

Race and Gender on the Brain: Electrocortical Measures of Attention to the Race and Gender of Multiply Categorizable Individuals

Tiffany A. Ito and Geoffrey R. Urland
University of Colorado at Boulder

The degree to which perceivers automatically attend to and encode social category information was investigated. Event-related brain potentials were used to assess attentional and working-memory processes on-line as participants were presented with pictures of Black and White males and females. The authors found that attention was preferentially directed to Black targets very early in processing (by about 100 ms after stimulus onset) in both experiments. Attention to gender also emerged early but occurred about 50 ms later than attention to race. Later working-memory processes were sensitive to more complex relations between the group memberships of a target individual and the surrounding social context. These working-memory processes were sensitive to both the explicit categorization task participants were performing as well as more implicit, task-irrelevant categorization dimensions. Results are consistent with models suggesting that information about certain category dimensions is encoded relatively automatically.

Categorization is widely viewed as a sensible and efficient way to order and simplify a complex social world (Bodenhausen & Macrae, 1998; Brewer, 1988; Fiske & Neuberg, 1990; Macrae & Bodenhausen, 2000), but when applied to people, the process itself might not be so simple. This is because people can often be classified simultaneously along multiple dimensions. A complete understanding of social categorization must address how such multiply categorizable individuals are assigned to social group membership. Toward this end, some have suggested that all categories relevant to an impression-formation target are activated in parallel when multiply categorizable individuals are perceived (Bodenhausen & Macrae, 1998; Macrae & Bodenhausen, 2000; Macrae, Bodenhausen, & Milne, 1995). Others have suggested that this automatic activation occurs more narrowly for certain well-practiced dimensions such as race, gender, and age (Brewer, 1988; Bruner, 1957; Fiske & Neuberg, 1990; Stangor, Lynch, Duan, & Glass, 1992). Regardless of the specific categories afforded automatic status, both assumptions have important implications: To the extent that beliefs and affective reactions are associated with social categories, automatic activation of social category judgments would make this information readily available to influence subsequent processing stages.

Although much has been learned about the ease with which stereotypes and prejudice are activated, less is known about the

categorization process itself and the degree to which it occurs automatically. This may be attributable in part to the difficulty in measuring early attentional processes. It is not obvious, for instance, that perceivers would be able to accurately report when social category information becomes available and the relative importance of different dimensions at very early stages of processing. Consequently, when categorization has been studied, it has sometimes been assessed indirectly through outcomes such as stereotype activation (e.g., Macrae et al., 1995). Categorization has also been studied more directly by assessing the speed with which categorization judgments are made (Stroessner, 1996; Zarate, Bonilla, & Luevano, 1995; Zarate & Smith, 1990), but such measures are limited to measuring a single conscious outcome of the entire categorization process.

The present article investigates more directly how and when social category information becomes available when multiple dimensions of categorization are available, using the specific dimensions of race and gender. We focus on the contribution of two sources of attention: bottom-up and top-down factors. Assumptions that these dimensions are processed automatically imply that attention is directed to race and gender information early in processing on the basis of more bottom-up, stimulus-driven features. At the same time, recent findings on the malleability of automatic processes in person perception (for a review, see Blair, 2002) have suggested that more top-down processes such as processing goals and contextual features can also quickly influence categorization. The goal of this article is to investigate the time course and interplay of these two sources of attention using event-related brain potentials (ERPs).

The advantage of ERPs in this context is that they provide an on-line assessment of multiple information-processing operations as they unfold across time. That is, a single waveform can be recorded as participants perform a judgment about a specific stimulus. Multiple components of the waveform can then be examined to assess the degree to which different aspects of processing have occurred, with the time course of the compo-

Tiffany A. Ito and Geoffrey R. Urland, Department of Psychology, University of Colorado at Boulder.

This research was supported by National Institute of Mental Health Grant 1R03 MH61327 to Tiffany A. Ito and by a National Science Foundation Graduate Research Fellowship to Geoffrey R. Urland. We are grateful to Demis Glasford, Eve Jensen, and Erin Thompson for assistance with data collection and to the CU Stereotyping and Prejudice Lab for helpful comments on this article.

Correspondence concerning this article should be addressed to Tiffany A. Ito, Department of Psychology, University of Colorado, 345 UCB, Boulder, Colorado 80309-0345. E-mail: tito@psych.colorado.edu

nents providing information about the ordering of these operations.

To examine the degree to which race and gender information is automatically encoded, we selectively directed participants' attention to one of the dimensions as they viewed pictures of individuals varying on both dimensions. We then examined responses of two classes of ERP components. The first contains three components (N100, P200, N200) that are associated with early attentional effects. The amplitude of these early components increases as a function of the extent to which attention is directed at some feature of an external stimulus. Importantly, these attentional effects have varied both as a function of bottom-up attentional processes (Luck & Hillyard, 1994) as well as top-down processes such as task relevance (Hillyard & Munte, 1984; Luck & Hillyard, 1994; Ritter, Simson, & Vaughan, 1983; Wijers, Mulder, Okita, Mulder, & Scheffers, 1989). Thus, an examination of these components allowed us to assess (a) whether race and gender are attended to early in processing and (b) the degree to which such attention varies as a function of simple stimulus features or more complex factors such as task relevance and contextual influences.

The second class of ERP components we examined is associated with working-memory operations that occur somewhat later in processing. The particular component we examined, the P300, has been sensitive to task-relevant categorization decisions (Cacioppo, Crites, Berntson, & Coles, 1993; Donchin, 1981; Ito & Cacioppo, 2000; Ito, Larsen, Smith, & Cacioppo, 1998). Specifically, P300 amplitude increases to stimuli that differ from preceding stimuli along task-relevant dimensions. This sensitivity has led to the belief that P300 amplitude reflects updates to working memory that maintain an accurate mental model of the external environment (Donchin, 1981). On the basis of this, we expected the P300 to index working-memory effects of the dimension to which participants are explicitly attending. More importantly, however, we have shown that the P300 is also sensitive to task irrelevant categorization processes (Ito & Cacioppo, 2000), suggesting that the P300 also indexes working-memory processes that operate on a more implicit, bottom-up basis. This suggests that the P300 may also be used to assess spontaneous sensitivity to race and gender information by assessing reactions to task-irrelevant social category information. Consequently, when participants are instructed to attend to race, we can assess P300 responses as a function of target gender to index the degree to which gender is spontaneously incorporated in working memory. Similarly, when participants are instructed to attend to gender, we can assess P300 responses as a function of target race to index the degree to which race is spontaneously encoded. In the case of both classes of ERP components, we were also interested in comparing responses to race and gender cues to assess the relative influence of these two dimensions.

Experiment 1

Method

Participants and Overview

Thirty-six students (13 women) from the University of Colorado participated for partial fulfillment of class requirements. Thirty-one identified their ethnicity as White, 1 as Asian, 1 as Hispanic, 2 as other, and 1 declined to indicate an ethnicity.¹ All participants viewed pictures of Black

and White males and females, with some participants ($n = 21$) randomly assigned to categorize target individuals in terms of their race and others ($n = 15$) to categorize in terms of gender.

Materials

Yearbook pictures of 80 people who varied in both race (Black and White) and gender were selected on the basis of pilot testing in which ratings of attractiveness and judgments of race were obtained from an independent sample of participants ($N = 46$). Chosen stimuli were equated for attractiveness, and all had at least 80% agreement in race judgments. There were 20 pictures each in the categories of Black males, Black females, White males, and White females.

Procedure

Following procedures used in prior research on social perception (Ito, Thompson, & Cacioppo, 2002), pictures were shown to participants in sequences of five, with each picture presented for 1,000 ms. Participants were asked to think either about whether the person shown was male or female (gender categorization condition) or Black or White (race categorization condition). They registered their explicit categorization judgment after stimulus offset via an appropriately labeled keypad. After a 1,200-ms interstimulus interval, the next picture was shown. The word *pause* was shown on the screen after the fifth picture in the sequence until participants pressed a button to initiate presentation of the next sequence of pictures.

Picture sequences were arranged into four separate blocks, with the majority of the pictures coming from a single category within each block (e.g., White males). On some trials, all five pictures were drawn from this category. On other trials, one of the pictures came from one of the other categories (e.g., White female, Black male, or Black female).

ERPs were recorded during the presentation of a single *target* stimulus within each five-stimulus sequence. We refer to the nontarget stimuli as the *context*. Across the experiment, each category of pictures served as both the context and target an equal number of times, resulting in 16 possible types of trials. Specifically, target stimuli could be (a) the same race and gender as the context (e.g., a White male embedded in a White male context), (b) the same gender but different race as the context (e.g., a Black male embedded in a White male context), (c) the same race but different gender as the context (e.g., a White female embedded in a White male context), or (d) both a different race and gender than the context (e.g., a Black female embedded in a White male context). To prevent participants from anticipating when individuals of various races or genders would be shown, the target stimulus randomly appeared in either the third, fourth, or fifth position in a sequence. Presenting the stimuli in this manner allowed us to look at both simple target group membership effects as well as more complex relations between the target and preceding individuals. Participants viewed 320 total trials (20 of each trial type). This required each unique picture to be shown 20 times.

Psychophysiological Data Collection and Reduction

ERP data were recorded at sites over midline frontal (Fz), central (Cz), and parietal (Pz) areas using tin electrodes sewn into an elastic cap (Electro-Cap International, Eaton, OH). Miniature tin electrodes were also placed over the left and right mastoid. Active scalp sites were referenced on-line to the left mastoid. Additional miniature tin electrodes were placed above and below the left eye, and on the outer canthus of each eye to monitor vertical and horizontal eye movements, respectively. Electrode impedances were below 5 K Ω at all sites. ERP recordings were amplified

¹In both experiments, analyses conducted only on the participants reporting their race as White were identical to those in the main text.

with a gain of 500 by NeuroScan Synamps (Compumedics USA, El Paso, TX) model amplifiers with a bandpass of .1–30 Hz (12-dB roll-off) and digitized at 1,000 Hz. Data recording began 128 ms before picture onset and continued throughout the 1,000-ms picture presentation.

Off-line, the data were rereferenced to a computed average of the left and right mastoids and submitted to a regression procedure to remove the effects of vertical eye movements from the ERP, which can distort measurements from scalp sites (Semlitsch, Anderer, Schuster, & Presslich, 1986), then corrected to the mean voltage of the 128-ms prestimulus recording period. We next visually inspected the ERP data and deleted any trials with remaining ocular or other artifact (e.g., because of movement). Data from all sites for that trial were eliminated from further analysis if artifact was detected at any of the scalp sites.

Two sets of ensemble averages were constructed by aggregating for each participant at each scalp site the electrical activity associated with each of the 16 trial types, with the averages aligned on stimulus onset. One set of ensemble averages, calculated without further filtering of the data, was used to evaluate the shorter latency, faster frequency components, whereas another set of average waveforms constructed on data that were first filtered with a 9-Hz low-pass filter was used to evaluate the longer latency, slower frequency P300 component. The use of different filtering parameters was chosen as the best way to extract different types of signal (i.e., short- vs. long-latency components) from the background ERP noise (cf. Gehring, Gratton, Coles, & Donchin, 1992; Scheffers & Coles, 2000).

Visual inspection of the resulting waveforms revealed four distinct deflections: a negative-going deflection with a mean amplitude of 122 ms, a positive-going deflection with a mean amplitude of 176 ms, a negative-going deflection with a mean amplitude of 256 ms, and a positive-going deflection with a mean amplitude of 485 ms. We refer to these components on the basis of their polarity and latency as the N100, P200, N200, and P300,² respectively. The amplitude of each of the four components was quantified within each of the 16 condition averages constructed for each participant at each electrode. This was done on the 30-Hz filtered averages for the N100, P200, and N200 and the 9-Hz filtered averages for the P300. Amplitude was quantified by locating within each of the 16 condition averages constructed for each participant (a) the largest negative-going potential at Cz between 50 ms and 150 ms after stimulus onset, (b) the largest positive-going potential at Pz between 150 ms and 250 ms after stimulus onset, (c) the largest negative-going potential at Cz between 200 ms and 350 ms after stimulus onset, and (d) the largest positive-going potential at Pz between 350 ms and 900 ms after stimulus onset. The amplitude of each component at the other (nondefinitional) sites was scored as the largest potential with the appropriate polarity occurring within ± 100 ms of the component located in the definition electrode (cf. Cacioppo et al., 1993; Ito & Cacioppo, 2000; Ito et al., 1998, 2002).

Results and Discussion

Preliminary analyses involving sagittal scalp site showed that the N100 and N200 were maximal at the Cz electrode, the P300 was maximal at the Pz electrode, and the P200 was equally large at Pz and Cz. For ease of presentation, we focus our analyses of each component on the electrode at which activity was maximal. In the case of the P200, we report analyses on the Cz electrode. Each component was analyzed with separate 2 (categorization task: race or gender) \times 2 (participant gender: male or female) \times 2 (target race: Black or White) \times 2 (target gender: male or female) \times 2 (status of target on task-relevant dimension: same as context or different than context) \times 2 (status of target on task-irrelevant dimension: same as context or different than context) mixed-model analyses of variance. All factors except categorization task and participant gender were within subject. In no case did participant

gender alter the pattern of results. This factor is therefore not considered further.

A note about inferences from ERP signals is warranted. The amplitude is thought to reflect the degree to which a particular information-processing operation is engaged. We therefore used the amplitude of the components to infer the degree to which different types of attentional processes and working-memory operations were indicated by various components. The latency of a component's peak is thought to reflect time taken to perform the corresponding information-processing operation. Mean peak latencies are reported to provide an indication of when in time the process associated with a given component was complete.

N100

Given past research indicating that early components such as the N100, P200, and N200 are associated with early selective attention effects, we examined these components to assess the influence of simple stimulus cues indicating race and gender group membership and more complex factors such as task instructions and contextual factors in early processing stages. Analyses on the first distinct component, which peaked at approximately 122 ms, revealed sensitivity to race information in the form of a significant target race main effect, $F(1, 32) = 40.22, p < .0001$. N100s were larger to Blacks ($M = -7.39 \mu\text{V}$) than to Whites ($M = -5.92 \mu\text{V}$; see Figure 1, Panel A). No other effects were significant in the N100 analysis.

P200

There were both main effects of target race, $F(1, 32) = 17.35, p < .001$, and target gender, $F(1, 32) = 6.10, p < .05$, in the P200, which peaked at approximately 176 ms. Larger P200s were elicited by Blacks ($M = 8.87 \mu\text{V}$) than Whites ($M = 7.96 \mu\text{V}$; Figure 1, Panel A) and by males ($M = 8.59 \mu\text{V}$) than females ($M = 7.98 \mu\text{V}$; Figure 1, Panel B).

N200

There was a target gender main effect, $F(1, 32) = 11.32, p < .01$, with larger N200s to female ($M = -4.37 \mu\text{V}$) than male ($M = -3.50 \mu\text{V}$) targets (Figure 1, Panel B).³

Summary of Early Attentional Effects

There was clear evidence that attention is covertly directed to race and gender cues at early processing stages. Race effects were

² Although the P300 had a latency longer than 300 ms, we use the P300 name because its scalp distribution and response to psychological processes mirrors the classic P300 component.

³ There was also a Task \times Status of Target on Task-Relevant Dimension interaction for the N200, $F(1, 32) = 5.42, p < .05$. This interaction appears to be attributable to the target's status on the attended dimension having opposite effects in the two task conditions. For participants attending to gender, a difference in target gender relative to the context increased N200 amplitude ($M = 3.21 \mu\text{V}$ vs. $M = 2.81 \mu\text{V}$). By contrast, for participants attending to race, a difference in target race relative to context decreased N200 amplitude ($M = 4.29 \mu\text{V}$ vs. $M = 4.99 \mu\text{V}$). In neither task condition were the simple effects of target status on task-relevant dimension significant. However, the size of the same-as-context/different-from-context comparison between the two tasks was marginally significant, $F(1, 34) = 3.33, p = .08$.

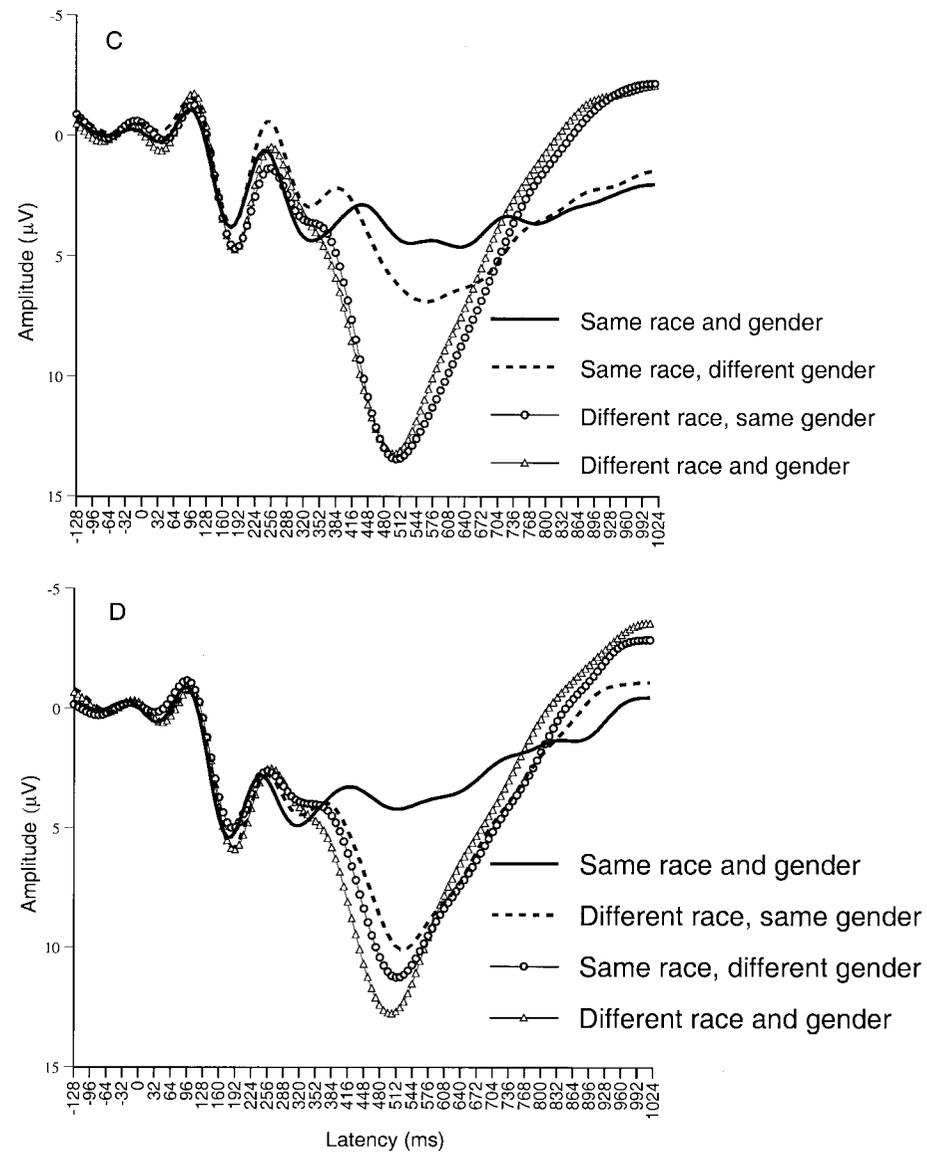
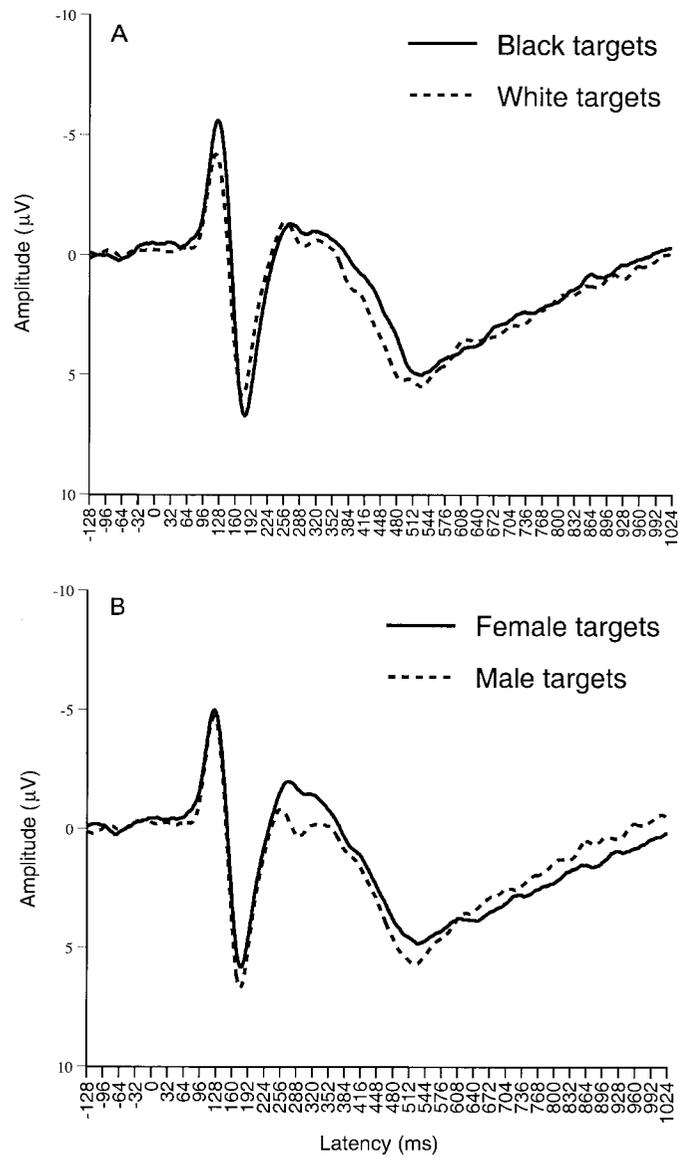


Figure 1. Early and late component effects from Experiment 1. Panels A and B show waveforms at the central area (Cz) collapsed across all factors except target race (Panel A) or target gender (Panel B). Data in Panels A and B were processed with a 30-Hz low-pass filter. Panels C and D show waveforms at the parietal area (Pz) as a function of target status on task-relevant and task-irrelevant dimensions for participants attending to race (Panel C) and gender (Panel D). Data in Panels C and D were processed with a 9-Hz low-pass filter.

the first to emerge, with larger N100s and P200s to Black than to White faces. Gender effects emerged in the P200, with larger responses to males than to females. This reversed in the N200, where responses were larger to females. Interestingly, no systematic effects of categorization task or the target's relation to the preceding context stimuli emerged in these earlier components. This indicates that at these earlier stages of visual processing, effects were primarily attributable to properties of the specific stimulus, regardless of the task being performed or the more complex relation between the present and preceding stimuli. On the basis of past research, however, we expect the P300 to be sensitive to these latter effects.

P300

Explicit categorization effects. Recall that in past P300 research, P300 amplitude has increased when a stimulus differs from the preceding context along task-relevant dimensions, reflecting context-updating processes that occur in response to explicit attention to task-relevant stimulus features. Consistent with this, P300s were larger when a target individual's social category membership differed from the individuals pictured in the preceding context on the task-relevant dimension ($M = 14.98 \mu\text{V}$) as compared with when it matched the preceding context ($M = 8.56 \mu\text{V}$), $F(1, 32) = 199.41, p < .0001$. Note that these means collapse across the target's status on the task-irrelevant dimension.

These explicit (i.e., task-relevant) context-updating effects were also qualified by type of task, $F(1, 32) = 10.20, p < .01$. Although participants attending to race and gender both showed the explicit context-updating effects, the effect was larger for participants attending to race, $F(1, 34) = 9.23, p < .01$ (see top panel of Table 1).⁴

Implicit categorization effects. Prior P300 research has often examined categorization of nonsocial stimuli, such as tones of different pitches or patterns of different colors. The effects just reported involving target status on the task-relevant dimension therefore extend prior P300 research to show that the processes it indexes are sensitive to social category judgments. Of greater theoretical relevance to the issue of automatic attention to social

category information is the P300's sensitivity to a target individual's status on the social dimension that is not task relevant. A significant main effect of target status on the task-irrelevant dimension was observed, $F(1, 32) = 26.18, p < .0001$, as was the interaction with task, $F(1, 32) = 6.84, p < .01$. The means in the bottom panel of Table 1 show that P300s were increased when a target differed from the context along the task-irrelevant dimension, although this difference was only marginally significant for participants performing the race categorization task (i.e., the implicit race effect was significant whereas the implicit gender effect was marginal). Note that these effects collapse across the target's status on the task-relevant dimension.

Relative magnitude of explicit and implicit categorization of race and gender. The preceding analyses indicate that (a) P300 is sensitive to explicit social categorization decisions, (b) P300 is sensitive to implicit social categorization decisions, and (c) racial group membership has more impact than gender group membership. The latter is supported by the larger explicit categorization effects for participants performing the race categorization task and the larger implicit categorization effects for participants performing the gender categorization task. All three of these effects are also supported by a significant Target Status on Task-Relevant Dimension \times Target Status on Task-Irrelevant Dimension interaction, $F(1, 32) = 25.58, p < .0001$, and the three-way interaction involving task, $F(1, 32) = 8.05, p < .01$.

The waveforms in Panels C and D of Figure 1 show P300 responses from participants performing the race and gender tasks, respectively. The smallest P300s in both task conditions were elicited by targets whose race and gender were identical to the preceding context pictures. A significantly larger P300 was elicited by targets who differed from the context on only the task-irrelevant condition, both for participants attending to race (in which case this comparison involves targets who were the same as the preceding context in terms of race but different in terms of gender), $F(1, 20) = 12.29, p < .005$, and for participants attending to gender (in which case this comparison involves targets who were the same as the preceding context in terms of gender but different in terms of race), $F(1, 14) = 100.84, p < .001$. These comparisons provide the strongest evidence that participants were modifying their mental models of the environment in response to social category membership information that was not directly task relevant.

Although participants were clearly demonstrating implicit categorization effects, the explicit categorization effects were of a larger magnitude. When attending to race, a target who differed from the preceding context in terms of only race attracted greater processing than someone who differed only in terms of gender, $F(1, 20) = 174.97, p < .0001$. Similarly, for participants attending to gender, a target who differed from the preceding context in

Table 1
Mean P300 Amplitude as a Function of Task and Target Status on Task-Relevant and Task-Irrelevant Dimensions

Task	Same as context		Different from context	
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
Target status on task-relevant dimension				
Race categorization	8.48 _a	3.31	15.80 _b	3.97
Gender categorization	8.85 _a	3.09	13.53 _b	3.81
Target status on task-irrelevant dimension				
Race categorization	11.76 _a	3.47	12.52 _b *	3.51
Gender categorization	9.95 _a	3.45	12.43 _b	3.38

Note. Means within the same row with different subscripts differ from each other at $p < .0001$.

* $p = .067$.

⁴ For the P300, the effects of target status on the task-relevant dimension were also moderated by target gender, $F(1, 32) = 7.03, p < .05$. This effect revealed that although larger responses to targets who did versus did not differ along the task-relevant dimension were present for both male and female targets, the effect was larger for male targets. The mean difference in P300 amplitude between targets who represented a change on task-relevant dimensions versus those who did not was $M = 6.94 \mu\text{V}$ for male targets but only $M = 5.50 \mu\text{V}$ for female targets.

terms of only gender attracted more processing than someone who differed only in terms of race, $F(1, 14) = 8.04, p < .05$.

We also tested whether a stimulus indicating the need to update one's mental model in terms of both race and gender information would be associated with the largest P300s. In neither task condition did a target differing from the context in both race and gender elicit P300s that differed significantly from those elicited by targets who differed along only the task-relevant dimension ($F_s < 1$ in both task conditions).

Finally, consistent with earlier indications that racial category information has a larger impact on working-memory representations than gender category information, Panels C and D of Figure 1 show that the implicit categorization effect attributable to target race was marginally larger than that attributable to target gender, $F(1, 34) = 3.94, p = .055$. No other comparisons between tasks were significant.

Other effects. The only other significant effect in this analysis was a target race main effect in which Whites elicited larger P300s ($M = 12.22 \mu V$) than Blacks ($M = 11.31 \mu V$), $F(1, 32) = 6.81, p < .05$. This effect is at odds with the larger N100s and P200s to Blacks and suggests that differential processing to a particular racial group may change as processing unfolds. This is addressed in greater detail in the General Discussion section.

Summary

Analyses on the earlier components indicate the quickness with which race and gender are encoded. The N100 and P200 analyses indicate that more attention is directed to Blacks than to Whites by approximately 122 ms and is sustained through at least 176 ms. Conversely, working-memory effects of task relevance and the more complex relation between the group memberships of the individual being viewed and preceding (context) individuals do not emerge until much later (around 483 ms). At this point, working-memory models were clearly influenced by the social category information to which participants were explicitly attending as well as social category information that was not explicitly task relevant. Finally, there were several asymmetries indicating greater attention to racial than gender cues. Race effects emerged earlier in time than gender effects, and in the P300, working-memory updating effects were larger as a function of race than of gender.⁵

Experiment 2

The effects of race and gender on early attentional processes and later working-memory representations are clearly consistent with the assumption that social perceivers automatically attend to these dimensions. We had not, however, expected the consistent bias in favor of racial cues. Although greater sensitivity to race than to gender has been obtained in some contexts (Stroessner, 1996), there are other instances in which sensitivity to gender has dominated race (Gardner, MacIntyre, & Lalonde, 1995; Kurzban, Tooby, & Cosmides, 2001; Stangor et al., 1992; Zarate et al., 1995; Zarate & Smith, 1990). Moreover, gender is likely to be a more potent cue biologically than race (Kurzban et al., 2001). We were therefore interested in better understanding the source of the asymmetry favoring race in Experiment 1. Differences in visual salience for racial versus gender information are one possible source. Skin tone differences between our Black and White targets may have

created larger luminance contrast effects than existed between male and female targets. To assess the degree to which these kinds of physical differences were responsible for the findings in Experiment 1, Experiment 2 used an identical design except that the color photos were converted to grayscale and equated for brightness and contrast across target race and gender. Use of grayscale images also allowed us to assess the extent to which early attention to these cues occurs even when the stimuli are rendered less potent.

Method

Participants

Thirty-six students (21 women) from the University of Colorado participated for partial fulfillment of class requirements. Thirty-two identified their ethnicity as White, 2 as Asian, and 2 as Hispanic. Half of the participants were randomly assigned to perform the racial categorization task and the other half the gender categorization task. One participant who performed the gender categorization task had high impedances at the Cz electrode and was dropped from the non-P300 analyses.

Materials

Pictures were converted to grayscale and equated for luminance and contrast using Adobe Photoshop 7.0. First, all color information was discarded by converting the picture from red-green-blue (RGB) to grayscale mode. Next, we created groupings of four pictures representing one picture from each category that were matched on attractiveness. Within each group, pictures were adjusted until they had equivalent values for brightness, contrast, and color intensity as provided by Photoshop's histogram feature, which provides the total pixels at each intensity level on a black-to-white continuum.

Procedure and Psychophysiological Data Collection and Reduction

All procedures and aspects of data collection and analysis were identical to Experiment 1. The N100, P200, and N200 latencies were very similar to Experiment 1 (121 ms, 180 ms, and 264 ms, respectively), but the P300 latency here was delayed (546 ms as compared with 485 ms in Experiment 1). This is consistent with prior research indicating that stimulus degradation slows processes indexed by the P300 (McCarthy & Donchin, 1981; Smulders, Kenemans, Schmidt, & Kok, 1999). As in Experiment 1, participant gender did not alter the pattern of results and is not considered further.

Results and Discussion

N100

Despite the use of grayscale stimuli, the N100 showed sensitivity to target race in the form of a main effect of target race, $F(1,$

⁵ The larger effects for race than for gender might alternatively be interpreted as indicating that race is more difficult to process. However, factors such as increases in task difficulty, perceptual load, or stimulus degradation typically have their effects by increasing component latency or decreasing amplitude (Caryl & Harper, 1996; Donchin, 1981; McCarthy & Donchin, 1981; Smulders et al., 1999). Thus, larger P300 amplitudes as a function of race are more consistent with the interpretation that race has a larger impact on working memory. Moreover, supplemental latency analyses in both experiments revealed no systematic asymmetries to race versus to gender.

31) = 10.29, $p < .005$. As in Experiment 1, N100s were larger to Blacks ($M = -6.64 \mu\text{V}$) than to Whites ($M = -5.91 \mu\text{V}$; see Figure 2, Panel A). A significant Target Ethnicity \times Target Gender \times Task interaction, $F(1, 31) = 4.85$, $p < .05$, indicated larger responses to Blacks than to Whites in all conditions except when female targets were viewed by participants performing the gender categorization task.

P200

As with the N100 and Experiment 1, there was a significant main effect of target race indicating that P200s were larger to Blacks ($M = 8.48 \mu\text{V}$) than to Whites ($M = 7.31 \mu\text{V}$), $F(1, 31) = 22.41$, $p < .001$. Also replicating Experiment 1, the gender main effect was significant, $F(1, 21) = 5.97$, $p < .05$, with larger P200s to males ($M = 8.25 \mu\text{V}$) than to females ($M = 7.54 \mu\text{V}$; Figure 2, Panel B). Both were qualified by the Target Race \times Target Gender interaction, $F(1, 31) = 7.31$, $p < .05$, which revealed larger P200s to Black males ($M = 9.13 \mu\text{V}$) than to the other three types of targets ($M = 7.48 \mu\text{V}$), $F(1, 34) = 30.10$, $p < .001$.⁶

N200

N200 analyses revealed a main effect of target race, $F(1, 31) = 6.44$, $p < .05$, in which White targets elicited larger N200s ($M = -4.16 \mu\text{V}$) than Black targets ($M = -3.34 \mu\text{V}$). The target gender main effect from Experiment 1 was marginally significant here, $F(1, 31) = 3.60$, $p = .067$, with larger N200s to females ($M = -4.07 \mu\text{V}$) than to males ($M = -3.44 \mu\text{V}$). There was also a Target Race \times Target Gender interaction, $F(1, 31) = 15.62$, $p < .001$. White males were associated with larger N200s ($M = -4.38 \mu\text{V}$) than Black males ($M = -2.51 \mu\text{V}$), $F(1, 31) = 18.39$, $p < .001$, but responses to Black females ($M = -4.18 \mu\text{V}$) and White females ($M = -3.95 \mu\text{V}$) did not differ.

P300

Explicit and implicit categorization effects. Explicit and implicit categorization effects were again seen. The main effect of target status on the task-relevant dimension showed that P300s were larger to target individuals whose social category membership differed from versus were the same as the individuals pictured in the preceding context on the task-relevant dimension ($M_s = 12.96 \mu\text{V}$ and $7.27 \mu\text{V}$), $F(1, 32) = 99.89$, $p < .0001$. The main effect of target status on the task-irrelevant dimension revealed that P300 responses were larger when a target individual differed from versus was the same as the preceding context individuals along the task-irrelevant dimension ($M_s = 11.38 \mu\text{V}$ and $8.86 \mu\text{V}$), $F(1, 32) = 36.06$, $p < .0001$. Unlike Experiment 1, neither of these main effects was moderated by task.

Relative magnitude of explicit and implicit categorization of race and gender. There was also an interaction between target status on the task-relevant and task-irrelevant dimensions, $F(1, 32) = 48.50$, $p < .0001$. Although this effect was not qualified by task, as it was in Experiment 1, the waveforms in Figure 2 display the data separately for the race (Panel C) and gender (Panel D) categorization conditions for comparison purposes. As can be seen, the smallest P300s were elicited by targets whose race and

gender were identical to the preceding context pictures. A significantly larger P300 was elicited by targets who differed from the context on only the task-irrelevant dimension, indicating implicit context-updating effects, $F(1, 35) = 87.83$, $p < .0001$. Explicit context-updating effects were in turn larger, $F(1, 35) = 44.22$, $p < .0001$. In this experiment, responses to individuals who differed from the context along both social dimensions elicited significantly larger P300s than targets who differed along only the task-relevant dimension, $F(1, 35) = 8.17$, $p < .01$.^{7,8}

Summary

Despite the change from full-color to grayscale images, Experiment 2 largely replicated Experiment 1: (a) N100s were larger to Blacks than to Whites, (b) P200s were larger to Blacks than to Whites, (c) P200s were larger to males than to females, and (d) N200s were larger to females than to males. Although attention still appeared to be asymmetrically oriented to race in early processing stages, the bias in favor of race that was seen in the P300 in Experiment 1 was not obtained here. Updates to working memory therefore appear to not be differentially affected by race and gender cues when grayscale images are used.

General Discussion

The two experiments reported here provide relatively direct evidence that when perceivers encounter multiply categorizable individuals, race and gender information are both activated at very early stages in processing. Across both experiments, sensitivity to race information was seen as early as 122 ms in the N100. Clear sensitivity to gender information emerged slightly later in the P200. Amplitude modulation in these early components has been associated with the covert orienting of attention in response to stimuli that are inherently attention grabbing (Luck & Hillyard, 1994). Although the simple effects of group membership were apparent in the early components, there was also evidence that perceivers became aware of more complex relations between the

⁶ Other significant P200 effects included the Target Status on Task-Relevant Dimension \times Task interaction, $F(1, 31) = 5.81$, $p < .05$. When participants were attending to race, targets whose race matched the preceding context elicited smaller P200s than those differing in race from the context, $F(1, 31) = 4.54$, $p < .05$. For participants performing the gender task, a target's relation to the gender of the context did not affect the P200. There was also a Target Status on Task-Irrelevant Dimension \times Target Gender interaction, $F(1, 31) = 8.46$, $p < .01$, such that larger responses to males than to females occurred when the target differed from the context along the task-irrelevant dimension but not when the target and context matched on the task-irrelevant dimension.

⁷ Even though these effects were not moderated by task, we performed the same comparison as in Experiment 1 to compare the size of the implicit categorization effect between the two task conditions. Implicit responses to race among participants attending to gender did not significantly exceed implicit responses to gender among participants attending to race ($F < 1$).

⁸ The Target Status on Task-Relevant Dimension \times Target Gender interaction was also significant in the P300 analysis, $F(1, 32) = 4.79$, $p < .05$. Responses were larger to all targets that differed from the context along the task-relevant dimension (consistent with the main effect reported in the text), but this difference was larger for pictures of males than of females.

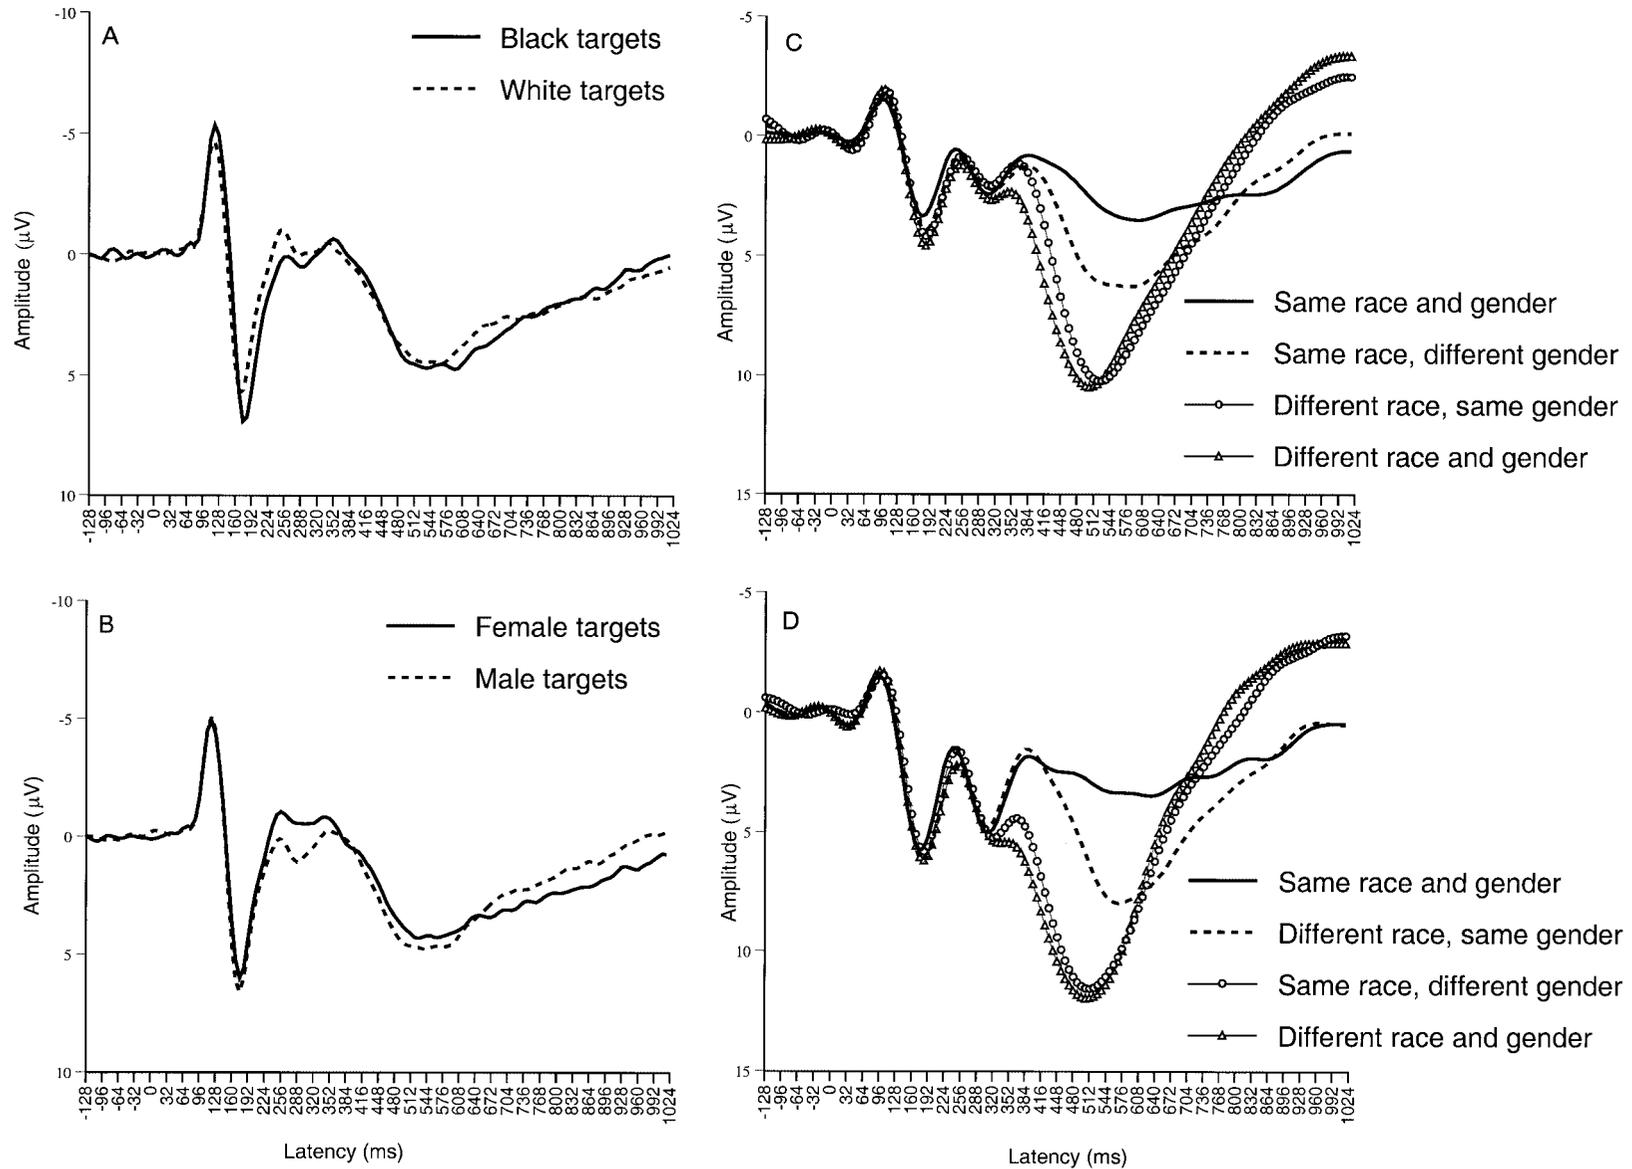


Figure 2. Early and late component effects from Experiment 2. Panels A and B show waveforms at the central area (Cz) collapsed across all factors except target race (Panel A) or target gender (Panel B). Data in Panels A and B were processed with a 30-Hz low-pass filter. Panels C and D show waveforms at the parietal area (Pz) as a function of target status on task-relevant and task-irrelevant dimensions for participants attending to race (Panel C) and gender (Panel D). Data in Panels C and D were processed with a 9-Hz low-pass filter.

social category memberships of a target individual and the individuals who preceded him or her in slightly later processing stages. Not surprisingly, participants were explicitly sensitive to these relations, as seen in the modulation of P300 amplitude as a function of a target's relation to its preceding context along the social category dimension to which participants were instructed to attend. This finding is consistent with past findings of P300 modulation by task-relevant dimensions. More importantly, participants in both experiments were also implicitly aware of these relations, as indicated in the modulation of P300 amplitude as a function of the relation between a target and its preceding context along the task-irrelevant social category dimension. Participants' sensitivity to context information in the P300 suggests that the covert orienting of attention to race and gender information (as shown in the N100, P200, and N200) was used to build a more complex representation that incorporated not only the target's race and gender but also the group membership of preceding individuals.

The replication of the early race effects in Experiment 2 argues against the possibility that our earlier race than gender effects were due to luminance and contrast effects creating a greater degree of physical differentiation for race than gender. Nevertheless, the color status of the photos did influence the working-memory effects. In Experiment 1, the working-memory processes indexed by the P300 were clearly larger in response to race than gender information; explicit categorization effects were larger for participants attending to race than to gender, and implicit categorization effects were larger for those attending to gender than to race. This asymmetry was not obtained in Experiment 2. Together with the results from the earlier components, this suggests that race is attended to earlier than gender even when the visual salience between racial categories in the form of luminance and contrast is equated, but more complex relations between the racial group membership of individuals in the environment have a processing advantage over comparable gender information only when people are perceived in their more naturalistic condition (i.e., in full color).

Early attentional effects for race were always in the direction of more attention—larger responses to the less normative racial group. In the P300 for Experiment 1 and N200 for Experiment 2, the pattern reversed, with greater attention to Whites than to Blacks. Similarly larger responses to Whites than to Blacks in a negative-going component that peaked at about 280 ms were obtained in an earlier study (Ito et al., 2002). This occurred for White participants who were making evaluative judgments of individuals and was interpreted as possibly reflecting the spontaneous direction of deeper processing to in-group or more normative group members, consistent with other research on in-group encoding biases (Chance & Goldstein, 1981; Levin, 2000; Malpass & Kravitz, 1969). Thus, attention may be initially directed at the out-group or less normative racial group, but is eventually preferentially directed to the in-group or more normative group to facilitate the biases in encoding that have been observed in other contexts. Note that the absence of a sizable number of Blacks in any of these samples prevents us from knowing whether our results are a function of the target's group membership (White–Black), or the participants' relation to the target (in-group/out-group).

Implications for Models of Face Processing

Although these experiments were designed primarily to examine social categorization, the use of face stimuli makes them relevant to understanding how faces are processed more generally. ERPs have been used repeatedly to examine the structural encoding of faces and the retrieval of personal identity information. Structural encoding of face information occurs around 170 ms after a face is encountered, and the retrieval of personal identity occurs later, around 350 ms. The different latencies of these ERP components have supported theories arguing that the structural analysis of the face must precede all other aspects of face processing, including social categorization (Bruce & Young, 1986). Despite the inclusion of social categorization in these models, little empirical research has examined it. The present results therefore address a previously understudied area.

More importantly, though, the N100 race effects we obtained occurred earlier than when the structural face-processing effect typically occurs. Methodological differences are important to note. All of our stimuli were faces, whereas faces are typically shown only occasionally in the face-processing literature. This suggests two interesting possibilities. At the very least, the firm view that face structural encoding cannot be completed until 170 ms may be situation specific. That our participants responded to face race at approximately 120 ms may indicate that some aspects of face perception can be facilitated with certain tasks or contexts. Alternatively, our results may indicate that social categorization need not depend on the complete structural analysis of the face (see also Mouchetant-Rostaing, Giard, Bentin, Aguera, & Pernier, 2000). This possibility is supported by the fact that features outside of the face can no doubt contribute to social categorization. Features such as skin tone or hair color–style may be modulating perception before the point at which full facial analysis is complete. Thus, our results suggest a revision to models of face processing to, at the minimum, allow for structural analysis of the face to be completed before 170 ms in some cases or to allow for the possibility that other aspects of face processing can occur before structural analysis is complete. Further examination of these types of issues should help to clarify both how faces are processed and how social category information is processed.

Relative Attention to Race and Gender

Our findings of earlier and larger race effects are at odds with some prior research in which gender categorization effects have been larger (Gardner et al., 1995; Stangor et al., 1992; Zarate et al., 1995; Zarate & Smith, 1990, but see Stroessner, 1996). Our effects could be specific to the stimuli we used, but we think that is unlikely given the use of pictures of 80 different individuals. Also, if anything, target gender was less ambiguous than target race. Virtually no pilot test participants found the gender of the targets ambiguous, but targets only had to be classified into their racial category by 80% of pilot test participants in order to be selected. Moreover, if anything, we might expect the occasion of a yearbook photo to heighten gender salience by directing the attention of the person being photographed to things like hair and makeup. Nevertheless, although our results speak to what we think occurs when people are perceived in a form in which they frequently present themselves, we also acknowledge that situational manipulations

might speed the effects of gender (e.g., if gender salience were increased by showing men in football and women in cheerleading uniforms or if an individual had a processing goal of looking for a potential date).

Prior instances of larger gender than race effects have been obtained using heterogeneous measures of categorization. Stangor et al. (1992) assessed categorization by comparing memory confusions across target gender and racial lines in the who-said-what paradigm. Gardner et al. (1995) assessed both stereotypicality judgments and memory errors. Zarate et al. (1995) and Zarate and Smith (1990) used a category verification task in which a picture of a target individual was followed by a category label probe (e.g., *Black*, *White*, *male*, or *female*). Categorization was measured as the speed with which participants could verify the racial or gender group membership of the target. In Zarate and Smith, verification was faster for gender group membership than for racial group membership. In Zarate et al., similarly larger effects for gender than for race were obtained, but only for White participants. For Hispanic participants, effects were larger for race than for gender. In addition, in Zarate et al., the effects for the White participants depended on the stimuli having been processed in a prior categorization task.

Of these heterogeneous measures, some clearly focus on more complex social judgments, such as ratings of stereotypicality and recall for statements made by specific targets. This may suggest that race has larger effects in terms of simple attention and early working-memory effects but that perceivers find gender a more meaningful distinction for more thoughtful and complex social judgments. It is also the case that our participants probably had much greater experience with gender than with racial diversity. Relative lack of racial diversity may have made the racial out-group targets in the experiments particularly salient, thereby directing more attention to race than to gender. However, one could also argue that greater practice with gender categorization should facilitate attention to gender. Thus, the effect that greater gender than racial diversity would have on attention is not clear, but the difference in our early processing effects and other findings based on more complex judgments warrant further research.

Our results are also relevant to perspectives that predict differential attention to specific social groups. According to the White male norm framework, male gender group membership and White racial group membership are perceived in American culture as more normative "default values." Deviations from these normative social groups draw attention (Smith & Zarate, 1992; Stroessner, 1996). The larger N100s and P200s to Blacks than to Whites are consistent with this and also fit with prior ERP research in which faces made less prototypical through digital editing elicited larger responses in early components (Halit, de Haan, & Johnson, 2000). Interestingly, though, the White male norm predicts greater attention to Black females, who differ from the norm along both race and gender dimensions, which was not obtained here. Instead, P200s were the largest to Black males in Experiment 2, and although the Target Race \times Target Gender interaction was not significant in Experiment 1, a post hoc contrast revealed that P200s were significantly larger to Black males ($M = 9.42 \mu\text{V}$) than to all other targets ($M = 8.05 \mu\text{V}$), $F(1, 35) = 8.08$, $p < .01$. Our results are more consistent with Eagly and Kite's (1987) analysis that the higher status of men results in their being disproportionately viewed as protagonists (see also Sidanius & Pratto, 1999). This in

turn makes the attributes of a social group more likely to be ascribed to male than to female members. To the extent that stereotypes about Blacks include aspects of aggressiveness and violence, Black males should be perceived as the most potentially threatening among the targets shown. Thus, the combination of greater perceived agency for men and the more threatening stereotypes about Blacks may be responsible for our effects.

Conclusion

Recent research on the ease with which evaluations and cognitions associated with social groups can be activated has focused attention on how the activation and application of this information can be modified. The present results help to further characterize this process by providing information about how quickly social category membership information becomes available. They also raise some interesting issues for future research. First, we do not yet know whether the early attentional effects we addressed can be modified. For instance, it would be interesting to know if manipulations that have been shown to affect the activation and application of prejudice and stereotypes do so only by operating on relevant evaluations and cognitions or whether these manipulations operate earlier at the stage of social categorization. A second issue that could be addressed is a more precise mapping of the four ERP components that were associated with processing social category information. Social categorization may be composed of distinct component processes. Identifying these components and their mapping onto these ERP components will further illuminate the process by which people order their social world.

References

- Blair, I. V. (2002). The malleability of automatic stereotypes and prejudice. *Personality and Social Psychology Review*, 6, 242–261.
- Bodenhausen, G. V., & Macrae, C. N. (1998). Stereotype activation and inhibition. In R. S. Wyer Jr. (Ed.), *Stereotype activation and inhibition* (pp. 1–52). Mahwah, NJ: Erlbaum.
- Brewer, M. C. (1988). A dual process model of impression formation. In R. Wyer & T. Srull (Eds.), *Advances in social cognition* (Vol. 1, pp. 1–36). Hillsdale, NJ: Erlbaum.
- Bruce, V., & Young, A. W. (1986). A theoretical perspective for understanding face recognition. *British Journal of Psychology*, 77, 305–327.
- Bruner, J. S. (1957). On perceptual readiness. *Psychological Review*, 64, 123–151.
- 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.
- Caryl, P. G., & Harper, A. (1996). Event related potentials (ERPs) in elementary cognitive tasks reflect task difficulty and task threshold. *Intelligence*, 22, 1–22.
- Chance, J. E., & Goldstein, A. G. (1981). Depth of processing in response to own- and other-race faces. *Personality and Social Psychology Bulletin*, 7, 475–480.
- Donchin, E. (1981). Surprise! . . . Surprise? *Psychophysiology*, 18, 493–513.
- Eagly, A. H., & Kite, M. E. (1987). Are stereotypes of nationalities applied to both men and women? *Journal of Personality and Social Psychology*, 53, 451–462.
- Fiske, S. T., & Neuberg, S. L. (1990). A continuum of impression formation, from category-based to individuating processes: Influences of in-

- formation and motivation on attention and interpretation. *Advances in Experimental Social Psychology*, 23, 1–73.
- Gardner, R. C., MacIntyre, P., & Lalonde, R. (1995). The effects of multiple social categories on stereotyping. *Canadian Journal of Behavioural Science*, 27, 466–483.
- Gehring, W. J., Gratton, G., Coles, M. G. H., & Donchin, E. (1992). Probability effects on stimulus evaluation and response processes. *Journal of Experimental Psychology: Human Perception and Performance*, 18, 198–216.
- Halit, H., de Haan, M., & Johnson, M. H. (2000). Modulation of event-related potentials by prototypical and atypical faces. *NeuroReport*, 11, 1871–1875.
- Hillyard, S. A., & Munte, T. F. (1984). Selective attention to color and location: An analysis with event-related brain potentials. *Perception and Psychophysics*, 36, 185–198.
- Ito, T. A., & Cacioppo, J. T. (2000). Electrophysiological evidence of implicit and explicit categorization processes. *Journal of Experimental Social Psychology*, 36, 660–676.
- 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.
- Ito, T. A., Thompson, E., & Cacioppo, J. T. (2002). *Reactions to ingroup and outgroup faces: Tracking component processes using event-related brain potentials*. Manuscript submitted for publication.
- Kurzban, R., Tooby, J., & Cosmides, L. (2001). Can race be erased? Coalitional computation and social categorization. *Proceedings of the National Academy of Sciences USA*, 98, 15387–15392.
- Levin, D. T. (2000). Race as a visual feature: Using visual search and perceptual discrimination tasks to understand face categories and the cross-race recognition deficit. *Journal of Experimental Psychology: General*, 129, 559–574.
- Luck, S. J., & Hillyard, S. A. (1994). Electrophysiological correlates of feature analysis during visual search. *Psychophysiology*, 31, 291–308.
- Macrae, C. N., & Bodenhausen, G. V. (2000). Social cognition: Thinking categorically about others. *Annual Review of Psychology*, 51, 93–120.
- Macrae, C. N., Bodenhausen, G. V., & Milne, A. B. (1995). The dissection of selection in person perception: Inhibitory processes in social stereotyping. *Journal of Personality and Social Psychology*, 69, 397–407.
- Malpass, R. S., & Kravitz, J. (1969). Recognition for faces of own and other race. *Journal of Personality and Social Psychology*, 13, 330–334.
- McCarthy, G., & Donchin, E. (1981, January 2). A metric for thought: A comparison of P300 latency and reaction time. *Science*, 211, 77–80.
- Mouchetant-Rostaing, Y., Giard, M., Bentin, S., Aguera, P., & Pernier, J. (2000). Neurophysiological correlates of face gender processing in humans. *European Journal of Neuroscience*, 12, 303–310.
- Ritter, W., Simson, R., & Vaughan, H. G. (1983). Event-related potential correlates of two stages of information processing in physical and semantic discrimination tasks. *Psychophysiology*, 20, 168–179.
- Scheffers, M. K., & Coles, M. G. H. (2000). Performance monitoring in a confusing world: Error-related brain activity, judgments of response accuracy, and types of errors. *Journal of Experimental Psychology: Human Perception and Performance*, 26, 141–151.
- 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.
- Sidanius, J., & Pratto, F. (1999). *Social dominance: An intergroup theory of social hierarchy and oppression*. New York: Cambridge University Press.
- Smith, E. R., & Zarate, M. A. (1992). Exemplar-based model of social judgment. *Psychological Review*, 99, 3–21.
- Smulders, F. T. Y., Kenemans, J. L., Schmidt, W. F., & Kok, A. (1999). Effects of task complexity in young and old adults: Reaction time and P300 latency are not always dissociated. *Psychophysiology*, 36, 118–125.
- Stangor, C., Lynch, L., Duan, C., & Glass, B. (1992). Categorization of individuals on the basis of multiple social features. *Journal of Personality and Social Psychology*, 62, 207–218.
- Stroessner, S. J. (1996). Social categorization by race or sex: Effects of perceived non-normalcy on response times. *Social Cognition*, 14, 247–276.
- Wijers, A., Mulder, G., Okita, T., Mulder, L. J. M., & Scheffers, M. (1989). Attention to color: An analysis of selection, controlled search, and motor activation, using event-related potentials. *Psychophysiology*, 26, 89–109.
- Zarate, M. A., Bonilla, S., & Luevano, M. (1995). Ethnic influences on exemplar retrieval and stereotyping. *Social Cognition*, 13, 145–162.
- Zarate, M. A., & Smith, E. R. (1990). Person categorization and stereotyping. *Social Cognition*, 8, 161–185.

Received May 24, 2002

Revision received May 2, 2003

Accepted May 2, 2003 ■