Research Article

ON WILDEBEESTS AND HUMANS: The Preferential Detection of Negative Stimuli

Ap Dijksterhuis[1] and Henk Aarts[2]

[1]*University of Amsterdam, Amsterdam, the Netherlands, and* [2]*University of Utrecht, Utrecht, the Netherlands*

Abstract—*On the basis of a functional perspective, we hypothesized that negative stimuli are detected faster than positive stimuli. In Experiment 1, participants were subliminally presented with positive and negative words or with no words at all. After each presentation, participants were asked whether they had seen a word. They detected negative words more accurately than positive words. In Experiment 2, participants were subliminally presented with negative or positive words. After each presentation, they were asked whether the presented word was positive or negative. Negative words were correctly categorized more often than positive words. Experiment 3 showed that although participants correctly categorized negative words more often than positive words, they could not guess the meaning of the words better than would be expected by chance. The results are discussed against the background of recent findings on basic affective processes.*

The structure of every organic being is related, in the most essential yet often hidden manner, to that of all other organic beings . . . from which it has to escape.
—Darwin (1859/1996, p. 127)

The image that Darwin's quote brings to mind may be that of an African savanna. In such an environment, the relation between physical features of predators and prey is evident. Cheetahs and lions are fast, and their potential prey reflect this capacity: Wildebeests and gazelles are fast too, and the ones that are not are devoured. However, Darwin wanted to emphasize the "hidden" capacities. Although it is of obvious importance for a wildebeest to be able to run fast, it is at least as important for a wildebeest to detect an approaching lion or cheetah quickly. The wildebeest's perceptual and affective systems should therefore be shaped in relation to its environment. And this should be true for all animals: At times, all animals are confronted with threatening stimuli, and it is of utmost importance to detect these stimuli as fast as possible.

Two phenomena related to the perception of positive and negative stimuli have been studied extensively: *automatic evaluation* and *automatic vigilance*. Automatic evaluation refers to the capacity to evaluate incoming stimuli automatically. The importance of this capacity is reflected in the lack of its boundary conditions. Humans evaluate all stimuli (Bargh, Chaiken, Raymond, & Hymes, 1996; but see Fazio, Sanbonmatsu, Powell, & Kardes, 1986) regardless of an intention to evaluate (Hermans, De Houwer, & Eelen, 1994). In addition, evaluation does not require conscious awareness of the meaning of the stimulus (Bargh, Litt, Pratto, & Spielman, 1989; Greenwald, Klinger, & Liu, 1989; Murphy & Zajonc, 1993; see also De Houwer, Hermans, & Spruyt, 2001). Automatic evaluation is obviously functional. A quick categorization of stimuli allows for the rapid onset of appropriate behavior (i.e., approach or avoidance).

Address correspondence to Ap Dijksterhuis, Social Psychology Program, University of Amsterdam, Roetersstraat 15, 1018 WB Amsterdam, the Netherlands; e-mail: sp_dijksterhuis@macmail.psy.uva.nl.

Research on automatic vigilance demonstrates that negative stimuli demand more attention than positive stimuli. Various researchers have shown that the processing of negative words interferes with other information processing to a greater extent than does the processing of positive words (Pratto & John, 1991; Wentura, Rothermund, & Bak, 2000; Williams, Matthews, & MacLeod, 1996). This effect is also highly functional: It means that whenever negative stimuli are encountered, individuals are forced to process them more elaborately than other stimuli.

In the study we report here, we tested a new hypothesis derived from this functional perspective. It would be highly functional for a negative stimulus to be detected as fast as possible, whereas this is less important for a positive stimulus. Should a wildebeest not detect a negative stimulus even faster than a positive stimulus? After all, the evolutionary pressure on detecting negative stimuli quickly should be stronger than the pressure on detecting positive stimuli quickly. Being a few hundred milliseconds late in detecting a lion is extremely dangerous, whereas being a little late in detecting edible vegetation is not so problematic. Hence, we hypothesized that it requires less stimulus input or less stimulus exposure to detect a negative stimulus than to detect a positive stimulus. It should be noted that the effect we hypothesized is fundamentally different from automatic vigilance. Automatic vigilance does not mean that negative stimuli are detected faster than positive stimuli. Instead, it refers to the fact that once detected, negative stimuli receive more attention than positive stimuli.

Our use of the term "detection" requires explanation. We aimed to assess the moment that a presented stimulus is evaluatively categorized more accurately than would be expected on the basis of chance, even though other properties of the stimulus (such as its meaning) are not accessible to consciousness. Our operationalization was based on an experiment by Bargh et al. (1989; see also Marcel, 1983, for his use of the same paradigm) in which words were presented to participants. The presentation durations varied but were always subliminal. After the presentation of each word, participants were asked for an evaluative judgment (Was the word positive or negative?) and a semantic judgment (Which of two words is a synonym of the presented word?). When stimuli were presented for brief durations, the proportion of words evaluated correctly was above chance level, but the participants still could not identify their meaning. Our hypothesis entails that such categorizations are made faster for negative than for positive stimuli.

The question is how such an asymmetry between the detection of negative and positive stimuli might arise. One could object that categorization of a stimulus as positive or as negative takes place at one point in time. After all, are positive and negative not two sides of the same coin? Recent work demonstrates that this is not the case, however. Responses to negative stimuli are largely independent from reactions to positive stimuli (e.g., Cacioppo, Crites, Berntson, & Coles, 1993; Davidson, Ekman, Saron, Senulis, & Friesen, 1990). Cacioppo et al. (1993) argued that evaluation should not be conceptualized in terms of one positive-negative dimension, but rather should be conceptualized in terms of two independent dimensions. LeDoux (1996) presented a convincing evolutionary argument: Systems that are responsible for eliciting different emotions developed largely independently of each other. The human system for detecting stimuli eliciting fear, for instance, developed independently from, say, the system responsible for happiness (see also Öhman & Wiens, 2001). Work by Davidson et al. (1990) demonstrates that whereas positive stimuli evoke more activity in the left than in the right hemisphere, negative stimuli evoke the opposite pattern. This knowledge leads to what we propose as the underlying mechanism for our hypothesis: The threshold for eliciting negative affect may well be lower than that for instigating positive affect. That is, it is possible that a briefly presented negative stimulus evokes detectable negative affect, whereas a positive stimulus presented for the same duration does not yet elicit detectable positive affect.

The research that comes closest to a test of our hypothesis was conducted by Hansen and Hansen (1988). In their experiments, participants were presented with an array of happy and angry faces, either one happy face amidst a number of angry faces or vice versa. The task was to locate the deviating face as quickly as possible. Participants were faster to locate the angry face amidst the happy ones than they were to locate the happy face amidst the angry ones. The authors concluded that angry faces were recognized faster than happy ones and that angry faces, as it were, grabbed attention by "popping out of the crowd." However, their experiments did not directly test our hypothesis because they did not assess detection. Participants took a long time to locate the target faces (almost 2 s on average), and indeed, the task forced participants to consciously recognize the target stimulus. The literature on affective priming demonstrates that evaluative categorization occurs before conscious recognition (e.g., Bargh et al., 1989; De Houwer, Baeyens, & Eelen, 1994).

A close look at the literature on affective priming (e.g., Bargh, Chaiken, Govender, & Pratto, 1992; Bargh et al., 1996; De Houwer et al., 2001; Fazio, 2001; Fazio et al., 1986; Glaser & Banaji, 1999; Greenwald et al., 1989; Hermans et al., 1994; Klauer, 1998; Klauer, Rossnagel, & Musch, 1997) does indirectly support our hypothesis. In an affective priming experiment, on each trial participants are presented with a positive or negative prime, followed by a positive or negative target. Responses to target words are faster if prime and target are evaluatively congruent (both positive or both negative) than if they are incongruent (one negative, the other positive). This paradigm enables a distinction between four prime-target pairs: positive-positive, positive-negative, negative-positive, and negative-negative. This in turn allows one to independently assess the impact of a negative prime (by comparing negative-negative pairs with negative-positive pairs) and the impact of a positive prime (by comparing positive-positive pairs with positive-negative pairs). The vast majority of experiments show a larger effect for negative primes than positive primes.[1]

Another common finding in the automatic-evaluation literature is that participants respond faster to positive than to negative stimuli. That is, researchers often report a main effect of valence of target.[2] One could infer from this effect that positive stimuli are detected faster than negative stimuli. However, the typical response in automatic-evaluation research requires conscious recognition. If anything, faster recognition of positive than of negative stimuli follows from automatic vigilance. The response to negative stimuli is slower because the greater information processing they elicit interferes more with selecting and executing a response. This explanation is in line with the literature on perceptual defense. In perceptual-defense research, participants are presented with words and are asked to verbalize them (e.g., Bootzin & Natsoulas, 1965; Broadbent & Gregory, 1967; Eriksen, 1963; McGinnies, 1949). Experiments often show that participants are particularly slow to verbalize negative, taboo words. This effect can be explained by vigilance (Blum, 1954; Kitayama, 1990): The attention that taboo words demand interferes with verbalization.

In sum, negative stimuli call for more attention than positive stimuli, and tasks requiring conscious recognition of stimuli usually show reactions to negative stimuli are slower than reactions to positive stimuli. It is known that a stimulus can be categorized as positive or negative with an accuracy greater than chance before it is consciously recognized (e.g., Bargh et al., 1989), but whether a negative stimulus is detected faster than a positive one remains unclear. We wanted to shed light on this possibility.

OVERVIEW OF THE EXPERIMENTS

In three experiments, participants were subliminally presented with positive and negative words. In Experiment 1, in half the trials a positive or a negative word was shown, and in the remaining trials no words were shown. Participants were asked to indicate whether they thought a word had appeared or not. We predicted that participants would correctly indicate a word was flashed more frequently for negative than for positive words. In Experiment 2, in all trials words were shown. In half the trials the word was negative, and in the remaining half the word was positive. Participants were asked whether they thought the word presented was positive or negative. In this case, we predicted that participants would correctly categorize negative words as negative more often than they correctly categorized positive words as positive. Experiment 3 was the same as Experiment 2 with one exception. In addition to making an evaluative judgment (Is the word positive or negative?), participants were asked to make a semantic judgment, by indicating which of two words was a synonym of the presented word. We predicted that participants would more often correctly categorize negative words as negative than positive words as positive, but would not be able to make accurate semantic judgments.

1. We compared the effects of positive and negative primes by calculating the difference in response time between positive-positive pairs and positive-negative pairs on the one hand and between negative-negative pairs and negative-positive pairs on the other hand. Work in which only the distinction between evaluatively congruent versus evaluatively incongruent was made could not be included. For our analysis, we used the experimental conditions of Bargh et al. (1992, 1996), De Houwer et al. (2001), Fazio et al. (1986), Greenwald et al. (1989), Hermans et al. (1994), and Klauer et al. (1997). In 27 experimental conditions, negative primes had more impact than positive primes. In 12 cases, positive primes had more impact. Thirteen cases were categorized as neutral (with a difference of 10 ms or less).

2. We assessed the responses to negative and positive targets regardless of prime using the same experimental conditions as for our analysis of the impact of prime (footnote 1), excluding the data from Bargh et al. (1996), who used a pronunciation task. In all 34 cases, participants responded faster to positive targets than to negative targets.

EXPERIMENT 1

Method

Participants

Twenty-five undergraduate students from the University of Nijmegen, Nijmegen, the Netherlands, participated in the experiment. They received 5 Dutch guilders ($2) in return.

Procedure and materials

Upon entering the laboratory, participants were led to a cubicle and seated in front of a computer. All instructions were provided by the computer. The participants were told that the experimenters were interested in how long a word needed to be presented for people to be able to recognize it. They were told that they would repeatedly see a row of six *X*s on the screen. This row would remain on the screen for 500 ms. Immediately after the row of *X*s disappeared, either a word would be flashed for 13.3 ms or the screen would remain blank for 13.3 ms. Participants were explicitly told that in 50% of the trials a word would appear and in 50% no word would appear. Immediately afterward, the row of *X*s would appear again and remain on the screen for 500 ms. After the postmask had disappeared, participants had to indicate whether they thought a word had been presented or not, by pressing "1" (word) or "2" (no word). They were explicitly told that it was highly likely that they would not be able to see a word at all, because of the short exposure duration. They were asked to guess. The words and the masks appeared on the center of the screen in Chicago 14 font.

Participants received 60 trials, 30 trials in which no word was presented, 15 trials in which a positive word was presented, and 15 trials in which a negative word was presented. The trials were presented in random order. The 15 positive and the 15 negative words were selected on the basis of pilot testing of 151 words. Thirty-five participants evaluated 3- to 6-letter words on a scale ranging from 1 (*extremely negative*) to 9 (*extremely positive*). The words selected for use in the experiment were all of extreme valence ($M > 7.8$ or $M < 2.3$).[3] All words were of medium frequency, but the mean frequency of the positive words was slightly higher than that of the negative words as positive words generally occur more frequently (Zajonc, 1968). One should note, though, that any influence of this difference in frequency on our results would be in the direction opposite to our hypothesis.

After participants had finished the 60 trials, they were asked whether they had been able to correctly identify some of the words. None of the participants had. The vast majority indicated that they had not seen anything flash at all and that they felt that they merely guessed throughout the experiment. Participants were all thoroughly debriefed.

3. The positive words were *baby* (baby), *gein* (fun), *geluk* (happiness), *katje* (kitten), *lach* (smile), *lente* (spring), *lief* (sweet), *reis* (trip), *strand* (beach), *thuis* (home), *vriend* (friend), *vrij* (free), *zacht* (soft), *zomer* (summer), and *zon* (sun). The negative words were *angst* (fear), *bom* (bomb), *bruut* (rude), *coma* (coma), *dief* (thief), *dood* (dead), *gemeen* (mean), *haai* (shark), *hel* (hell), *kanker* (cancer), *oorlog* (war), *pest* (plague), *polio* (polio), *slang* (snake), and *wapen* (weapon).

Results and Discussion

The proportions of positive and negative words correctly categorized as words were calculated for each participant. As expected, negative words were categorized more often as words ($M = .545$, $SD = .27$) than were positive words ($M = .401$, $SD = .21$), $F(1, 24) = 5.20$, $p < .032$.

Experiment 1 confirmed our hypothesis. With the same brief stimulus exposure, participants were better able to detect negative words than positive words, although detection of negative words was not better than chance. Whereas it is clear that negative words were detected with greater accuracy than positive words, it is not clear whether positive words were sometimes detected as well because the 40% correct identification of positive words could simply be the consequence of participants' knowledge that in 50% of the cases a word was presented. It should be noted in this regard that the false alarm rate (i.e., incorrect responses when no word was presented) was .335. This number did not differ significantly from the hit rate for positive words.

In addition, it is not clear whether participants could categorize the words as positive or negative. The experimental design allows us to conclude that on some occasions stimuli were categorized as words, but does not support a conclusion as to whether these words were evaluated as positive or negative. Experiment 2 was designed to explore whether participants sometimes detected positive words and whether they could categorize the valence of the stimuli.

EXPERIMENT 2

Method

Participants

Fifty-six undergraduate students from the University of Nijmegen participated in the experiment. They received 5 Dutch guilders ($2) in return.

Procedure and materials

The procedure for Experiment 2 was similar to the procedure for Experiment 1, except that a word was shown in every trial. Participants were told that 50% of the time the word would be positive and 50% of the time the word would be negative. After the postmask disappeared, participants had to indicate whether they thought the word was positive or negative by pressing the appropriate key. Allocation of the choices "positive" and "negative" to the buttons "1" and "2" was counterbalanced across participants. Again, participants were explicitly told that it was highly likely that they would not be able to see a word at all, because of the short exposure duration, but that they should guess nonetheless.

In total, participants received 30 trials. In 15 trials a positive word was presented, and in 15 trials a negative word was presented. The trials were presented in random order. The positive and the negative words were the same as in Experiment 1.

After participants had finished the trials, they were asked whether they had been able to correctly identify some of the words. None of the participants had. Again, the vast majority indicated that they had not seen anything flash at all and that they felt that they merely guessed throughout the experiment.

Results and Discussion

For all participants, the proportions of correctly categorized positive and negative words were calculated. Confirming our hypothesis, the proportion of correctly identified negative words (M = .563, SD = .22) was higher than the proportion of correctly identified positive words (M = .480, SD = .23), $F(1, 55) = 5.27$, $p < .027$. Moreover, the proportion of negative words categorized correctly also differed significantly from chance, $t(56) = 2.14$, $p < .041$.

As in Experiment 1, negative words were categorized as words more accurately than positive words. On the basis of Experiment 2, we can conclude that negative words were evaluatively categorized more accurately than positive words. In fact, our presentation conditions (13.3 ms in Chicago 14 font, with pre- and postmask) critically differentiated between the detection of positive and negative words.

In both experiments, participants detected negative words with greater accuracy than positive words. But what exactly was detected? We maintain that what participants displayed was a detectable affective response to the words and that the threshold for such a response is lower for negative than for positive affect. Furthermore, we argue that such a response is instigated before participants are able to report the meaning of the word (see also Bargh et al., 1989). The findings of the first two experiments, however, are also consistent with an alternative explanation. It is possible that negative words draw attention to their content (as could be predicted from the work on vigilance) and that the meaning of a negative word is detected more easily than the meaning of a positive word. If this is the case, our participants may have been able to categorize negative words more accurately than positive words not because of detectable affect but because they detected the meaning of the negative words better than the meaning of the positive words. In Experiment 3, our aim was to show that the effect we observed was due to detection of affective responses to the words, and not due to better detection of the meaning of negative than positive stimuli.

EXPERIMENT 3

Method

Participants

Thirty-one undergraduate students from the University of Amsterdam, Amsterdam, the Netherlands, participated in the experiment. They received 5 Dutch guilders ($2) in return.

Procedure and materials

The procedure for Experiment 3 was the same as the procedure for Experiment 2 with one exception. In Experiment 2, participants were asked to categorize each word with respect to its valence. In Experiment 3, we also asked participants to categorize the word in terms of its meaning. That is, after each word was presented, participants were asked two questions. The question assessing correct evaluation was administered the same way as in Experiment 2. The question assessing meaning was based on the study by Bargh et al. (1989) discussed earlier. Two words appeared on the screen. One was a synonym or close synonym of the presented word, and the other word was unrelated. However, both alternatives always had the same valence as the stimulus word. Participants were asked to guess which of the two words was the synonym of the presented word by pressing one of two keys.[4]

During debriefing, participants indicated that they had not seen anything flash on the screen at all and that they felt that they had merely guessed throughout the experiment.

Results and Discussion

For all participants, the proportions of correctly evaluated negative and positive words were determined. Confirming our hypothesis, the proportion of correctly identified negative words (M = .577, SD = .13) was higher than the proportion of correctly identified positive words (M = .513, SD = .11), $F(1, 30) = 4.55$, $p < .041$. Moreover, the proportion of correctly categorized negative words differed reliably from chance, $t(31) = 3.36$, $p < .005$. In addition, we determined the proportion of correct responses to the questions assessing meaning, and no differences were found between responses to negative and positive words (both Ms = .520), $F(1, 30) = 0.00$, n.s. In sum, although participants were again better in detecting negative words than positive words, this was due to the affective response that was detected and not to superior semantic processing.

GENERAL DISCUSSION

Throughout this article, we have emphasized a functional perspective on the processing of negative and positive stimuli. Whereas fast detection of negative information is often crucial, fast detection of positive information is much less so. This should have led to more evolutionary pressure on the development of a system to quickly detect negative information relative to positive information. Our experiments confirmed this reasoning: People need less stimulus input to detect a negative stimulus than to detect a positive stimulus.

We concede that we used our evolutionary argument rather loosely. Not all negative stimuli are actually threatening, and not all positive stimuli are appetitive. For instance, a word such as *shark* is both negative and threatening, whereas a word such as *boredom* is only negative. It would be interesting to see whether our results hold for all negative and all positive stimuli. Some researchers might suggest that it would. Cacioppo, Gardner, and Berntson (1999), for instance, suggested that people first process the valence of a stimulus and only a little later determine other aspects, such as the potential threat of a stimulus that is categorized as negative. It is also possible, however, that only threatening negative stimuli are detected faster than other stimuli. Recently, Wentura et al. (2000) published an interesting set of studies on automatic vigilance showing that the categorization of positive versus negative is unnecessarily crude. They used personality traits as stimuli and distinguished between "other-relevant traits" (such as brutal or aggressive) and "possessor-relevant traits" (such as depressive). Whereas negative other-relevant traits pose a potential threat to a perceiver, possessor-relevant traits do not. Indeed, negative other-relevant traits were particularly attention demanding, and negative possessor-relevant traits were not. Given that our current work is essentially based on the same functional perspective, it is plausible

4. Note that the order in which the two questions were presented could not be counterbalanced. Presentation of the semantic question would give away the answer to the evaluative question, so the evaluative question always had to come first.

that the distinction made by Wentura et al. applies also to the initial detection of stimuli. Further research could shed light on this interesting issue.

Acknowledgments—The authors would like to thank Karl Christoph Klauer, John Skowronski, and Pamela Smith for very helpful comments on an earlier draft of the manuscript.

REFERENCES

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

Bargh, J.A., Chaiken, S., Raymond, P., & Hymes, C. (1996). The automatic evaluation effect: Unconditional automatic attitude activation with a pronunciation task. *Journal of Experimental Social Psychology*, *32*, 104–128.

Bargh, J.A., Litt, J., Pratto, F., & Spielman, L.A. (1989). On the preconscious evaluation of social stimuli. In A.F. Bennett & K.M. McConkey (Eds.), *Cognition in individual and social contexts: Proceedings of the XXV International Congress of Psychology* (pp. 357–370). Amsterdam: Elsevier.

Blum, G.S. (1954). An experimental reunion of psychoanalytic theory with perceptual vigilance and defense. *Journal of Abnormal and Social Psychology*, *49*, 94–98.

Bootzin, R.R., & Natsoulas, T. (1965). Evidence for perceptual defense uncontaminated by response bias. *Journal of Personality and Social Psychology*, *5*, 461–468.

Broadbent, D.E., & Gregory, M. (1967). Perception of emotionally toned words. *Nature*, *215*, 581–584.

Cacioppo, J.T., Crites, S.L., Jr., Berntson, G.G., & Coles, M.G.H. (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. (1999). The affect system has parallel and integrative processing components: Form follows function. *Journal of Personality and Social Psychology*, *76*, 839–855.

Darwin, C. (1996). *The origin of species*. New York: Gramercy. (Original work published 1859)

Davidson, R.J., Ekman, P., Saron, R.D., Senulis, J.A., & Friesen, W.V. (1990). Approach-withdrawal and cerebral asymmetry: Emotional expression and brain physiology. *Journal of Personality and Social Psychology*, *58*, 330–341.

De Houwer, J., Baeyens, F., & Eelen, P. (1994). Verbal conditioning with undetected US presentations. *Behavior Research and Therapy*, *32*, 629–633.

De Houwer, J., Hermans, D., & Spruyt, A. (2001). Affective priming of pronunciation responses: Effects of target degradation. *Journal of Experimental Social Psychology*, *37*, 85–91.

Eriksen, C.W. (1963). Perception and personality. In J.M. Wepman & R.W. Heine (Eds.), *Concepts of personality* (pp. 31–62). Chicago: Aldine.

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

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

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., Klinger, M.R., & Liu, T.J. (1989). Unconscious processing of dichotically masked words. *Memory & Cognition*, *17*, 35–47.

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.

Hermans, D., De Houwer, J., & Eelen, P. (1994). The affective priming effect: Automatic activation of evaluative information in memory. *Cognition and Emotion*, *8*, 515–533.

Kitayama, S. (1990). Interaction between affect and cognition in word perception. *Journal of Personality and Social Psychology*, *58*, 209–217.

Klauer, K.C. (1998). Affective priming. *European Review of Social Psychology*, *8*, 63–107.

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

LeDoux, J. (1996). *The emotional brain*. New York: Simon & Schuster.

Marcel, A.J. (1983). Conscious and unconscious perception: An approach to the relations between phenomenal experience and perceptual processes. *Cognitive Psychology*, *15*, 238–300.

McGinnies, E.M. (1949). Emotionality and perceptual defense. *Psychological Review*, *56*, 471–482.

Murphy, S., & Zajonc, R. (1993). Affect, cognition and awareness: Affective priming with suboptimal and optimal stimuli. *Journal of Personality and Social Psychology*, *64*, 723–729.

Öhman, A., & Wiens, S. (2001, June). *The concept of an evolved fear module as a challenge to cognitive theories of anxiety.* Paper presented at Feelings and Emotions: The Amsterdam Symposium, Amsterdam.

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.

Wentura, D., Rothermund, K., & Bak, P. (2000). Automatic vigilance: The attention-grabbing power of approach- and avoidance-related social information. *Journal of Personality and Social Psychology*, *78*, 1024–1037.

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

Zajonc, R.B. (1968). Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology Monographs*, *9*(Pt. 2), 1–27.

(RECEIVED 8/6/01; REVISION ACCEPTED 3/13/02)