressing one of two keys on the keyboard, marked "flower" or "food." For half of the participants (see top of Table 3), half of the primes were related to flowers (and, hence, were semantically related to half of the target words). However, the other half of the primes were the same neutral primes used in our previous studies. Because our task is no longer evaluative, it is useful to refer to these stimuli as *object primes*. These object primes were not related to either of the two classes of target words (food or flowers). An analogous set of conditions applied to the other half of the participants (see bottom of Table 3), except that they were presented with equal numbers of food-related primes and object primes. Once again, note that the food primes were semantically related to half of the target words, but the object primes were not related to either class of targets.

Predictions

Extension of the response-mapping framework to this paradigm leads to two predictions. First, in the flower context (top of Table 3) participants will map flower responses onto flower primes, and in the food context (bottom of Table 3) participants will map food responses onto food primes. Second—and more important—the object primes will end up being mapped with whatever response option is "left over" after the other option is assigned to the flower or food primes. In the flower prime context, participants will map

Table 3
Mean Error Rates as a Function of Prime and Target Word Type for Experiment 4

Context and prime type	Target word type: Flower	Target word type: Food	Difference
Flower priming context			
Object primes	15.69	13.41	2.28**
Flower primes	12.29	17.88	−5.59**
Food priming context			
Object primes	14.50	15.09	−0.59
Food primes	18.82	12.35	6.47**

** $p < .01$.

the *food* response onto the object primes. This means that the presentation of a towel prime (for example) should facilitate responses toward food-related words. Conversely, in the food priming context, participants will map the *flower* response onto object primes. Hence, in this latter condition, that same towel prime will now facilitate responses toward flower words.

It is worth emphasizing that our design provides a far more rigorous test of our model than could be offered by an evaluative priming paradigm. In evaluative priming paradigms, it can be extremely difficult to rule out the possibility that there could have been some change of evaluative meaning in the priming stimuli. In contrast, semantic meaning is more clear-cut. Unless your doctor has put you on a very strange diet, towels are not food, nor are they closely related to food, regardless of the context in which they are presented. The upshot of these considerations is that this design provides a much stronger test case of our response-mapping framework than any of the preceding experiments.

Method

Participants and Design

Participants were 113 undergraduate students who participated in return for partial course credit. Every participant completed two priming blocks: In one block, participants were asked to classify words that were either flower related or food related and viewed primes that were either pictures of household objects or food. In the other block participants made the same word judgments, but the primes were either household objects or flower pictures. Half the participants received the object–food priming block first, and the other half received the object–flower priming block first. The order of blocks produced no significant effects (all Fs < 1), so the following analyses collapse across this variable.

Stimuli

As in previous experiments, primes were selected from the IAPS (Lang et al., 2008). There were a total of 18 primes. Of these, six consisted of pictures of common household objects. In addition to the object primes, participants were presented with six food primes (e.g., pictures of French fries, hamburgers, ice cream) and six flower primes (e.g., pictures of bouquets, fields of flowers). Target words included five food-related words (*hungry*, *eat*, *taste*, *dinner*, *full*) and five flower-related words (*tulip*, *rose*, *bloom*, *daisy*, *grow*). Hence, this study consisted of three classes of primes (objects, food, flowers) and two classes of target words (food, flowers).

Procedure

As in Experiments 1 and 3, participants were told that they would see a briefly presented picture that would be followed by a word. Their task was to ignore the picture and identify target words as either food-related or flower-related as fast as possible, using appropriately labeled keyboard keys. The position of the keys was counterbalanced between subjects. Each participant completed two priming blocks, one that contained object and food primes (hereinafter the *food priming condition*) and another than contained object and flower primes (henceforth the *flower priming condition*). The blocks were presented in counterbalanced order and separated by a brief break. Preliminary analyses revealed no effect of counterbalancing order (all Fs < 1.0) and hence our analyses collapse over this factor. The timing of the trials was identical to that in Experiments 1 and 3. Within each task block, primes and targets were presented in randomized order. Each prime was paired with each target one time, for a total of 120 trials per block (12 primes $\times$ 10 target words) and 240 trials total (120 trials $\times$ 2 blocks). After completing the two task blocks, participants were debriefed and dismissed.

Results

Error indices were calculated in a manner analogous to Experiments 1 and 3, except that in this case, errors consisted of erroneously categorizing food-related words as flowers and flower-related words as food. This resulted in two error indices for each prime: one for errors on flower-related words and another for errors on food-related words.

Analyses of Object Primes

We predicted that the object primes would facilitate responses either toward flowers or food, depending on the context in which these object primes appeared. As seen in Table 3, our prediction was confirmed. The object primes, when presented with other flower primes, tended to facilitate food relative to flower judgments, whereas the opposite was true of object primes presented with other food primes. This pattern was confirmed by analyses of the object prime only, in which we submitted the error data for the object primes to a 2 (task context: flower vs. food) $\times$ 2 (word type: flower vs. food) multivariate analysis. Results revealed the predicted two-way Task Context $\times$ Word Type interaction, $F(1, 112) = 7.10$, $p < .01$. Follow-up analyses within each of the two priming contexts also confirmed expectations. In the flower priming condition, object primes facilitated food relative to flower judgments in the flower priming condition (difference in error on food vs. flower words), $F(1, 112) = 9.18$, $p < .01$. The opposite pattern arose in the food priming condition, although the difference in food versus flower error did not reach significance in this case ($p > .10$). Hence, even though the object primes and target words were semantically unrelated, the object primes still appeared to influence participants' responses to the words as if they were related.

Analyses of Flower and Food Primes

We predicted and found that the effects obtained for the object primes, reported above, could coexist independently of, and in addition to, the types of semantic priming effects typically observed in this area. This was confirmed by further analyses of the flower and food primes only in combination with the flower and food words. A 2 (prime: flower vs. food) $\times$ 2 (word type: flower vs. food) within-subjects ANOVA revealed the expected two-way interaction, $F(1, 112) = 106.10$, $p < .001$. Simple effects tests confirmed that this interaction was driven by flower primes facilitating flower judgments relative to food judgments, $F(1, 112) = 46.87$, $p < .001$, and food primes doing just the opposite, $F(1, 112) = 52.40$, $p < .001$.

Discussion

The present experiment was designed to test a highly counterintuitive prediction of our model, namely, that semantic priming effects can occur on trials in which the prime and target have absolutely no semantic relation to one another. This is exactly what we found. In one task condition, household objects facilitated responses to food-related words, whereas in another condition household objects facilitated responses to flower-related words. Our response-mapping framework can easily account for such findings, but we are not aware of any other theoretical model of priming that could have specifically anticipated these sorts of effects.

It is worth noting in passing that the facilitation effects elicited by the flower and food primes were much larger than effects elicited by the object primes. The respective difference in effect size does not qualify our conclusions with regard to the response-mapping hypothesis, but it does potentially provide insight into the contributions of various mechanisms to overall priming effects. Indeed, response-mapping processes can apparently occur in addition to, and independently from, the effects of semantic priming.

Experiment 5

Experiments 3 and 4 provided convincing evidence in favor of the response-mapping framework. However, an alternative explanation still remains, and the purpose of this experiment was to rule out this alternative. According to the response-mapping model, the neutral prime is associatively paired with a good or bad response. Another possibility, however, is that the neutral primes are not paired with any response at all, but instead, that these primes reveal an underlying response bias.

In order to understand this alternative view, consider a case in which half of the primes are positive and half are neutral. In this case, the positive primes could be said to dominate the pattern of responding, causing participants to use the positive key more than they otherwise would. Recall, however, that participants have a general "set" or expectation that they will end up using the available response options about equally (Wedell et al., 2007). As a result, participants might try to correct for the fact that they are using the positive key "too much" by forming a response bias favoring the negative key whenever the dominant prime does not appear. This perspective makes a prediction similar to our own framework, namely, that trials containing the neutral prime would elicit the negative response option. However, this interpretation differs from the response-mapping framework, because it does not need to assume that there is any actual mapping of responses onto the neutral primes at all.

In order to address this alternative hypothesis, we conducted an experiment that was similar to Experiment 2, but this time the task included a small number of control trials, in which no primes were presented at all. (Hereinafter, we refer to such trials as *no-prime controls*.) If the contrast effect for neutral primes is the result of an overall response bias, and not response mapping, then the contrast effect should emerge for both the neutral primes and on these no-prime control trials. However, if the neutral primes actually become associated with a particular response, and it is actually this mapping between prime and response that drives the contrast effect, then a contrast effect should emerge for the neutral primes, but not for trials that contain no prime.

Method

Participants and Design

Participants were 43 undergraduate students who participated in return for partial course credit. There was one between-subjects factor pertaining to the nature of extreme primes that were presented in the task, which were either positive or negative. There was also one within-subjects factor, pertaining to the nature of the primes that were presented in each task (positive/negative, neutral, and no-prime controls).

Stimuli and Procedure

All participants were randomly assigned to complete one of two AMP tasks. In that task, participants viewed either (a) positive, neutral, and control primes or (b) negative, neutral, and control primes. This experiment was identical to Experiment 2, save for two details. First, each of the four unambiguous and four neutral primes were presented nine times (rather than six times, as in Experiment 2). Second, this experiment included nine trials in which no prime was presented at all, the no-prime controls. For these latter trials, the screen remained blank for the amount of time that a prime would normally appear. Hence, in this experiment, there were a total of 81 trials (8 prime images $\times$ 9 presentations, plus 9 control trials).

Results and Discussion

A 2 (task context: positive vs. negative) $\times$ 3 (prime type: unambiguous, neutral, and control) mixed-model ANOVA revealed a main effect of prime type, $F(2, 40) = 5.61$, $p < .01$, indicating that, collapsing across task context, the control trials elicited relatively more pleasant responses than either the unambiguous or neutral primes. This main effect was qualified by a significant two-way interaction, $F(2, 40) = 5.41$, $p < .01$. As depicted in Figure 4, this interaction was driven by the unambiguous and neutral primes, which successfully replicated the expected contrast effect. Further planned comparisons confirmed that the neutral primes elicited relatively more pleasant responses in the negative context than in the positive context, $F(1, 42) = 4.21$, $p < .05$.

Of greater interest was the test of an alternative "response bias" explanation of contrast effects, which predicts that the contrast effect will emerge for both neutral and control trials. The response-mapping hypothesis, however, predicts that the contrast effect will emerge for the neutral primes, but not the control primes. The data clearly support the response-mapping hypothesis. The control primes showed no evidence of contrast. They elicited the same proportion of pleasant responses, regardless of the task context, $F(1, 42) = 0.08$, $p > .70$. If any response bias was revealed in this experiment, it is that in the absence of priming, participants tend to favor the "pleasant" key. However, this pleasant key bias was not qualified by the task context, and therefore the response bias explanation cannot account for the contrast effects observed in the present research.

Experiment 6

One of the emerging themes of our research is that primes that *appear* to reflect strong preexisting automatic evaluations may, in fact, be evaluatively neutral. Thus far we have demonstrated this principle mostly with neutral and/or inanimate stimuli. For this

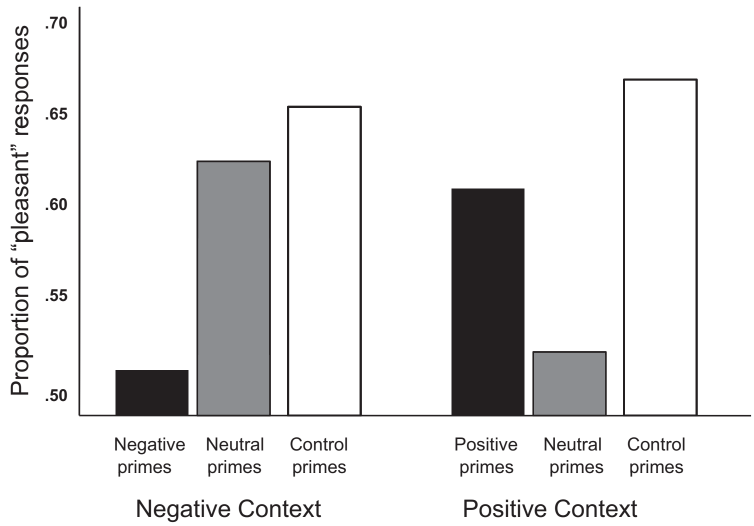

Figure 4. Experiment 5 results. This figure displays the proportion of "pleasant" responses following each of the primes, by task context. Greater number of "pleasant" responses following a prime indicates that the prime elicited a relatively positive response bias.

reason, it may not be immediately apparent how our work is relevant to the majority of work in the social psychological literature. Researchers in that area have traditionally been more concerned with detecting implicit attitudes with socially meaningful stimuli, especially with the goal of detecting the presence of outgroup prejudice. As we describe below, the issues raised in our research are very much related to this work.

To illustrate, consider an evaluative priming task designed to test for the presence of prejudice toward the elderly among young participants. A typical approach would be to design a priming task in which participants are presented with rapidly alternating trials containing old (outgroup) versus young (ingroup) faces. Suppose that the task showed evidence of negative priming effects with the old faces and positive priming effects with the young faces. One might conclude, in this case, that the experiment demonstrated evidence of automatic negative prejudice toward the elderly. According to our framework, however, this conclusion may not necessarily be correct. In particular, it could be that the young faces are indeed associated with strong preexisting positive associations but that the old faces are actually *neutral* (i.e., possess few if any associations at all). The unfavorable priming result for the elderly faces could be the result of a contrast effect, driven by the juxtaposition of the elderly faces against the young faces. Note, however, that the contrast effect in question does not entail an actual change in evaluative meaning, but rather reflects the assignment of the negative response to the elderly primes in the context of that particular task.

This ambiguity pervades much of the evaluative priming literature, because outgroup primes are nearly always "pitted" against ingroup primes. As other scholars have noted, this problem can be somewhat ameliorated by designing a task in which ingroup primes are excluded from the task altogether (cf. Karpinski & Steinman, 2006). However, we take this line of reasoning several steps further. According to our framework, valuable information about the presence of preexisting evaluative associations can be obtained by manipulating the type of context in which the prime appears and then observing how that prime "acts" across context. Using this strategy, one can obtain much better leverage in understanding the nature of that prime (and the associations that may or may not be linked to it) than if one measured the properties of that prime in only one context.

In this experiment we explored this idea by randomly assigning participants to conditions in which they were presented with elderly primes in positive, negative, and neutral priming contexts. We expected that valuable leverage regarding the attitudinal characteristics of the elderly primes could be discerned by examining the extent to which priming effects elicited by them would—or would not—vary across condition. As we showed in our earlier studies, stimuli that already possess strong, prepotent evaluative responses tend not to show context effects (i.e., they elicit much the same type of response regardless of the valence of the other primes). If we found that the elderly primes elicited negative priming effects across all conditions, this would provide compelling evidence that, at least for the participants in our sample, the elderly primes do seem to possess strong negative associations. Conversely, if we found that the elderly primes elicited negative priming effects only when placed in the context of positive primes, this would suggest that our young participants do not, in fact, hold any meaningful evaluative associations with the elderly at all.

Method

Participants and Design

A total of 35 undergraduate students participated in return for partial course credit. In a between-subjects design, participants were randomly assigned to one of three conditions pertaining to whether the primes presented in the task consisted of (a) elderly primes and negative primes (positive context), (b) elderly primes and positive primes (negative context), or (c) elderly primes and neutral primes (neutral context).

Stimuli and Procedure

In this study we used a version of the AMP (Payne, Cheng, et al., 2005) that was similar in most respects to the task we used in Experiment 2, with a few key modifications. The timing of trials, instructions, and targets were exactly the same as Experiment 2. In addition, the positive, negative, and neutral primes were identical to those used in our earlier study. As for the elderly primes, we used two female and two male elderly pictures (both with neutral expressions) that were cropped above the shoulder in a manner similar to images used in the race-based priming literature (e.g., Greenwald et al., 1998).

In all task conditions, exactly half of the primes were elderly. The nature of the nonelderly primes varied across condition and consisted either of positive, negative, or neutral primes. The conditions were varied between subjects, so that for one group of participants the primes included elderly and negative pictures, for another group they included elderly and positive pictures, and for a third group they included elderly and neutral pictures. Each prime was presented nine times, for a total of 72 trials in each condition (9 presentations $\times$ 8 primes). After completing the task, participants were debriefed, thanked for their participation in the task, and dismissed.

Results

Data were processed in the same manner as in Experiments 2 and 5. This resulted in computation of three indices for the elderly primes corresponding to the condition in which they appeared (positive context, negative context, and neutral context). In addition, we also computed indices for the context stimuli themselves (i.e., the positive, negative, and neutral primes).

Figure 5 displays the relevant pattern of means for these indices across the positive context (left), neutral context (middle), and negative context (right). Setting aside the elderly primes for the moment, it is worth noting at the outset that the contextual primes, themselves, elicited the expected pattern of results. The positive primes elicited strong positive bias, the negative primes elicited strong negative bias, and the neutral primes fell in between these two extremes. This pattern was responsible for a significant linear trend across the contextual primes, $t(1, 32) = 2.75$, $p < .05$.

Of greater interest, the degree to which the elderly primes elicited positive or negative bias—or no bias at all—depended strongly on the nature of the primes with which they were paired. In particular, the elderly primes elicited strong negative bias when they were embedded amongst positive primes, strong positive bias when they appeared with negative primes, and no discernable bias at all when they were embedded amongst neutral primes. This pattern was responsible for a significant linear trend as well, $t(1, 32) = 2.14$, $p < .05$. Clearly, the types of priming effects elicited by the elderly primes varied as a function of experimental condition.

Although this experiment was not a completely crossed design, data analyses can be also conceptualized as a 2 $\times$ 3 factorial in which one factor—varied within subjects—pertained to the type of prime that appeared in the task (i.e., elderly vs. nonelderly) and the other factor—varied between subjects—refers to the valence of the nonelderly prime (positive, negative, or neutral). These analyses revealed a two-way interaction, $F(2, 32) = 5.83$, $p < .01$, which was driven by both the difference in bias elicited by the contextual primes, as well as the change in bias elicited by the elderly primes across experimental contexts.

Additional planned comparisons revealed that when the elderly primes elicited negative bias, they were just as negative as the unambiguously negative primes ($F < 1.0$). When the elderly primes elicited positive bias, they were just as positive as the unambiguously positive primes ($F < 1.0$). Finally, the elderly primes in the neutral context elicited bias that was no different from that of the neutral primes ($F < 1.0$).

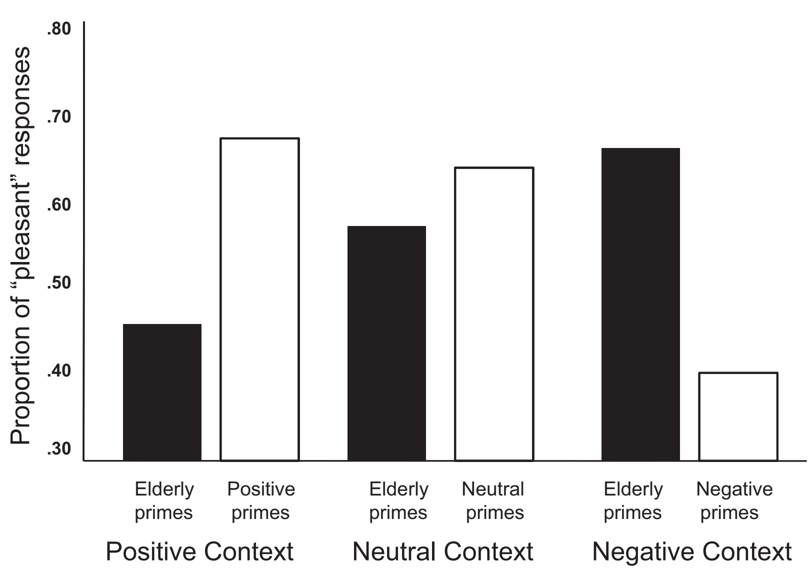

Figure 5. Experiment 6 results. This figure displays the proportion of "pleasant" responses following each of the primes, by task context. Greater number of "pleasant" responses following a prime indicates that the prime elicited a relatively positive response bias.

Discussion

This experiment further demonstrates the validity and utility of our framework, showing its implications for a set of issues of central concern in the social psychological literature over the last 20 years. Like the neutral primes that we examined in earlier experiments, the elderly primes were highly malleable, eliciting strong positive and negative bias depending on the context in which they appeared. Hence, these findings show that our earlier demonstration of contrast extends to the same sort of meaningful social stimuli of interest to social psychologists.

Of equal interest, elderly primes elicited no bias at all when they appeared in the context of neutral primes. This finding provides an excellent example of the appeal and utility of our "contextual" approach. At least with this current sample of participants, our findings suggest that the elderly primes are not "intrinsically negative" in the same sense that unambiguously negative primes are intrinsically negative. If that were so, then we would have found the elderly primes to elicit reliable negative priming effects when they appeared among the neutral primes. However, this did not happen. Taken in combination, the findings of this study demonstrate that the nature and valence of evaluative associations cannot easily be understood by examining the effects of that prime within a single priming context. Rather, far more information can be gleaned by examining the effects of that prime across contexts.

Before turning to the general discussion, it is important to address a potential misunderstanding of our findings. It would be foolish to conclude that all previous demonstrations of automatized outgroup prejudice reflect methodological artifacts, and we do not make this claim. Certainly, there may be circumstances in which outgroup primes may be associated with preexisting negative associations. This could be true even for certain types of elderly primes, depending on the nature of those priming stimuli. For example, it is likely that certain types of elderly faces are extremely stereotypical (e.g., strongly evoke the image of the "grumpy old man"; cf. Livingston & Brewer, 2002), and these kinds of faces may indeed be more likely to possess preexisting negative associations. These considerations may also depend on the type of subject population used and the personal experiences of the participants. While we acknowledge these caveats, they do not qualify the validity of our main point, which is that evidence for automatic priming effects can be highly misleading when such evidence is obtained when a single type of priming context is used. By observing the nature of how primes "act" in different types of contexts, one can learn a great deal more about the nature of the primes—and of priming effects more generally—than otherwise would be the case.

General Discussion

One of the more well-supported principles in the psychological literature is that judgments of any given stimulus are highly dependent on the context in which it appears (Helson, 1947; Volkmann, 1951). However, much of the research supporting this principle has been obtained using explicit measures. By contrast, relatively little is known about these considerations as they might apply to implicit measures, especially those that rely on priming techniques. Across six experiments we explored these considerations using a variety of different types of priming paradigms that involved a broad range of different types of priming stimuli and task objectives.

In pursuing these objectives, we drew from two distinct and heretofore "unconnected" psychological literatures, including classic work on judgmental contrast as well as theory and research on implicit attitudes. This integrative approach enabled us to generate several provocative insights and contributions to these literatures, the nature of which can be summarized as follows:

1. As discussed earlier, explicit measures are relatively limited in their ability to reveal the nature of any associations that might exist with any given stimulus. This limitation is important, because systematic "shifts" in explicit ratings do not clearly indicate whether there has been any actual change in the underlying meaning of the stimulus (cf. Petty & Cacioppo, 1981; Sherman et al., 1978). Use of implicit measures provided us with far more traction on these matters, showing that strong contrast effects can occur, even in the absence of any true change in the underlying representation of the target stimulus.

2. The evidence we obtained for automatized contrast effects enabled us to gain valuable traction on an issue that has recently emerged as a vexing issue for priming theorists, namely, the difficulties in devising competitive tests of spreading activation versus response-based models. The extant literature has had a great deal of difficulty teasing these two models apart because they often make similar if not identical predictions. As it turns out, our own experimental paradigm provided excellent leverage in testing the viability of these models, and our results strongly supported a response-based model, at least within the AMP and evaluative priming paradigms.

3. Our findings go beyond merely validating the basic assumptions of response-based models. As previous theorists have noted (Ferguson & Bargh, 2003) response-based models of priming are generally silent on the processes by which responses become "assigned" to the priming stimuli in the first place. To address this gap, we proposed and validated a response-mapping model that provides a theoretical account of our contrast effects and also addresses several ambiguities in response-based models.

4. Finally, our research revealed that automatic priming effects can arise for different reasons, and sometimes these reasons have little to do with the intrinsic features of the priming stimulus itself. In our view, this point may be of most interest to social psychologists, particularly those who have come to rely on priming paradigms to reveal evidence of "automatized" implicit attitudes. One of the more important implications of our research is that "automatic priming effects" in priming paradigms do not necessarily provide evidence that the primes in question actually possessed any preexisting associations at all.

Taken as a whole, our findings appear to be fairly generalizable. Our results cannot be attributed to an idiosyncratic feature of a particular type of priming task, nor are they specific to evaluative priming paradigms, as we showed support for our model in the semantic priming domain as well. It is also worth noting that we obtained support for our model regardless of whether the primes were "social" or "nonsocial" in nature. Hence, given the heterogeneous nature of our tasks and the diversity of targets and primes to which participants responded, our confidence in the validity of our findings is bolstered moreso than if we had relied on only a single type of paradigm to test the predictions of our model.

A Closer Look at the Response-Mapping Process

As we briefly noted earlier in our article, our framework assumes that participants treat the primes and targets as belonging to distinct classes of stimuli, at least with respect to the response-mapping process. That is, they map the primes and the targets separately; the targets are sorted into response-based categories, and then in a different step, the primes are separately sorted into response-based categories. At first it may seem implausible that primes and targets are mapped separately. However there are three characteristics of these priming tasks that make it likely that this assumption is correct. First, the task instructions explicitly tell participants to treat the targets and primes as separate entities, by asking them to classify the targets but do nothing with the primes. Second, the primes are always presented at the beginning of a trial, whereas the targets are presented at the end, which increases the likelihood that the primes and targets are distinguished from each other. Third, in our experiments the primes were pictures whereas the targets were words, so even the stimuli themselves created circumstances in which participants are likely to treat the primes and targets as separate groups of stimuli. Taken in combination, these considerations make it more likely that participants would apply their categorization scheme to the targets and primes separately.

There could, in principle, be other ways of cognitively organizing primes and targets during this response-mapping process. For example, it is possible that participants treat the primes and targets as belonging to the same stimulus set. To illustrate, consider a task in which the stimuli include positive and neutral primes and positive and negative targets. If participants treated the primes and targets as one large group of stimuli, the positive primes and targets would be mapped onto the positive response, and the negative targets would be mapped onto the negative response. However, the neutral prime would (according to this alternative view) presumably remain unmapped, because the neutral primes are neither as positive as the unambiguously positive targets and primes nor as negative as the unambiguously negative targets. Hence, if the primes and targets are treated as one large group of stimuli, then the neutral primes would not be expected to produce any priming effects at all. We cannot logically rule out the possibility that participants would ever employ this alternative sort of response-mapping process, and there could well be cases in which this sort of conceptualization could account for the obtained pattern of results. However, this alternative obviously cannot account for the pattern of results obtained in our own research, given the strong priming effects we observed for the neutral primes.

A Consideration of "Contrast" in Different Types of Implicit Priming Paradigms

Earlier in this article we briefly alluded to two other priming paradigms in which researchers have used the term "contrast," including research and theory on double priming paradigms and also reverse priming effects. Although these paradigms may seem to have some connections to our own work, the overlap between paradigms is more illusory than real. In the space below we discuss each of these other paradigms in detail.

In the implicit literature, reverse priming effects are one type of effect that are commonly referred to as "contrasts" (e.g., Banse, 2001; Glaser & Banaji, 1999; Klauer, Robnagel, & Musch, 1997; Wentura, 1999; for an excellent review, see Klauer et al., 2009). Under normal circumstances, positive primes facilitate responses to positive targets, whereas negative primes facilitate responses to negative targets. A reverse priming effect is a circumstance in which a prime shows a pattern of facilitation that is the opposite of what one would normally expect. That is, a positive prime would facilitate responses to negative targets, whereas negative primes would facilitate responses to positive targets (e.g., Glaser & Banaji, 1999; Klauer et al., 1997; Spruyt, Hermans, De Houwer, Vandromme, & Eelen, 2007). Reverse priming effects are known to occur at intermediate stimulus onset asynchronies (SOAs; Klauer et al., 1997), in pronunciations tasks (Glaser & Banaji, 1999), when primes are masked immediately after presentation (Banse, 2001), and also in high-anxiety participant groups (Maier, Berner, & Pekrun, 2003). At present, there are a number of explanations as to why these sorts of reverse priming effects occur (Eimer & Schlaghecken, 1998; Klauer et al., 2009), and the literature has yet to gain clarity on this issue. Whatever the explanation for reverse priming effects, however, they are distinct from the effects reported in this article, insofar as reverse priming effects represent a circumstance in which primes facilitate a response that is the opposite of what is normally expected. This is different from the focus of our research, which was concerned with the mapping of responses onto neutral stimuli with no a priori associations.

Implicit contrast effects have also been observed in double priming procedures (e.g., Deutsch & Gawronski, 2009; Gawronski et al., 2005). In that paradigm, two primes are presented in close succession and are followed by a target. The typical finding in this paradigm is that the second prime produces a large priming effect, but only when the first prime is of the opposite valence.[3] For example, a negative prime that is preceded by a positive prime will produce a large negative priming effect, but this effect disappears when the negative prime is preceded by another negative prime. Double priming effects were originally explained via an attentional mechanism (Deutsch & Gawronski, 2009). However, recent research suggests that double priming effects are actually more similar to the sorts of reverse priming effects that occur at intermediate (approximately 500 ms) SOAs (Klauer et al., 2009). For example, when the two primes are evaluatively congruent, the second prime possesses no new evaluative

[3] In their application of the double priming paradigm to the AMP, Deutsch and Gawronski (2009) found evidence for additive effects of the two primes. This might suggest that the AMP could have some properties that differ from sequential priming tasks. However, the data from this article did not address why this difference would have occurred, leading these researchers to conclude that it is an "open question which particular features of the two measures were responsible for the obtained dissociation" (Deutsch & Gawronski, 2009, p. 11). At any rate, if there are any differences between sequential priming paradigms and the AMP, the results reported by Deutsch and Gawronski are circumscribed within the context of *double priming paradigms*, and one should not necessarily assume that those differences would generalize to other experimental paradigms. Indeed, as the results of Experiments 1 and 2 clearly show, the emergence of contrast in our (single priming) paradigm was very similar across the two types of priming tasks.

information. As a result, the "psychological SOA" for a congruent trial is twice as long as that for an incongruent trial (for a discussion of these and related issues, see Klauer et al., 2009). Hence, the "contrast effect" that is observed in double priming procedures is probably best understood as another instance of reverse priming which, for reasons noted above, is concerned with issues rather different from those under concern in our own research.

It is also worth noting that reverse priming and response mapping are not necessarily incompatible with each other. That is, response mapping should, according to our current conceptualization of the model, occur in addition to and independently of reverse priming. For example, neutral primes that are placed in a negative task context should elicit positive bias at short SOAs (as we demonstrated with our own research), but the reverse effect should be obtained for the neutral primes at intermediate SOAs. Of course, future research will be needed to confirm these predictions, and we believe that it will be important to explore these sorts of reverse priming effects within the context of our own contrast paradigm.

Unresolved Issues and Directions for Future Research

Although our data offer considerable leverage on some important issues, there are several important issues that merit further investigation. First and most important, we do not in any way maintain that spreading activation processes are a fiction and that response-based processes can explain all manifestations of priming effects. Indeed, one of our own studies in this article makes this point nicely. Recall that in Experiment 4, the priming effects elicited by the object primes coexisted along with semantic priming processes involving (for example) the tendency for flower-related primes to facilitate responses to flower-related targets. Moreover, previous research in domains other than those under concern here (e.g., lexical decision task, pronunciation tasks) has provided strong and compelling evidence that the priming effects demonstrated in those paradigms are due to spreading activation processes (Balota, Paul, & Spieler, 1999; for a review, see Klauer & Musch, 2003). Future research is clearly needed to delineate the types of paradigms in which response-mapping processes are likely to apply and whether such processes are (or are not) likely to occur independently of spreading activation processes.

Second, in our research, as in many other priming paradigms, participants were instructed to ignore the primes. However there are some paradigms in which participants are asked to perform a task involving the primes, such as recite the prime word aloud (Fazio, Sanbonmatsu, Powel, & Kardes, 1986). In such cases, instantiating a specific goal with regard to the primes may, in turn, make it less likely that participants impose the target-related goal onto the primes. Hence, another issue that merits further investigation is whether response mapping applies to situations in which participants are not asked to ignore the primes but instead are given a dual task that involves both the targets and the primes.

Third, response mapping may, at first blush, seem difficult to reconcile with some past research showing that priming effects depend on associative strength. For example, in a series of experiments by Fazio et al. (1986), strongly valenced primes showed strong priming effects, whereas weakly valenced primes showed weak effects. Response mapping might seem to have difficulty accounting for such effects, because it assumes that the weak primes will be mapped onto a response along with the strong primes, or alternatively, that the weak primes will not be mapped at all.

In our view, there are at least two plausible explanations that may resolve this apparent inconsistency. First, in Fazio et al.'s (1986) experiments, participants were asked to recite the prime aloud. As we discussed previously, response mapping may be less likely to influence responding when participants have a distinct task goal with regard to the primes. Second, remember that one of the central tenets of response mapping is that participants impose the target classification scheme onto the primes. In other words, response mapping relies on participants' ability to take a broad, response-based classification scheme and apply it to the prime stimuli. Hence, it is possible that response-mapping processes are limited to situations in which the primes are easily sorted into a small number of coherent categories. For example, in our experiments the primes always fell into two very distinct and coherent categories (e.g., flowers and objects). By contrast, in the Fazio et al. experiments the primes included a broad array of unrelated words (e.g., *aquarium*, *cake*, and *recession*) that did not form any coherent categories. As a result, it is possible that participants in the Fazio et al. experiments could not map the primes because they could not easily categorize them. Under these sorts of circumstances, we expect that other processes (such as spreading activation) may play a more dominant role in producing priming effects. Clearly there are quite a few complex issues herein, and future research will be needed to clarify them.

The final issue raised by our research concerns the fact that the context effects obtained in our paradigm emerged in the form of contrast rather than assimilation. In our paradigm, assimilation effects would reflect the tendency for neutral primes to take on properties similar to, as opposed to the opposite of, the contextual primes. We have already begun to focus on this issue in our own recent work in our laboratory. Consistent with previous models of assimilation and contrast (Schwarz & Bless, 1992), assimilation seems to be more likely to occur when the experiment is able to blur the categorical distinction between the "target" and "context" stimuli. These and other interesting issues raised by our research clearly merit further empirical attention. Such work is likely to yield valuable insight into the various ways in which contextual factors play an important role in human judgment.

References

Balota, D. A., Paul, S., & Spieler, D. (1999). Attentional control of lexical processing pathways during word recognition and reading. In S. Garrod & M. Pickering (Eds.), *Language processing* (pp. 15–57). Hove, United Kingdom: Psychology Press.

Banse, R. (2001). Affective priming with liked and disliked persons: Prime visibility determines congruency and incongruency effects. *Cognition & Emotion, 15,* 501–520.

Bargh, J. A. (1994). The four horsemen of automaticity: Awareness, intention, efficiency and control in social cognition. In R. S. Wyer & T. K. Srull (Eds.), *Handbook of social cognition* (pp. 1–40). Hillsdale, NJ: Erlbaum.

Bargh, J. A., & Chartrand, T. L. (1999). The unbearable automaticity of being. *American Psychologist, 54,* 462–479.

Biernat, M. (2005). *Standards and expectancies: Contrast and assimilation in judgments of self and others.* New York: Routledge.

Biernat, M., & Manis, M. (1994). Shifting standards and stereotype-based judgments. *Journal of Personality and Social Psychology, 66,* 5–20.

Blair, I. V., Ma, J. E., & Lenton, A. P. (2001). Imagining stereotype away: The moderation of implicit stereotypes through mental imagery. *Journal of Personality and Social Psychology, 81,* 828–841.

Brickman, P., Coates, D., & Janoff-Bulman, R. (1978). Lottery winners and accident victims: Is happiness relative? *Journal of Personality and Social Psychology, 36,* 917–927.

Dasgupta, N., & Greenwald, A. G. (2001). On the malleability of automatic attitudes: Combating automatic prejudice with images of admired and disliked individuals. *Journal of Personality and Social Psychology, 81,* 800–814.

De Houwer, J. (2006). What are implicit measures and why are we using them? In R. W. Wiers & A. W. Stacy (Eds.), *The handbook of implicit cognition and addiction* (pp. 11–28). Thousand Oaks, CA: Sage.

Deutsch, R., & Gawronski, B. (2009). When the method makes a difference: Antagonistic effects on "automatic evaluations" as a function of task characteristics of the measure. *Journal of Experimental Social Psychology, 45,* 101–114.

Devine, P. (2001). Implicit prejudice and stereotyping: How automatic are they? Introduction to the Special Section. *Journal of Personality and Social Psychology, 81,* 757–759.

Eimer, M., & Schlaghecken, F. (1998). Effects of masked stimuli on motor activation: Behavioral and electrophysiological evidence. *Journal of Experimental Psychology: Human Perception and Performance, 24,* 1737–1747.

Eiser, J. R. (1986). *Social psychology: Attitudes, cognition, and social behavior.* Cambridge, United Kingdom: Cambridge University Press.

Eiser, J. R. (1990). *Social judgment.* Pacific Grove, CA: Brooks/Cole.

Fazio, R. H. (2001). On the automatic activation of associated evaluations: An overview. *Cognition & Emotion, 15,* 115–141.

Fazio, R. H., Jackson, J. R., Dunton, B. C., & Williams, C. J. (1995). Variability in automatic activation as an unobtrusive measure of racial attitudes: A bona fide pipeline? *Journal of Personality and Social Psychology, 69,* 1013–1027.

Fazio, R. H., Sanbonmatsu, D. M., Powel, M. C., & Kardes, F. R. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology, 50,* 229–238.

Ferguson, M. J., & Bargh, J. A. (2003). The constructive nature of automatic evaluation. In J. Musch & K. C. Klauer (Eds.), *The psychology of evaluation: Affective processes in cognition and emotion* (pp. 169–188). Hillsdale, NJ: Erlbaum.

Gawronski, B., Deutsch, R., & Seidel, O. (2005). Contextual influences on implicit evaluation: A test of additive versus contrastive effects of evaluative context stimuli in affective priming. *Personality and Social Psychology Bulletin, 31,* 1226–1236.

Glaser, J., & Banaji, M. R. (1999). When fair is foul and foul is fair: Reverse priming in automatic evaluation. *Journal of Personality and Social Psychology, 77,* 669–687.

Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. (1998). Measuring individual differences in implicit cognition: The Implicit Association Test. *Journal of Personality and Social Psychology, 74,* 1464–1480.

Helson, H. (1947). Adaptation-level as frame of reference for prediction of psychophysical data. *American Journal of Psychology, 60,* 1–29.

Karpinski, A., & Steinman, R. B. (2006). The single category Implicit Association Test as a measure of implicit social cognition. *Journal of Personality and Social Psychology, 91,* 16–32.

Klauer, K. C., & Musch, J. (2003). Affective priming: Findings and theories. In J. Musch & K. C. Klauer (Eds.), *The psychology of evaluation* (pp. 7–36). Mahwah, NJ: Erlbaum.

Klauer, K. C., Robnagel, C., & Musch, J. (1997). List-context effects in evaluative priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23,* 246–255.

Klauer, K. C., Teige-Mocigembad, S., & Spruyt, S. (2009). Contrast effects in spontaneous evaluations: A psychophysical account. *Journal of Personality and Social Psychology, 96,* 265–287.

Lang, P. J., Bradley, M. M., & Cuthbert, B. N. (2008). *International affective picture system (IAPS): Affective ratings of pictures and instruction manual* (Tech. Report A-8). Gainesville: University of Florida.

Larsen, R. J., Mercer, K. A., & Balota, D. A. (2006). Lexical characteristics of words used in emotional Stroop experiments. *Emotion, 6,* 62–72.

Livingston, R. W., & Brewer, M. B. (2002). What are we really priming? Cue-based versus category-based processing of facial stimuli. *Journal of Personality and Social Psychology, 82,* 5–18.

Maier, M. A., Berner, M. P., & Pekrun, R. (2003). Directionality of affective priming: Effects of trait anxiety and activation level. *Experimental Psychology, 50,* 116–123.

Neely, J. H. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. In D. Besner & G. W. Humphreys (Eds.), *Basic processes in reading: Visual word recognition* (pp. 264–336). Hillsdale, NJ: Erlbaum.

Ostrom, T. M. (1970). Perspective as determinant of attitude change. *Journal of Experimental Social Psychology, 6,* 280–292.

Parducci, A. (1965). Category judgment: A range-frequency model. *Psychological Review, 72,* 407–418.

Payne, B. K., Cheng, C. M., Govorum, O., & Stewart, B. D. (2005). An inkblot for attitudes: Affect misattribution as implicit measurement. *Journal of Personality and Social Psychology, 89,* 277–293.

Payne, B. K., Jacoby, L. L., & Lambert, A. L. (2005). Attitudes as accessibility bias: Dissociating automatic and controlled components. In R. Hassin, J. Bargh, & J. Uleman (Eds.), *The new unconscious* (pp. 393–420). Oxford, United Kingdom: Oxford University Press.

Petty, R. E., & Cacioppo, J. T. (1981). *Attitudes and persuasion: Classic and contemporary approaches.* Dubuque, IA: Brown.

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

Schwarz, N., Groves, R. M., & Schuman, H. (1998). Survey methods. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *The handbook of social psychology* (pp. 143–179). New York: McGraw-Hill.

Sherif, M., & Hovland, C. I. (1961). *Social judgment: Assimilation and contrast effects in communication and attitude change.* New Haven, CT: Yale University Press.

Sherman, S. J., Ahlm, K., Berman, L., & Lynn, S. (1978). Contrast effects and their relationship to subsequent behavior. *Journal of Experimental Social Psychology, 14,* 340–350.

Simpson, D. D., & Ostrom, T. M. (1976). Contrast effects in impression formation. *Journal of Personality and Social Psychology, 34,* 625–629.

Spruyt, A., Hermans, D., De Houwer, J., Vandromme, J., & Eelen, P. (2007). On the nature of the affective priming effect: Effects of stimulus onset asynchrony and congruency proportion in naming and evaluative categorization. *Memory & Cognition, 35,* 95–106.

Steele, C. M., & Ostrom, T. M. (1974). Perspective-mediated attitude change: When is indirect persuasion more effective than direct persuasion? *Journal of Personality and Social Psychology, 29,* 737–741.

Upshaw, H. S. (1969). Stimulus range and the judgmental unit. *Journal of Experimental Social Psychology, 5,* 1–11.

Volkmann, J. (1951). Scales of judgment and their implications for social psychology. In J. H. Rohrer & M. Sherif (Eds.), *Social psychology at the crossroads* (pp. 273–294). New York: Harper.

Wedell, D. H., Hicklin, S. K., & Smarandescu, L. O. (2007). Contrasting models of assimilation and contrast. In D. A. Stapel & J. Suls (Eds.), *Assimilation and contrast in social psychology* (pp. 45–74). New York: Psychology Press.

Wentura, D. (1999). Activation and inhibition of affective information: Evidence for negative priming in the evaluation task. *Cognition & Emotion, 13,* 65–91.

Received October 15, 2008
Revision received February 23, 2009
Accepted February 24, 2009 ■