Journal of Personality and Social Psychology
2006, Vol. 90, No. 2, 210–220

0022-3514/06/$12.00 DOI: 10.1037/0022-3514.90.2.210

Being Bad Isn't Always Good: Affective Context Moderates the Attention Bias Toward Negative Information

N. Kyle Smith
Ohio Wesleyan University

Jeff T. Larsen
Texas Tech University

Tanya L. Chartrand
Duke University

John T. Cacioppo
University of Chicago

Heather A. Katafiasz and Kathleen E. Moran
Ohio Wesleyan University

Research has demonstrated that people automatically devote more attention to negative information than to positive information. The authors conducted 3 experiments to test whether this bias is attenuated by a person's affective context. Specifically, the authors primed participants with positive and negative information using traditional (e.g., subliminal semantic priming) and nontraditional (e.g., social interactions) means and measured the amount of attention they allocated to positive and negative information. With both event-related brain potentials (Experiment 1) and the Stroop task (Experiments 2 and 3), results suggest that the attention bias to negative information is attenuated or eliminated when positive constructs are made accessible. The implications of this result for other biases to negative information and for the self-reinforcing nature of emotional disorders are discussed.

Keywords: attention bias, event-related brain potentials, positive priming

Researchers have documented many information processing biases toward negative information. For example, people give more weight to negative traits than positive traits in impression formation tasks (Anderson, 1965; Peeters & Czapinski, 1990), they dislike losses more than they like equally large gains (Kahneman & Tversky, 1984), and they make more causal attributions for negative events than for positive events (Peeters & Czapinski, 1990). This differential emphasis on negative stimuli manifests itself not only in self-report investigations but in those measuring behavior (Spence & Segner, 1967) as well as brain activity in response to positive and negative stimuli (Ito, Larsen, Smith, & Cacioppo, 1998). In reviewing the literature on the relationship between the valence of a stimulus and the magnitude of the responses generated, Baumeister, Bratslavsky, Finkenauer, and Vohs (2001) stated, "We have found bad to be stronger than good in a disappointingly relentless pattern . . . this difference may be one of the most basic and far-reaching psychological principles" (p. 362).

N. Kyle Smith, Heather A. Katafiasz, and Kathleen E. Moran, Department of Psychology, Ohio Wesleyan University; Jeff T. Larsen, Department of Psychology, Texas Tech University; Tanya L. Chartrand, Department of Psychology, Duke University; John T. Cacioppo, Department of Psychology, University of Chicago.

We thank Todd Amdor, Laura Browne, and Phing Dong for their help during data collection.

Correspondence concerning this article should be addressed to N. Kyle Smith, Department of Psychology, Ohio Wesleyan University, Delaware, OH 43015. E-mail: nksmith@owu.edu

All of the just-mentioned research demonstrates that there is an extremity bias in information processing. Cacioppo, Gardner, and Berntson (1997) have suggested that this is a basic operating parameter of the evaluative system. This bias arises, they stated, because for each unit of input to the positive and negative evaluative systems, the negative evaluation system responds with a larger output. Therefore, regardless of whether you are learning information about a new person, receiving feedback on an academic pursuit, or deciding whether to bet on a sporting event, negative information will tend to play a larger role in your decision than one might normatively expect.

In addition to these evaluative and elaborative biases, negative stimuli have also been shown to elicit more attention than positive stimuli. That is, apart from the extremity bias, there is also an *attention bias.* Early work in this domain relied on measures such as the amount of time voluntarily allocated to processing positive and negative information (e.g., Fiske, 1980; Graziano, Brothen, & Berscheid, 1980). For example, Graziano et al. (1980) found that, when participants were given the option to hear either positive or negative feedback about themselves, they chose to listen to the negative feedback for a significantly longer amount of time.

As interest in automatic processes has increased and methodologies to study automatic processes have developed, research surrounding the attention bias has shifted to assessing the extent to which negative stimuli draw attention automatically. For example, Hansen and Hansen (1988; see also Öhman, Lundqvist, & Esteves, 2001) gave participants the task of locating an evaluatively discrepant face in an array of faces. Consistent with the attention bias, participants were faster at picking out angry faces embedded within a grid of happy faces than they were at picking out happy faces embedded within a grid of angry faces. Further, as the size of the grid increased, the time to locate a discrepant happy face increased, whereas the time to locate a discrepant angry face remained constant. This suggests that participants' searches for the angry faces proceeded efficiently and in parallel, whereas searches for happy faces were performed serially.

A question that remained unanswered, however, concerned whether the attention bias was automatic when people were not given an evaluative task. That is, because Hansen and Hansen (1988) asked people to locate angry or happy faces, it is possible that the attention bias was goal-dependently automatic. Whether participants would continue to show the attention bias when not given explicit instructions to evaluate stimuli was still unclear.

To address this question, Pratto and John (1991) used the emotional Stroop task to assess the attention demands of positive and negative stimuli. In this paradigm, participants are shown positive and negative words written in different colors and asked to name the color in which the word was written as quickly as possible. Differences in color-naming latencies for positive and negative words index the extent to which participants' attention is being automatically drawn away from the color-naming task and focused on the word being presented. Because participants are presumably focused on the color-naming tasks and ignoring the evaluative nature of the words (a supposition supported by participants' self-reports solicited during the debriefing), they are not in an evaluative set, and any attention differences observed in this paradigm should, therefore, be deemed preconsciously automatic.

Pratto and John's (1991) results showed that negative words had longer color-naming latencies than positive words, suggesting that negative words were automatically drawing more attention than positive words. From these data, Pratto and John (1991) concluded that people are automatically vigilant for negative information in their surroundings.

One interpretational difficulty associated with Hansen and Hansen's (1988) and Pratto and John's (1991) work is that both relied on reaction time measures as a way of tapping attention allocation. Given that many mental processes take place between the attention allocation phase of processing and the overt response that these two studies measured, it is possible that the observed differences were due not to attention at all but rather to systematic differences in one or more downstream processes. For example, Taylor (1991) suggested that processing a negative stimulus tends to result in, among other things, an initial mobilization of resources to respond to that stimulus. If our response tendencies are stronger for negative information than positive information, it stands to reason that in tasks in which responding to negative stimuli facilitates our task, we would respond more quickly to negative information (as in Hansen & Hansen, 1988, and Öhman et al., 2001), whereas in tasks in which responding to negative stimuli interferes with our task, we would respond more slowly to negative information (as in Pratto & John, 1991).

To rule out this alternative explanation, Smith, Cacioppo, Larsen, and Chartrand (2003) used event-related brain potentials (ERPs) to measure the attention allocation stage of information processing. ERPs are measures of brain activity that are time locked to a stimulus event and, therefore, index the information processing steps taken as a result of the stimulus. ERPs can be conceptually broken down into a smaller number of components, each thought to represent a specific information processing function. One of these components, the P1 (so named because of its positive polarity and its mean latency of approximately 100-ms poststimulus onset), indexes the number of neurons allocated to process a given stimulus in the extrastriate visual areas (V. P. Clark & Hillyard, 1996). That is, larger P1s indicate that more neurons are being recruited to process a visual stimulus, suggesting that that the stimulus is receiving more attention. Because of its extremely early latency and its location in the visual areas of the brain, the P1 is relatively insensitive to downstream processing. Therefore, Smith et al. (2003) reasoned, the P1 would be a good measure to test the hypothesis that negative stimuli receive more attention than positive stimuli.

In two experiments, Smith et al. (2003) showed participants positive, negative, and neutral pictures. In Experiment 1, participants viewed positive and negative pictures by block. In the first (counterbalanced) block, participants saw a positive picture for 11 s. Embedded within this positive picture 5, 7, or 9 s later was a smaller, target picture that was either positive (50% of the time) or negative (50% of the time). The valence of the background picture was switched in a second block. ERPs were recorded to the embedded positive and negative pictures. In Experiment 2, participants viewed predominantly neutral stimuli with occasional positive and negative stimuli interspersed. ERPs were again recorded to the positive and negative pictures. In both studies, negative pictures elicited significantly larger P1s than did positive pictures, suggesting that negative stimuli received more attention than positive stimuli at this very early stage of information processing.

To summarize, an obligatory attention bias to negative information has been demonstrated using diverse methods. This bias seems to be preconsciously automatic and has been shown to modulate processing within 100 ms of exposure to a negative stimulus.

Various theoretical formulations have attempted to explain why biases toward negative information exist. Some have suggested that negative information differs along many dimensions from positive information (e.g., the extent to which it violates our usually positive expectations, the diagnostic value of negative behaviors) and that one of these differences that correlates with valence is responsible for the enhanced value given to negative information (Skowronski & Carlston, 1989). Others have suggested that selection pressures during evolutionary times were such that organisms that weighted negative information more heavily than positive information were more successful and, therefore, prospered (e.g., Cacioppo et al., 1997; Hansen & Hansen, 1988; Peeters & Czapinski, 1990; Pratto & John, 1991). Regardless of the rationale, it seems clear that, whether because of incidental features correlated with valence or valence itself, negative information plays a larger role in our information processing and behavior than does positive information.

This is not to say that negative stimuli are always given an advantage over positive stimuli. For example, the hedonic contingency model (Wegener & Petty, 1994) has suggested that when in positive moods people preferentially process positive information in an effort to maintain that positive mood. Other researchers (e.g., Walker, Vogl, & Thompson, 1997) have suggested that the unpleasantness of negative memories decreases at a greater rate than the pleasantness of positive memories. Taylor (1991) has suggested that although negative events have larger initial reactions, these reactions are often dampened more quickly and severely than reactions to positive or neutral events. However, researchers (e.g., Baumeister et al., 2001) have argued that these are mere exceptions to the general rule that negative information dominates positive information at most tasks.

The Moderating Effect of Context

As the amount of evidence supporting negativity biases has mounted, and as theoretical accounts tying the biases to evolutionary pressures have accumulated, researchers have become more willing to suggest that the biases to negative information may be unconditional. In the strongest statement to effect, Baumeister et al. (2001) wrote, "Given the large number of patterns in which bad outweighs good, however, any reversals are likely to remain as mere exceptions" (p. 362). However, there are several compelling arguments to suggest that this is perhaps an oversimplified view.

It is tempting to think that whatever is evolutionary is immutable. However, just as evolution favors particular static, anatomical structures (e.g., opposable thumbs) and behaviors (e.g., the cough reflex) that enhance fitness, it is also true that evolution favors flexibility. For example, pupils that can expand and contract based on the amount of light in the environment are much more useful than pupils that are fixed at a moderate level of openness. Applying the idea of valuing flexibility to the attention bias, it is certainly true that there are environments in which it is quite beneficial to preferentially attend to negative information. For example, in an environment in which the consequences of negative events are more dire than the consequences of positive events are beneficial (e.g., being killed by a predator vs. capturing and eating a prey animal), it can be adaptive to overattend to negative information. However, in environments where there are no dangers, or where the threat posed by the dangers is in line with or less than the potential benefits produced by the boons (e.g., a mating situation in which successfully courting a female results in passing on genes, whereas being attacked by a male rival results in a non-life-threatening injury), it is of little benefit (and may even be maladaptive) to overattend to negative information.

Given that (a) there are environments in which displaying the attention bias is adaptive and other environments in which it is of little benefit and (b) situational factors (e.g., the ratio of positive to negative forces in the environment and the relative severity of the potential positive and negative outcomes) can indicate the value of the attention bias in a particular environment, it stands to reason that it would be adaptive for organisms to modulate the size of the bias based on the demands of the particular environment. To investigate this phenomenon, therefore, it would be beneficial to find a construct that is sensitive to the level of positive and negative stimuli in the environment and able to modulate a person's attention.

One low-effort way in which organisms could keep track of the status of danger in their environment is by using the accessibility of positive and negative constructs in memory. That is, negative constructs predominating in memory could serve as an indicator that danger is possible and, therefore, guide our attention to negative information. Alternatively, positive constructs predominating could indicate that there is little risk of danger and that negative information need not be weighted or attended to as heavily. This theorizing is consistent with a long tradition of work suggesting that increasing the accessibility of a stimulus or a category of stimuli (e.g., by means of a priming procedure) makes people preferentially attend to the more accessible stimuli (e.g., Neely, 1977). Therefore, in Experiment 1, we examined whether manipulating the probability of positive and negative stimuli, and thereby manipulating the accessibility of positive and negative information, could modulate the attention bias to negative information.

Experiment 1

It is well known in social psychology that accessible attitudes can guide attention toward the relevant attitude objects (e.g., Roskos-Ewoldsen & Fazio, 1992). There is also evidence that increasing the accessibility of one member of a valence category (e.g., a positive word) facilitates responses to other semantically related members of that valence category that are semantically unrelated to the prime word (Bargh, Chaiken, Govender, & Pratto, 1992). Taking these two results together, it is possible that increasing the accessibility of specific positive constructs in memory should increase the attention paid to other positive stimuli in the environment regardless of their semantic relationship to the primed constructs. Therefore, we hypothesize that the attention bias to negative information will be attenuated by increasing the accessibility of positive constructs in memory.[1]

We exposed participants to blocks of pictures in which one valence category (positive or negative) was primed. To manipulate the categorical accessibility of the stimuli, within each block stimuli of one valence category (e.g., negative) were presented frequently, whereas stimuli of the other valence category (e.g., positive) were rare. Following up the ERP work of Smith et al. (2003), we hypothesized that when negative pictures were presented frequently, P1 amplitude would be larger to negative pictures than positive pictures, indicating that negative pictures received more attention. However, when positive pictures were primed, we hypothesized that the attention bias would be attenuated or reversed.

Method

Participants. Twenty-seven Ohio State University introductory psychology students participated in this study for partial course credit. Of the 27 students, 1 did not complete the procedure, and the data from another were unusable because of excessive artifact. Therefore, statistical analyses were performed on data from the remaining 25 participants.

Materials. Thirty normatively positive and 30 normatively negative pictures were selected from the International Affective Picture System (Lang, Bradley, & Cuthbert, 1995) for use in the study. Three pictures were selected from each group to serve as target stimuli; the remaining 27 served

[1] Repeated exposure to a stimulus can often have contradictory effects on attention. Just as one can predict that increased exposure will lead to increased accessibility and, therefore, increased attention, one can also predict that increased exposure will lead to habituation or boredom and, therefore, decreased attention. Johnston et al. (e.g., Johnston & Hawley, 1994) have suggested that both of these do occur but at different points during the information processing stream. They suggest that initially more attention is given to primed stimuli and that in subsequent processing these stimuli are given less attention, and novel stimuli are instead processed. Given that our work is focusing on the initial, obligatory allocation of attention, it is most likely that we will only see the increased attention associated with priming as opposed to the subsequent decrease in attention.

as context stimuli. All four groups of pictures were equated on evaluative extremity and self-reported arousal using norms collected by Ito, Cacioppo, and Lang (1998).

Procedures. On arrival at the lab, research participants were given an overview of the procedure. Participants then had electrodes attached to their heads for recording ERPs. Participants were seated in a sound-attenuating, electrically shielded room and given instructions for the task. They were told that they would be watching pictures on a computer monitor, and their task would be to indicate, by button press, whether they thought each picture was positive or negative.

Stimuli were presented to the participants in two counterbalanced blocks. Across both blocks, pictures were presented in sequences of five, with each picture having a 1-s presentation time and a 1-s interstimulus interval. Within each block, participants were exposed to either predominantly negative pictures with rare positive pictures or predominantly positive pictures with rare negative pictures.[2] ERPs were recorded during each of the rare (unprimed) stimulus presentation and an equal number of presentations of the frequent (primed) valence category (Table 1). Each of the sequence types in Table 1 was repeated 10 times so that 30 ERP waveforms would be available for each experimental condition. The hand with which participants indicated a stimulus was positive was counterbalanced.

Psychophysiological data collection and cleaning. ERP data were recorded from 28 tin electrodes mounted in an elastic cap (ElectroCap International, Eaton, OH). The locations of the electrodes were based on an expanded version of the international 10–20 electrode placement system. Only four of these sites, the midline sites Fz, Cz, Pz, and Oz, were analyzed for the current report. Each of these sites was referenced online to the left mastoid; the right mastoid was recorded for later offline rereferencing. Additionally, two tin cup electrodes were placed above and below participants' left eye to record vertical eye movements and eyeblinks. Interelectrode impedances were all less than 10 kΩ.

Electrical activity from the scalp and eyes was amplified by a NeuroScan (El Paso, TX) SynAmps amplifier, which applied a bandpass filter between 0.1 Hz and 30 Hz (12-dB roll-off). The data were then digitized by a computer at 1000 Hz and stored on the hard drive for later analysis. Each recording epoch began 128 ms before the target picture appeared on the screen and continued for the entire duration of the presentation.

To eliminate extraneous noise, several offline signal processing techniques were performed. First, the data were rereferenced to a mathematically simulated linked ears reference. Next, to remove baseline differences, the average amplitude of each electrode for each trial over the 128-ms baseline period was set to zero. Further, to remove artifacts caused by eyeblinks, variance correlated with vertical electro-oculographic activity was removed from the electroencephalographic signal (Semlitsch, Anderer, Schuster, & Presslich, 1986). Subsequently, trials were individually inspected and excluded if nonneurogenic artifact was present.

Table 1
Stimulus Presentation Orders for Experiment 1

Trial type	Stimulus position 1	2	3	4	5
1: positive block	Pos	Pos	**Pos**	Pos	Pos
2: positive block	Pos	Pos	Pos	**Pos**	Pos
3: positive block	Pos	Pos	Pos	Pos	**Pos**
4: positive block	Pos	Pos	**Neg**	Pos	Pos
5: positive block	Pos	Pos	Pos	**Neg**	Pos
6: positive block	Pos	Pos	Pos	Pos	**Neg**
1: negative block	Neg	Neg	**Neg**	Neg	Neg
2: negative block	Neg	Neg	Neg	**Neg**	Neg
3: negative block	Neg	Neg	Neg	Neg	**Neg**
4: negative block	Neg	Neg	**Pos**	Neg	Neg
5: negative block	Neg	Neg	Neg	**Pos**	Neg
6: negative block	Neg	Neg	Neg	Neg	**Pos**

Note. Bold typeface indicates stimuli during which event-related brain potentials were recorded. Pos = positive stimulus presentation; Neg = negative stimulus presentation.

Data reduction and analysis. After deleting trials with artifacts, each participant's trials were aggregated based on trial type. That is, for each participant, one average ERP waveform was constructed for each cell of the 2 (stimulus valence: positive or negative) × 2 (prime valence: positive or negative) design.

To quantify the P1, a principal components analysis (PCA) was performed on the data. PCA partitions the total variance of the ERP into a small number of underlying components by analyzing the patterns of covariation between time points in the average waveforms. As recommended by Donchin and Heffley (1978) and Van Boxtel (1998), we calculated the PCA on a covariance matrix, extracted six components, and used a varimax rotation on the extracted components.

The output of the PCA consists of two matrices. One, the component loading matrix, shows the strength of the influence each component has at each time point. Just as in performing a PCA on a survey one might find that some portion of the variability in responses to Questions 1, 3, and 7 on a questionnaire is attributable to an underlying component (that, on further thought an investigator might label *introversion*), the loading matrix of a PCA says that the electrical activity at some set of time points is attributable to an underlying component. For example, the loading matrix (see Figure 1) shows that Component 4 has an influence over time points in the 600-ms range, with its impact decreasing the further away from 600 ms one looks. To interpret the psychological relevance of a given component, researchers assess the timing, polarity, and scalp distribution and compare these with previously established components to infer its function.

The second matrix, the component score matrix, assesses the extent to which each component is present in each average waveform entered into the PCA. This is analogous to, in the previous survey example, a set of scores that indicate the extent to which each participant's data contain a given component (i.e., how introverted is Participant 5). Because, in this PCA, these scores indicate how much of an ERP component was present in a given participant's condition average, analyzing them as a dependent variable will allow an inference to be made about the extent to which the psychological process that the component represents is occurring in participants in that condition. Therefore, our analysis strategy was to find a component in the time window of the P1, check its scalp distribution and polarity to see whether it matched the known parameters of the P1 (i.e., that it has an occipital maximum and is positive going at occipital sites), and, if so, analyze the component scores for that component to see in which experimental conditions the P1 was maximal.

Applying this strategy, we found a single component matching the latency (active between 100 and 150 ms poststimulus onset), scalp distribution (maximal over the occipital lobe), and polarity (positive going over the occipital lobe) of the P1. Therefore, we analyzed the extent to which the

[2] It is important to note that, although, within each block, one valence category was more frequent than another, the frequency of each picture was held constant regardless of valence. That is, although positive stimuli were presented 270 times and negative stimuli 30 times in the positive block, each individual stimulus was presented exactly 10 times. Therefore, the priming effects that we are hypothesizing are a differential priming not simply because of different numbers of exposures to a given picture but rather because of different numbers of exposures to a given valence category.

presence of this component was modulated by the experimental conditions. See Figure 1 for a graph of component loading versus time.[3]

Results

To examine whether the priming manipulation affected the amount of attention drawn by the stimuli, the P1's component scores, which index the extent to which the P1 is in each average waveform, were analyzed in a 2 (stimulus valence: positive or negative) × 2 (prime valence: positive or negative) general lineal model (GLM).[4] There was a significant main effect of prime valence such that stimuli in the positive prime condition elicited larger P1s ($M = 1.27$, $SE = 0.167$) than stimuli in the negative prime condition ($M = 1.10$, $SE = 0.154$), $F(1, 24) = 7.02$, $p < .05$. However, this main effect was qualified by a significant Stimulus Valence × Prime Valence interaction, $F(1, 24) = 22.5$, $p < .001$ (see Figure 2). To break down this interaction, we performed simple-effects tests comparing the P1s elicited by positive and negative stimuli within each priming condition. In the negative prime block, the data replicated the effects found in Smith et al. (2003). That is, negative pictures elicited larger P1s ($M = 1.23$, $SE = 0.153$) than did positive pictures ($M = 0.958$, $SE = 0.163$), $t(24) = 4.07$, $p < .001$. In contrast, when positive constructs were primed, this pattern reversed. That is, positive pictures elicited marginally larger P1s ($M = 1.34$, $SE = 0.174$) than negative pictures ($M = 1.194$, $SE = 0.167$), $t(24) = 2.01$, $p = .056$. The grand average ERPs at the occipital scalp site that illustrate this effect are shown in Figure 3.

In prior research, studies conducted in our laboratory (e.g., Cacioppo, Crites, Berntson, & Coles, 1993; Ito et al., 1998) have focused on the late positive potential (LPP), a component that is maximal over the parietal lobe, occurs between 400 and 700 ms poststimulus onset, and has an amplitude that is proportional to the evaluative discrepancy between the evaluative context and the current target stimulus. To demonstrate that the current work is consistent with prior work and to examine the possibility that the amplitude of the LPP and the P1 may be associated, we examined the LPP component in the current data.[5]

After picking peaks, the peak amplitudes were analyzed using a 2 (stimulus valence: positive or negative) × 2 (prime valence: positive or negative) interaction. Consistent with prior literature, and our expectation that stimuli that mismatched the prime valence would yield the largest LPPs, we found a significant Stimulus Valence × Prime Valence interaction, $F(1, 24) = 97.59$, $p < .001$.

Breaking down this interaction, in the negative prime condition, evaluatively discrepant positive targets yielded larger LPPs ($M = 12.93\ \mu V$, $SE = 1.12\ \mu V$) than evaluatively consistent negative targets ($M = 5.96\ \mu V$, $SE = 0.90\ \mu V$), $t(24) = 6.67$, $p < .001$. In the positive prime condition, evaluatively discrepant negative targets yielded larger LPPs ($M = 10.75\ \mu V$, $SE = 1.24\ \mu V$) than evaluatively consistent positive targets ($M = 5.02\ \mu V$, $SE = 0.87\ \mu V$), $t(24) = 5.62$, $p < .001$. In short, consistent with prior work on the LPP, stimuli that varied from the evaluative context caused larger LPPs than those that were consistent with it.

After establishing that the LPP results are consistent with the prior findings in the literature, we sought to address whether the LPP and P1 are related. Because the LPP is largest when the valence of the target stimulus is inconsistent with the valence of its context and the P1 is largest when the valence of the target is consistent with the valence of its context, comparing raw and LPP and P1 amplitudes would not be expected to show a meaningful relationship. Therefore, we computed several indexes of the neg-

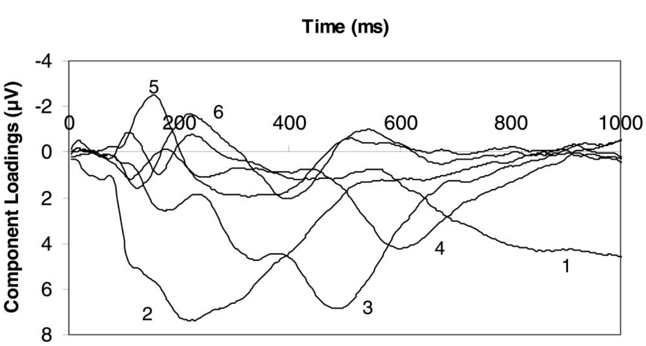

Figure 1. Component loadings as a function of time for Experiment 1. The components are numbered in order of the proportion of variance for which each accounted. The P1 is numbered Component 2.

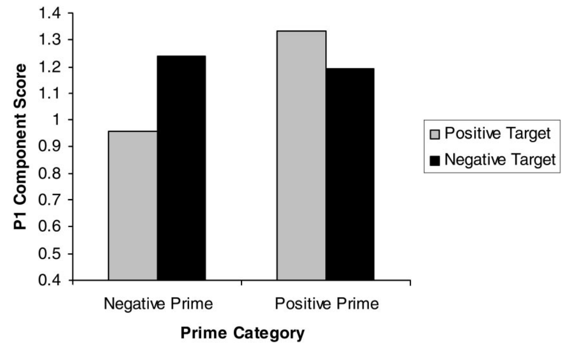

Figure 2. P1 component score as a function of stimulus valence and prime type in Experiment 1. When negative constructs were primed, negative stimuli evoked larger P1s (and therefore more attention) than positive stimuli. When positive constructs were primed, negative and positive stimuli did not differ in amplitude of P1s they elicited.

[3] The fact that the P1 component is broader than normal (lasting until 500 ms poststimulus onset) suggests that it and another component have common scalp distributions and responsiveness to the experimental conditions. Although this does affect the shape of the component extracted by the PCA, it does not qualify our interpretations of this component as representing the activity of the P1.

[4] The full design included the factors of block order and hand indicating a stimulus was positive; however, neither of these factors influenced the principal analyses, so they were excluded.

[5] Because the LPP is relatively large compared with other activity within the same time frame, it is not necessary to use PCA to analyze this component. To be consistent with prior research, therefore, we measured LPP amplitude by picking the largest peak within the 400–700 ms time window in each participant's four-stimulus valence by prime valence condition averages. Because the LPP is maximal at electrode site Pz, and that is its traditional site of LPP quantification, we analyzed the LPP at Pz. Analyses were also conducted analyzing data from the PCA using Component 4 as the LPP. The results of that analysis were consistent with those reported here.

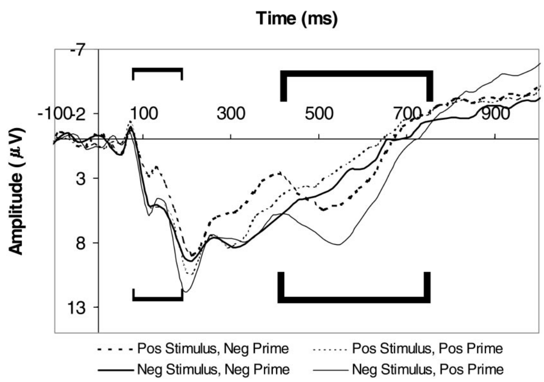

Figure 3. Averaged event-related brain potential waveforms at the midline occipital electrode to primed and unprimed positive and negative stimuli. The small square brackets highlight the time frame of the P1, and the large square brackets highlight the time frame of the late positive potential. Note that the amplitudes during the P1 time frame show the same pattern of activity as the means depicted in Figure 2.

ativity bias in each component to assess the relationship between the two.[6] None of the comparisons yielded a significant relationship between the LPP and P1 indexes (all $ps > .20$).

Discussion

In Experiment 1, we demonstrated that the attention bias to negative information (as indexed by the P1 component of the ERP) was eliminated when positive stimuli were made categorically frequent. These results show that whereas negative stimuli may generally draw more attention than positive stimuli, positive stimuli can draw more attention than negative stimuli in positive affective contexts.

One question that we can ask about these data is, how do they compare with our prior research into the attention bias? On the surface, this study seems similar to Smith et al.'s (2003) Experiment 1, in which participants saw positive or negative target pictures embedded within single positive or negative contextual pictures. In that study, regardless of the valence of the contextual (prime) picture, negative target stimuli elicited larger P1s, and hence received more attention, than positive target stimuli. Although the designs of the two studies, positive and negative target pictures preceded by positive and negative primes, seem similar on the surface, the sheer number of primes from one study to the second points to an important boundary condition of this effect. In Smith et al. (2003), two thirds of the pictures were of the prime valence and one third was of the target valence. In the current study, 90% of the pictures were of the prime valence and 10% were of the target valence. This suggests that, in order for the attention bias to be successfully attenuated, the balance of positive and negative information must be strongly skewed to the positive. Therefore, in Experiments 2 and 3, we used a strong priming manipulation before assessing the direction of participants' attention.

Several important features should be noted about this study. First, we can only conclude from these data that the attenuation of the attention bias, although automatic, is goal dependent (Bargh, 1989). Because participants were given an evaluative task in Experiment 1, it is unclear whether or not this attenuation of the bias occurs when participants are not consciously evaluating stimuli. However, based on the work suggesting that people automatically evaluate stimuli in their environment (e.g., Bargh et al., 1992), we believe that an explicit evaluative task is not necessary for this attenuation. We tested this issue in Experiment 2. Further, because participants were primed with either positive stimuli or negative stimuli in Experiment 1, it is unclear how our results relate to the attention biases documented by researchers who did not manipulate the affective context of their participants. To address this issue, in Experiment 2 we included a baseline condition in which participants were not primed. Finally, although the P1 component of the ERP has been used to index the attention bias previously, it would lend further credence to this attenuation to demonstrate it with a more traditional social psychological measure.

Using the emotional Stroop paradigm allows all of these issues to be addressed. Because the emotional Stroop task does not require an evaluative task, it is suitable for addressing the question of the automaticity of the attenuation of the attention bias. Further, using words as stimuli allows us to assess the generality of the priming attenuation of the attention bias.

Experiment 2

To manipulate the accessibility of positive and negative constructs in memory, we used parafoveal subliminal priming (Bargh, Bond, Lombardi, & Toda, 1986) with positive and negative prime words. Specifically, one group was primed with positive words and another with negative words; a third group was not primed and served as a control group. Each group then performed the emotional Stroop task.

On the basis of the findings of Experiment 1 and the results of Pratto and John (1991), we expected that both the group primed with negative words and the group that was not primed would show the attention bias toward negative information. That is, they should take longer to name the color of negative words than positive words. However, we expected that the group primed with positive words would not show this attention bias.

Method

Procedure overview. Participants were 35 introductory psychology students at Ohio State University. On arriving, participants were informed that they would be participating in a study on different types of reaction times. The positive- and negative-prime groups were told that they were going to perform a test that examined hand–eye reaction time. These two groups then completed the priming task. Next, all three groups completed the emotional Stroop task. Participants in the two priming conditions

[6] In the LPP, we can define the negativity bias as the difference in amplitude between a negative oddball and a positive oddball. An alternative operationalization would have the oddball amplitudes subtracted from the amplitudes of the pictures of their priming contexts. The difference between these baseline positive and negative oddball amplitudes could then be used as the negativity bias. For the P1, we prepared two operationalizations of the attention bias. First, we computed the difference between positive and negative oddballs in each prime context condition. Second, we computed the difference between positive and negative frequent stimuli in each prime context condition.

completed a funneled debriefing in which their perception of the subliminal primes was assessed. One participant was able to articulate a negative prime word, and his data were excluded. The remainder of the participants were unable to report any of the prime words, and most expressed surprise that words had been presented. Finally, all participants were debriefed and released.

Priming manipulation. The priming phase of the experiment was carried out on an IBM-compatible Pentium computer using DMDX software developed at Monash University and the University of Arizona by K. I. Forster and J. C. Forster. To ensure that stimuli were subliminal, they were presented between two and five degrees away from participants' visual fixation. Specifically, participants were placed 36 in. away from the screen, and stimuli were presented 1.9 in. away from the fixation point, along the diagonals of the screen such that the stimuli fell into one of the four quadrants of the computer screen. Participants were instructed to press a key as quickly as they could to indicate whether the stimulus appeared on the left or right.

The prime stimuli were presented for 68 ms and immediately covered by a mask of *X*s for another 68 ms. In both prime groups, participants viewed 20 presentations of each of four stimulus words in a random order. The prime words were taken from Chartrand and Bargh (2001; see Bargh & Chartrand, 1999). Positive primes were *sunshine, friends, Friday,* and *music,* whereas the negative primes were *cancer, crime, cockroach,* and *war.*

Emotional Stroop task. After the priming manipulation, all participants completed a computer-based emotional Stroop task. Because this is a task that measures a verbal reaction time, participants were fitted with a microphone connected to a headset. The microphone was positioned slightly above the nose to ensure that any verbalizations they made would be detected and that breathing would not cause any artifactual triggers.

During the task itself, participants were shown words written in one of five colors: yellow, purple, red, blue, and green. Each word's color was randomly determined with the provision that the same color could not appear twice in a row. The participants' task was to name the color in which the word was written as quickly as possible. Words were presented one at a time in the middle of the screen. The words stayed on the screen until a verbal response was detected by a software-based voice key.

Stimuli were presented and voice reaction times were collected with DMDX software. The 88 target words were taken from Pratto and John (1991; Study 3; see their Appendix A for a list of the stimuli).[7] To allow participants to adjust to the task, the first four trials were practice trials composed of neutral words.

Data reduction and cleaning. As in Pratto and John (1991), we excluded data from trials in which incorrect responses were given or the voice timer was triggered by a nonword noise (e.g., a cough). Reaction times less than 300 ms or greater than 1,500 ms were excluded as well. Overall, 3.4% of trials were excluded. From the cleaned data, we generated two average reaction times per participant, one for positive targets and one for negative targets. These reaction times were then analyzed in a 2 (target valence: positive or negative) $\times$ 3 (prime valence: positive, negative or no prime) mixed-model GLM; the latter factor was manipulated between groups.

Results and Discussion

To support our prediction that the attention bias would be present in the no-prime and negative prime conditions but not in the positive prime condition, we expected a significant Target Valence $\times$ Prime Valence interaction. That two-way interaction was significant, $F(2, 32) = 4.12$, $p < .05$ (see Figure 4). To break down this interaction, we chose to analyze the effects of target valence at each level of prime valence. First, replicating Pratto and John (1991), we found that in the no-prime condition negative targets resulted in longer reaction times ($M = 639$ ms, $SE = 21.7$) than positive targets ($M = 619$ ms, $SE = 19.3$), $t(12) = 2.71$, $p <$ .01. The same pattern of results was found in the negative prime condition as well (negative targets: $M = 655$ ms, $SE = 22.5$; positive targets: $M = 635$, $SE = 20.1$), $t(11) = 2.76$, $p < .05$. However, in the positive prime condition, no difference was found in the reaction times to positive versus negative targets (negative targets: $M = 656$ ms, $SE = 24.7$; positive targets: $M = 664$, $SE = 22.0$), $t(9) = 0.944$, *ns*.

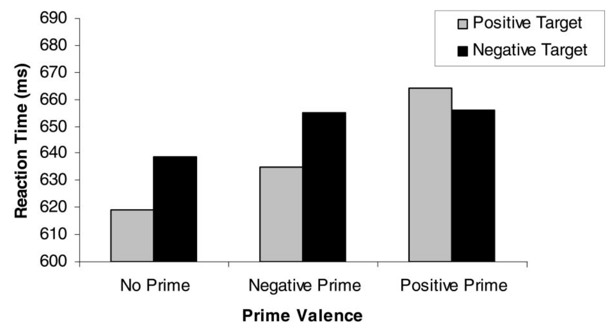

Figure 4. Reaction time as a function of word valence and prime type in Experiment 2. When no priming occurred and when negative constructs were primed, negative stimuli target words elicited longer reaction times than did positive target words. When positive constructs were primed, reaction times to positive and negative stimuli did not differ.

These results replicate and extend the findings from Experiment 1. First, they provide converging evidence from a very different paradigm that priming positive stimuli eliminates the attention bias toward negative information. Second, because participants were not given an evaluative task and they were not aware of the nature of the priming manipulation, these results show that the elimination of the attention bias is not dependent on having an evaluative goal.

A question that remains, however, concerns the nature of the manipulations used in the first two studies. Although both results are theoretically important, it is unlikely that a person is going to view a large number of positive pictures or be incidentally exposed to a large number of positive words in their daily life. These manipulations leave open the question of how likely is it that this bias could be attenuated in the everyday situations to which people are normally exposed. Addressing this question was the goal of Experiment 3.

Experiment 3

To provide a more externally valid manipulation of the accessibility of positive and negative thoughts, we had the experiment interact with the participant in either a positive or negative manner.

[7] The target words used by Pratto and John (1991, Experiment 3) crossed valence with trait frequency. Although they did not observe a significant interaction between the two variables ($p = .27$), their means suggested that the attention bias was larger for common words than for uncommon words. In our study, we presented both common and uncommon words. However, uncommon words yielded no effects in any experimental condition, suggesting that participants were not sufficiently familiar with the words for them to generate interference. Therefore, we are presenting the results generated by the 44 common words.

Method

Participants. Participants were 18 volunteers (10 women, 8 men; mean age = 20.8 years) recruited during summer school at Ohio Wesleyan University. The participants were not compensated for their participation.

Overview. Participants came into the lab and were informed they would be performing a color-naming task. Participants were randomly assigned to one of two conditions. In one condition the experimenter acted in a harsh manner toward the participants, and in the other condition the experimenter acted nicely toward the participants. All participants then performed a Stroop color-naming task, as in Experiment 2.

Procedure. Immediately after participants arrived at the lab, they were instructed to carefully read and sign an informed consent form describing the experiment. Participants were told that they would complete a color-naming task to test their verbal reaction times. Participants were questioned as to their age, gender, and native language. They were also asked whether they were colorblind. No colorblind participants were included in the experiment.

Participants then performed the Stroop task as in Experiment 2, with the exception that only the 44 common traits from Pratto and John (1991, Experiment 3) were used. After completing the Stroop task, participants took part in a manipulation check, in the form of an anonymous survey. This survey questioned participants about their level of enjoyment of the experiment and the likability of the experimenter, using two 7-point Likert scales. Participants were then probed for suspicion of the hypothesis. Finally, participants were debriefed (during which time they gave permission for their ostensibly anonymous surveys to be associated with their data), given a piece of candy to undo any residual negative mood resulting from the experimental manipulation, and released.

Experimenters. Both experimenters in this study were college-aged women. One always played the role of the friendly experimenter, the other the mean experimenter.

Meanness manipulation. The experimenters' behavior when interacting with a participant was a function of the participant's experimental condition. Participants in the nice experimenter condition were thanked repeatedly for their participation and were treated with kindness and respect. Participants were asked whether they were comfortable at multiple points throughout the procedure, and instructions were given in a friendly tone of voice. When adjusting the microphone to an optimal position near their mouth, the participants were thanked and told that they did well. After the instructions were given, participants were also asked whether they understood the experiment, and instructions were repeated if they had questions.

In the mean experimenter condition, participants were treated curtly and rudely. Participants were made to wait outside for 5 min before they were allowed to begin the experiment. Once the participants were in the experiment room, the door was slammed shut behind them. Instructions to the participants were framed as orders, instead of requests, and the tone of voice used was brusque and unfriendly. The instructions, "Remember, name the color and not the word," were repeated numerous times, with the experimenter remarking in a condescending fashion that she hoped that participants knew what each color looked like. When participants were adjusting the microphone, the experimenter twice told them that they were doing it incorrectly and told them to do it again. Once the Stroop task began, participants were treated equally regardless of experimental condition.

Results

This study was a 2 (trait valence: positive or negative) × 2 (experimenter style: friendly or mean) mixed-model design; the former factor was manipulated within subjects.

Manipulation check. To determine whether the experimenter style manipulation was effective, we compared participants' self-reported liking of the experimenter between the mean and friendly experimenter conditions. Suggesting that the experiment style manipulation was effective, a t test showed that participants in the mean experimenter condition reported liking the experimenter significantly less ($M = 4.9$, $SE = 1.74$) than those in the friendly experimenter condition ($M = 6.5$, $SE = .85$), $t(16) = -2.5$, $p = .024$.

The attention bias. Our main hypothesis was that the attention bias would be greater in the mean experimenter condition than in the friendly experimenter condition. Specifically, we expected a Trait Valence × Experimenter Style interaction. There were no significant main effects for trait valence or experimenter style (both Fs < 1). However, as predicted, there was a significant interaction between trait valence and experimenter style, $F(1, 16) = 4.740$, $p = .045$. Breaking down the interaction, in the mean experimenter condition, participants' color-naming latencies were longer for negative traits ($M = 667.4$, $SE = 34.1$) than positive traits ($M = 653.8$, $SE = 38.0$). Further, and consistent with our hypothesis, latencies in the friendly condition were longer for positive traits ($M = 682.2$, $SE = 34.0$) than for negative traits ($M = 669.7$, $SE = 30.5$; see Figure 5). Neither of these simple effects reached conventional significance levels.[8]

One possible alternative explanation for our results is that experimenters' behaviors could have depicted, and therefore directly primed, some of the traits used in the emotional Stroop task. For example, the friendly experimenter acted in an extroverted manner, which was a trait used in the Stroop task. Instead of altering participants' reaction times to positive and negative stimuli in general, it may be that reaction times were only increased to words that referred to traits that the experimenters displayed.

To rule out the hypothesis that behaviorally applicable traits yielded larger effects than nonapplicable traits, experimenters categorized the 44 trait adjectives used in the Stroop task as applicable or nonapplicable to their behavior. To evaluate the possibility that the attenuation effect only held for applicable traits, we performed a 2 (experimenter style: mean or friendly) × 2 (trait valence: positive or negative) × 2 (applicability of traits in experimenters behavior: applicable or nonapplicable) GLM. Although the three way-interaction was significant, $F(1, 16) = 5.668$, $p = .030$, it showed that the nonapplicable condition yielded a significant two-way interaction consistent with the experimental hypothesis, $F(1, 16) = 6.755$, $p = .019$, but the applicable condition did not, $F(1, 16) = 1.439$, $p = .248$. Therefore, an explanation

[8] Although the simple-effects test in Experiments 2 and 3 did not reach conventional levels of significance, it is important to note that our theoretical prediction was for a moderation of the attention bias. That is, the difference between the amount of attention paid to negative and positive stimuli should be smaller in the positive condition than the negative condition. This difference between differences prediction is most effectively tested with the overall interaction test. Adding more participants may make the null simple effects in Experiments 2 and 3 significant; however, this is not inconsistent with our predictions, because in both cases the interaction would still maintain that the attention bias is significantly smaller in the positive prime condition.

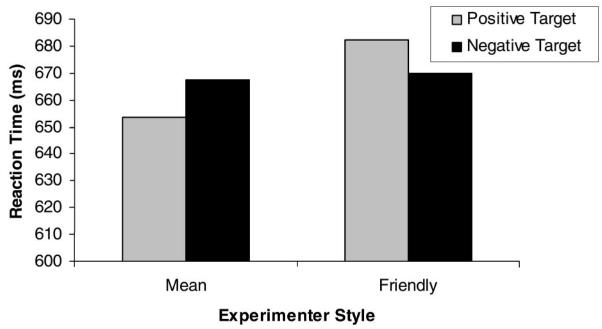

Figure 5. Reaction time as a function of word valence and experimenter style in Experiment 3. Participants' attention is biased toward negative traits in the mean experimenter condition and positive traits in the friendly experimenter condition.

suggesting that primed traits were driving the effect does not seem plausible.[9]

Discussion

Experiment 3 demonstrated that positive and negative social interactions have the power to bias attention toward affectively congruent information. This finding is an important addition to Experiments 1 and 2 because it suggests that the eliminating the attention bias does not require exotic manipulations.

There are two possible alternative explanations for the results of Experiment 3. The first, that the moderation of the attention bias was primarily driven by words that described the behaviors displayed by the experimenters, was ruled out previously. A second alternative explanation suggests that, although the experimenter style manipulation was altering the accessibility of positive and negative constructs, it may also have been altering other psychological variables within the participant, such as motivation level. However, the idea that motivational state is responsible for making people preferentially attend to valence-consistent information (i.e., people in positive motivational states would preferentially attend to positive information) is directly contradicted by the work of Rothermund (2003). In four studies, Rothermund (2003) showed that inducing valenced motivational states increased the attention paid to stimuli of the opposite valence. Further, if we want to conclude that participants in the positive social interaction condition had a higher level of motivation, it may be reasonable to conclude that they would have quicker reaction times than participants who had a lower level of motivation. This was not found to be the case. Therefore, this alternative explanation of the current results seems less tenable than an attentional explanation.

General Discussion

Across three experiments, using diverse manipulations and measures, we have replicated the attention bias to negative information and demonstrated that this bias can be eliminated when positive constructs are made accessible in memory. This finding is important for several reasons. First, it shows a moderator of the negativity attention bias, a bias that is often thought of as obligatory (Pratto & John, 1991).

In this vein, Experiment 3 calls into question the supposition that most of the time people's attention is biased toward negative information. If the attention bias is as malleable as other affective constructs, such as mood, then perhaps social psychologists need to begin to reconsider how we think of the bias. This is especially true in light of the fact that many of our participants are required to participate in psychological research and are often less than overjoyed to be doing so. Their putative negative affective state may be the reason why they pay more attention to negative information in experimental contexts.

A second interesting implication of this research is that, given that directly priming positive and negative constructs can moderate this bias, it also is possible that other preconscious mental states may do the same. For example, Chartrand and Jefferis (2003) have shown that certain goals (e.g., achievement goals) may be activated and pursued outside of awareness. When one succeeds or fails at these goals, one experiences positive and negative affect, respectively. Therefore, it is possible that one's attention can be biased by succeeding or failing at a goal one never knew one had.

The most important implication of this study, however, pertains to the role of attention bias as the mediator of other biases to negative information. That is, if one is paying a disproportionate amount of attention to negative information, it stands to reason that one would weight negative information heavily in decision making. For example, when reading a list of traits about a person, if one spends more time attending to negative traits than positive traits, it stands to reason that the negative traits will play a larger role in the final impression of the individual (consistent with Anderson, 1965). If the attention bias does mediate other biases to negative information, the finding that the attention bias can be eliminated suggests that it may be possible to eliminate other biases to negative information with similar manipulations.

Finally, these results may also have important implications for breaking down the self-maintaining nature of emotional disorders. It has been suggested that many mental disorders direct people's attention to the stimuli associated with the disorder, thus prolonging the disorder. For example, preferentially attending to panic-relevant stimuli in the environment has been shown to exacerbate panic disorder, making panic attacks more frequent. This, in turn, furthers the biasing of attention toward panic-relevant stimuli, and the cycle continues (D. M. Clark, 1988). Given that overattending to pathology-relevant stimuli can be a maintaining cause of psychopathology (Williams, Mathews, & MacLeod, 1996), manipulating the types of stimuli people attend to may help decrease the cyclic nature of the pathology. There is some evidence that the attention biases caused by mental disorders take time to accrue. For instance, Williams and Nulty (1986) have shown that emotional Stroop interference times are more related to depression status 12 months previously than current depression status, suggesting that chronic pathology-related attention biases take time to develop. Therefore, interventions designed to decrease the relative

[9] This should not be taken to imply that behavior-consistent trait adjectives had no effect overall. Of the 44 traits used in the study, 15 were deemed to be applicable to the experimenters. Therefore, the participant means for applicable words were based on fewer reaction times and were, therefore, more variable than the means for nonapplicable words. This may have been responsible for the three-way interaction detailed previously.

accessibility of negative thoughts (e.g., positive subliminal priming procedures, such as the one performed in Experiment 2) might be able to help stave off this accrual of chronic accessibility for depression-relevant stimuli and thus be beneficial in attenuating this aspect of the self-maintaining nature of depression. This hypothesis is consistent with the work of Fredrickson and Joiner (2002), who suggested that positive emotions can create a self-reinforcing cycle of well-being.

Although we contend that the bias to negative information is attenuated because of the increased accessibility of positive information, it is also possible that repeatedly presenting people with positive information is additionally changing some other variable that has been hypothesized to play a role in causing negativity biases. For example, repeated positive subliminal priming has been shown to improve mood (Bargh & Chartrand, 1999), and good moods have been associated with multiple psychological changes, including increased creativity and flexible thinking (Isen, 1999) and a decreased ability to carefully process information (Mackie & Worth, 1989).

Although it is possible that mood is playing a role in the effects described, several factors limit the usefulness of this alternative interpretation. First, many of the effects of mood (e.g., increased creativity and flexibility, as in Isen, 1999, or broad-minded coping; see Fredrickson & Joiner, 2002) are downstream processes that may not be able to affect the attention allocated to stimuli within 100 ms of initial exposure. Second, because mood itself may be a conscious interpretation of the unconscious accessibility of positive and negative information (Bargh & Chartrand, 1999), it is debatable whether mood causing the just-mentioned effects is an alternative explanation or whether the accessibility of positive and negative information is the mechanism through which mood may modify attention. A future study that dissociated the unconscious accessibility of positive and negative constructs (by means of subliminal priming) and mood (by explicit manipulation) may be successful at disentangling the influences of these two closely related constructs.

A second alternative explanation may suggest that psychological states previously implicated in biases to negative information (e.g., expectations about valence and diagnosticity of information) may also be altered by strongly skewing the type of stimuli to which participants are exposed. Based on prior work, these alternative mechanisms would predict effects opposite those found. That is, researchers have suggested that we pay more attention to negative information because it violates our expectations. In contexts in which most of the stimuli are positive, logic dictates that negative stimuli would be even more unexpected and thus demand more attention. Similarly, when most stimuli are positive, the diagnostic value of positive information would seem to be lessened, and thus it would merit less attention. Additionally, as argued previously, the time frame and physical location of these effects (i.e., processing that occurs in the extrastriate visual areas 100 ms poststimulus onset) seem to make it unlikely that higher order processes (e.g., diagnostic) would be able to explain these effects.

There is another possible interpretation of these data. To this point, we have been conceptualizing the attention bias as the amount of attention paid to positive stimuli subtracted from the amount of attention paid to negative stimuli. That is, we have been discussing the attention bias as a relative measure. It is also possible to think of the attention bias as an absolute, suggesting that negative information receives a large amount of attention regardless of the context. Looking across all three experiments, negative stimuli seem to be attracting a relatively consistent and relatively sizable amount of attention regardless of whether they are perceived in a positive or negative context (although, admittedly, in Experiments 2 and 3 this is a between-participants comparison and should be interpreted with caution). Although it does not qualify any of the effects discussed here, this interpretation suggests that the attention bias is always present; however, the affective context can serve to mask its appearance.

In addition to the implications listed previously, several other questions were left open by this research that should be addressed in the future. First, in each of the three experiments, positive and negative stimuli are presented in isolation. Although this is sufficient for establishing the moderation of the attention bias, it does not allow an assessment of the ability of positive stimuli to compete with simultaneously presented negative stimuli, a more ecologically valid situation.

A second interesting question concerns the longevity of these biases in attention. Similar to the psychopathology argument given previously, having attention biased toward a particular valence of stimuli would likely cause people to perceive and process more stimuli of that valence. This differential perception and processing would then, in turn, affect what constructs they have accessible. Therefore, one might expect these biases in attention to endure longer than the typical priming effect because the opportunities for self-reinforcement are more frequent than for other categories of stimuli.

In conclusion, the current work suggests a new conceptualization of the attention bias toward negative information. That is, rather than a hard and fast rule of evolution, the bias should be considered a default that may be overridden by situational forces.

References

Anderson, N. H. (1965). Averaging versus adding as a stimulus-combination rule in impression formation. *Journal of Personality and Social Psychology, 2,* 1–9.

Bargh, J. A. (1989). Conditional automaticity: Varieties of automatic influence in social perception and cognition. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 3–51). New York: Guilford Press.

Bargh, J. A., Bond, R. N., Lombardi, W. J., & Tota, M. E. (1986). The additive natures of chronic and temporary sources of accessibility. *Journal of Personality and Social Psychology, 50,* 869–878.

Bargh, J. A., Chaiken, S., Govender, R., & Pratto, F. (1992). The generality of the automatic attitude activation effect. *Journal of Personality and Social Psychology, 62,* 893–912.

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

Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad is stronger than good. *Review of General Psychology, 5,* 323–370.

Cacioppo, J. T., Crites, S. L., Berntson, G. G., & Coles, M. G. (1993). If attitudes affect how stimuli are processed, should they not affect the event-related brain potential? *Psychological Science, 4,* 108–112.

Cacioppo, J. T., Gardner, W. L., & Berntson, G. G. (1997). Beyond bipolar conceptualizations and measures: The case of attitudes and evaluative space. *Personality and Social Psychology Review, 1,* 3–25.

Chartrand, T. L., & Bargh, J. A. (2001). *Affective consequences of automatic evaluation: Subliminal exposure to valenced words affects mood.* Unpublished manuscript.

Chartrand, T. L., & Jefferis, V. E. (2003). Consequences of automatic goal pursuit and the case of nonconscious mimicry. In J. Forgas, K. Williams, & W. Von Hippel (Eds.), *Social judgments: Implicit and explicit processes* (pp. 290–305). New York: Cambridge University Press.

Clark, D. M. (1988). A cognitive model of panic attacks. In S. Rachman & J. D. Maser (Eds.), *Panic: Psychological perspectives* (pp. 71–89). Hillsdale, NJ: Erlbaum

Clark, V. P., & Hillyard, S. A. (1996). Spatial selective attention affects early extrastriate but not striate components of the visual evoked potential. *Journal of Cognitive Neuroscience, 8,* 387–402.

Donchin, E., & Heffley, E. (1978). Multivariate analysis of event related potential data: A tutorial review. In D. Otto (Ed.), *Multidisciplinary perspectives in event-related brain potential research* (EPA-600/9–77–043). Washington, DC: U.S. Government Printing Office.

Fiske, S. T. (1980). Attention and weight in person perception: The impact of negative and extreme information. *Journal of Personality and Social Psychology, 80,* 491–499.

Fredrickson, B. L., & Joiner, T. (2002). Positive emotions trigger upward spirals toward emotional well-being. *Psychological Science, 13,* 172–175.

Graziano, W. G., Brothen, T., & Berscheid, E. (1980). Attention, attraction, and individual differences in reaction to criticism. *Journal of Personality and Social Psychology, 38,* 193–202.

Hansen, C. H., & Hansen, R. D. (1988). Finding the face in the crowd: An anger superiority effect. *Journal of Personality and Social Psychology, 54,* 917–924.

Isen, A. M. (1999). Positive affect. In T. Dalgleish & M. Power (Eds.), *Handbook of cognition and emotion* (pp. 521–539). New York: Wiley.

Ito, T. A., Cacioppo, J. T., & Lang, P. J. (1998). Eliciting affect using the International Affective Picture System: Trajectories through evaluative space. *Personality and Social Psychology Bulletin, 24,* 855–879.

Ito, T. A., Larsen, J. T., Smith, N. K., & Cacioppo, J. T. (1998). Negative information weighs more heavily on the brain: The negativity bias in evaluative categorizations. *Journal of Personality and Social Psychology, 75,* 887–900.

Johnston, W. A., & Hawley, K. J. (1994). Perceptual inhibition of expected inputs: The key that opens closed minds. *Psychonomic Bulletin and Review, 1,* 56–72.

Kahneman, D., & Tversky, A. (1984). Choices, values and frames. *American Psychologist, 39,* 341–350.

Lang, P. J., Bradley, M. M., & Cuthbert, B. N. (1995). *International affective picture system (IAPS): Technical manual and affective ratings.* Gainesville: University of Florida, NIMH Center for the Study of Emotion and Attention.

Mackie, D. M., & Worth, L. T. (1989). Processing deficits and the mediation of positive affect in persuasion. *Journal of Personality and Social Psychology, 57,* 27–40.

Neely, J. H. (1977). Semantic priming and the retrieval from lexical memory: Evidence for facilitory and inhibitory processes. *Memory & Cognition, 4,* 648–654.

Öhman, A., Lundqvist, D., & Esteves, F. (2001). The face in the crowd revisited: A threat advantage with schematic stimuli. *Journal of Personality and Social Psychology, 80,* 381–396.

Peeters, G., & Czapinski, J. (1990). Positive-negative asymmetry in evaluations: The distinction between affective and informational negativity effects. In W. Stroebe & M. Hewstone (Eds.), *European review of social psychology* (Vol. 1, pp. 33–60). Chichester, UK: Wiley.

Pratto, F., & John, O. P. (1991). Automatic vigilance: The attention-grabbing power of negative social information. *Journal of Personality and Social Psychology, 61,* 380–391.

Roskos-Ewoldsen, D. R., & Fazio, R. H. (1992). On the orienting value of attitudes: Attitude accessibility as a determinant of an object's attraction of visual attention. *Journal of Personality and Social Psychology, 63,* 198–211.

Rothermund, K. (2003). Motivation and attention: Incongruent effects of feedback on the processing of valence. *Emotion, 3,* 223–238.

Semlitsch, H. V., Anderer, P., Schuster, P., & Presslich, O. (1986). A solution for reliable and valid reduction of ocular artifacts, applied to the P300 ERP. *Psychophysiology, 23,* 695–703.

Skowronski, J. J., & Carlston, D. E., (1989). Negativity and extremity biases in impression formation: A review of explanations. *Psychological Bulletin, 105,* 131–142.

Smith, N. K., Cacioppo, J. T., Larsen, J. T., & Chartrand, T. L. (2003). May I have your attention, please: Electrocortical responses to positive and negative stimuli. *Neuropsychologia, 41,* 171–183.

Spence, J. T., & Segner, L. L. (1967). Verbal versus nonverbal reinforcement combinations in the discrimination learning of middle- and lower-class children. *Child Development, 38,* 29–38.

Taylor, S. E. (1991). Asymmetrical effects of positive and negative events: The mobilization-minimization hypothesis. *Psychological Bulletin, 110,* 67–85.

Van Boxtel, G. (1998). Computational and statistical methods for analyzing event-related potential data. *Behavior Research Methods, Instruments & Computers, 30,* 87–102.

Walker, W. R., Vogl, R. J., & Thompson, C. P. (1997). Autobiographical memory: Unpleasantness fades faster than pleasantness over time. *Applied Cognitive Psychology, 11,* 399–413.

Wegener, D. T., & Petty, R. E. (1994). Mood management across affective states: The hedonic contingency hypothesis. *Journal of Personality and Social Psychology, 66,* 1034–1048.

Williams, J. M. G., Mathews, A., & MacLeod, C. (1996). The emotional Stroop task and psychopathology. *Psychological Bulletin, 120,* 3–24.

Williams, J. M. G., & Nulty, D. D. (1986). Construct accessibility, depression and the emotional Stroop task: Transient emotion or stable structure? *Personality and Individual Differences, 7,* 485–491.

Received July 8, 2004
Revision received June 3, 2005
Accepted June 6, 2005 ■