Title

Bearing False Witness Under Pressure: Implicit and Explicit Components of Stereotype-Driven Memory Distortions

Journal
Social Cognition, 21(3)

ISSN
0278-016X

Authors
Sherman, Jeffrey W
Groom, Carla J
Ehrenberg, Katja
et al.

Publication Date
2003-06-01

DOI
10.1521/soco.21.3.213.25340

Peer reviewed

Social Cognition, Vol. 21, No. 3, 2003, pp. 213-246

BEARING FALSE WITNESS UNDER PRESSURE: IMPLICIT AND EXPLICIT COMPONENTS OF STEREOTYPE-DRIVEN MEMORY DISTORTIONS

Jeffrey W. Sherman
Northwestern University

Carla J. Groom
University of Texas, Austin

Katja Ehrenberg and Karl Christoph Klauer
Rheinische Friedrich-Wilhelms-Universität, Bonn, Germany

This research asked why people falsely remember stereotype-consistent information when cognitive resources are depleted. A task adapted from Jacoby's (1991) process dissociation procedure assessed participants' ability to distinguish between distractor items and behaviors performed by a stereotyped target. A multinomial analysis revealed that when cognitive capacity was restricted, participants were less likely to base judgments of stereotype-consistent behaviors on recollection and more likely to respond based on the mere familiarity of the behaviors. Capacity depletion did not affect the basis for judging stereotype-inconsistent items, nor did depletion promote simple stereotype-consistent response bias. We discuss the implications for stereotyping and eyewitness testimony.

Test your movie trivia: In the movie *West Side Story*, a fight between Puerto Rican and Caucasian street gangs escalates suddenly to the point where one member of each gang is killed. One homicide was accidental, the other deliberate. Can you remember which gang was

This research was supported by grant MH59774 from the National Institute of Mental Health to the second author. We thank Riki Conrey for her helpful feedback on drafts of this article.

Correspondence regarding this article should be addressed to Jeff Sherman, Department of Psychology, Northwestern University, 2029 Sheridan Road, Evanston, IL, 60208-2710; E-mail: sherm@northwestern.edu.

responsible for the deliberate murder? If you decided that the Caucasian "Jets" were to blame, you would be right. However, if you were a witness under the pressure of a court cross-examination rather than reading this article, you might be biased by a stereotype that Hispanic men are aggressive and believe that you had witnessed a Puerto Rican "Shark" deliberately murder a Caucasian "Jet."

STEREOTYPES AND FALSE MEMORY

False memory phenomena have long been a concern of clinical psychologists (e.g., Brewin & Andrews, 1998; Memon & Young, 1997) and of cognitive psychologists (see Koriat, Goldsmith, & Pansky, 2000, for a review). More recently, social psychologists have warned that the content of false memory is often not arbitrary, but rather is shaped by expectancies arising from stereotypes about social groups. For example, when Lenton, Blair, and Hastie (2001) exposed participants to a list of occupations that related closely to either male or female gender stereotypes, they discovered that false alarms on a recognition test were highest for distractor items that were stereotype-consistent. Similarly, Slusher and Anderson (1987) reported that people found it difficult to distinguish between stereotypic pairs of occupations and traits that had actually been presented and pairs that had only been imagined. Even if perceivers accurately remember that a stereotypic event occurred, they may be unable to attribute that event to its appropriate source. Instead, these events are often falsely attributed to a target belonging to the stereotyped group (Bayen, Nakamura, Dupuis, & Yang, 2000; Mather, Johnson, & De Leonardis, 1999; Sherman & Bessenoff, 1999; Spaniol & Bayen, 2002). On the basis of a meta-analysis, Stangor and McMillan (1992) concluded that memory was indeed better for expectancy-mismatched than expectancy-matched information when memory was measured using recall or recognition. At the same time, there was a strong response bias toward expected information in recognition memory, yielding a high rate of false memories. These phenomena have obvious consequences for members of stereotyped groups, ranging from being inaccurately labeled with stereotypic traits to being accused of serious crimes on the basis of stereotype-tainted errors of memory.

THE ROLE OF COGNITIVE RESOURCES IN STEREOTYPE-BASED FALSE MEMORY

One important factor that has been shown to affect memory for stereotype-relevant information is the availability of cognitive resources. Capacity restrictions promote reliance on relatively effortless, heuristic processes relative to effortful, controlled processes, and as such are generally associated with increased reliance on stereotypes (for a review, see Sherman, Macrae, & Bodenhausen, 2000). One result is that the relative advantage for stereotype-inconsistent over stereotype-consistent information in recognition accuracy is increased when the information is encoded under a cognitive load (see Sherman & Frost, 2000, for a discussion of the differential effects of encoding load on recall vs. recognition; Sherman, Lee, Bessenoff, & Frost, 1998).

Sherman and Bessenoff (1999) showed that memory for stereotype-relevant information is also affected when resources are diminished during memory retrieval. They used a 2-day procedure to assess whether capacity restriction would exacerbate the tendency to misattribute stereotypic behaviors to a member of a stereotyped group. Because the present research was based on their procedure, it is worth describing in detail. On Day 1, participants saw two different series of behaviors. Each series contained 30 statements describing friendly, unfriendly, and neutral behaviors. The first list was introduced as merely material to be memorized. The second list derived purportedly from an interview with "Bob," who was a member of either a stereotypically unfriendly group (a skinhead) or a stereotypically friendly group (a priest). On Day 2, a recognition task tested memory for Bob's behaviors. Participants were instructed to respond "yes" to Bob's behaviors (List 2 items) but "no" to the made-up (List 1) items. Participants also needed to respond "no" to List 3 items, which comprised friendly and unfriendly behaviors that had not been presented on Day 1. To impose a cognitive load as they performed this recognition task, some participants rehearsed an 8-digit number. Results suggested that when resources were restricted in this way, people relied more heavily on their stereotypes. Participants' errors were biased systematically under load, such that behaviors were more likely to be misattributed to Bob when the behaviors were consistent with the stereotype applied to him. For example, "carried a switchblade knife" was more likely to be falsely attributed to Bob when Bob was a skinhead rather than a priest. However, this load-induced bias applied only to distractor items presented on Day 1 (i.e., List 1 items). Misattributions of novel (i.e., List 3) items to Bob were consistently low across high- and low-capacity conditions, and for both stereotype-consistent and inconsistent items. This three-way interaction suggested that the difficulty of distinguishing familiar (i.e., List 1) behaviors that Bob had not done from those he had, combined with the processing limitations imposed by cognitive load, prompted participants to use their stereotypes to help them judge which items should be associated with Bob. Thus, once-witnessed stereotypic behaviors may be particularly likely to be wrongly attributed to a member of the stereotyped group when the memory judgment is made under low-capacity conditions, for example, a witness under cross-examination.

These results demonstrated that stereotypes biased memory for stereotype-relevant information when cognitive capacity was restricted. However, Sherman and Bessenoff's (1999) procedure was not designed to answer questions regarding the processes implicated in this bias. Memory judgments may be derived from multiple sources, any of which alone or in combination could produce such errors. In this case, one possibility is that cognitive load influences the extent to which participants are willing or able to engage in conscious recollection of the source of stereotype-consistent behaviors. It may be particularly difficult to identify the source of consistent behaviors because perceivers often extract only the gist meanings of these behaviors and not their details (e.g., Sherman et al., 1998). Perceivers also may be less willing to exert the effort to retrieve the details of consistent behaviors because there is typically a lower standard for accepting the occurrence of stereotypic than counterstereotypic behavior (see Hamilton & Sherman, 1994, for a review).

A second possible explanation for the misattribution results is that cognitive load changes the contribution of more implicit memory processes to the judgment of stereotype-consistent items. Under load, the sheer familiarity of List 1 items may serve as evidence that those were behaviors that Bob performed when the item is consistent rather than inconsistent with his stereotype. Again, the differing standards of evidence required for consistent and inconsistent information may affect judgments in this way. Alternatively, it could be the case that stereotypic items feel more familiar than counterstereotypic ones, and that this difference is particularly pronounced under cognitive load when stereotype use is particularly prevalent. In this case, it is not a question of making differential use of equally familiar consistent and inconsistent behaviors; rather, the behaviors are actually experienced as more familiar when they are consistent.

Finally, a third potential explanation is that load encourages the use of stereotypes as mental short cuts in general, such that perceivers would have a response bias in favor of stereotype-consistent judgments when attentional resources were scarce. That is, participants may simply be more likely to guess that consistent items are old in these conditions. Each of these three possibilities alone or in combination could produce the stereotypic misattribution bias reported by Sherman and Bessenoff (1999).

If we wish to understand how best to reduce the prevalence of stereotypic memory biases, it is critical to distinguish among these explanations because they offer quite different solutions to the problem. If the effect is due to diminished conscious recollection of stereotypic behaviors, a controlled process, then errors may be reduced by increasing perceivers' accuracy motivation, both during initial encoding and retrieval. Increasing the motivation to form accurate, individuated impressions would encourage people to go beyond extracting the basic gist of stereotypic behaviors (e.g., Sherman et al., 1998) and encode specific details, thereby enhancing source recollection. Increasing accuracy motivation during retrieval would encourage people to adopt a more strict recollection criterion, reducing the incidence of false recollections of stereotypic behaviors.

If the effect is due to familiarity, the problem would be much less tractable. Familiarity influences judgments at an implicit, unconscious level (Jacoby, 1991) and, as such, it is difficult, if not impossible, to subjectively quantify its effects and correct for them. Indeed, increasing accuracy motivation may only exacerbate the influence of familiarity because people believe that the feeling is a valid memory cue. In fact, familiarity is a valid memory cue much of the time, when the only source of the familiarity is prior exposure. Thus, controlling the use of familiarity may be a poor solution. Rather, it would seem necessary to alter the extent to which stereotypic behaviors may feel more familiar than counterstereotypic ones. The only way to do that would be to change people's stereotypes, likely a much more difficult solution than changing their motivations.

Finally, if the effect is due to response bias, then the manner in which perceivers encode and retrieve behaviors is irrelevant to the stereotypic memory errors identified by Sherman and Bessenoff (1999), as is the extent to which familiarity drives memory judgments. In this case, the effect results from a simple bias that leads people to claim to have witnessed stereotypic events that they cannot recollect, that may not even feel familiar, and that never happened. Alhough this bias sounds particularly pernicious, correcting it may be a relatively simple matter. If people are warned that they are prone to such guessing biases, then they may correct for them, intentionally guessing in the opposite direction (Wegener & Petty, 1997). Of course, people typically are not privy to the strength of their biases, and overcorrection or undercorrection is a likely outcome. In this case, then, perhaps the most promising solution is to simply warn people not to guess if they cannot recollect the behavior and if it does not feel familiar.

Of the three possible explanations, the response bias interpretation seems the least likely because Sherman and Bessenoff found that stereotypicality and capacity did not influence misattributions of List 3 (novel) items, only misattributions of List 1(familiar) items. At the same time, other researchers have implicated response biases as the primary source of observed stereotypic misattributions (Bayen et al., 2000; Mather et al., 1999; Spaniol & Bayen, 2002). However, none of the past research was designed to separate response biases (which influence judgments about all items) from familiarity effects (which influence only old items). This observation provides even greater reason to consider and attempt to separate different processes that may underlie stereotypic misattribution effects.

SEPARATING THE ROLES OF RECOLLECTION AND FAMILIARITY: THE PROCESS DISSOCIATION PROCEDURE

Fortunately, a general technique known as the process dissociation procedure has been developed that can tease apart the roles of recollection, familiarity, and bias in memory tasks (Buchner, Erdfelder, & Vaterrodt-Plünnecke, 1995; Jacoby, 1991; Jacoby, Toth, & Yonelinas, 1993). First, we will describe the means of separating recollection and familiarity. Further on, we will describe the means of separating familiarity and bias. To separate recollection and familiarity effects, the process dissociation procedure requires two forms of a memory task, one of which allows recollection and familiarity to work together, and one of which places them in competition. To the extent that participants cannot override the influence of implicit familiarity effects on their responses, performance on the mismatched task will be poor compared with that on the matched task.

In one of the original studies using the procedure (Jacoby, 1991), participants studied two lists of words and were later given a recognition test. One version of the test instructed participants to identify as "old" all items that had appeared in either of the study lists, responding "new" only to items never seen before. Termed an inclusion task, controlled recollection was set to work in concert with feelings of familiarity. As Equation 1 shows, responses based on either or both processes would yield correct answers. In other words, the probability of a correct response equaled the probability of responding based on recollection, R, plus the probability of responding based on familiarity if recollection failed, $F(1 - R)$.

(1) P (Correct responses on the inclusion task) = $R + F(1 - R)$

A second, "exclusion" version of the test instructed participants to respond "old" only to words from the second study list. Words from the first list, as well as new distractor items, were to be labeled "new." Because both study lists evoked a feeling of familiarity, controlled attempts to identify items from List 2 alone now competed with the familiarity of List 1 items. As shown in Equation 2, the probability of an incorrect response to a List 1 item corresponded to the probability of responding based on familiarity when recollection failed.

(2) P (Incorrect responses on exclusion task) = F (1 - R).

Simple algebra reveals the two processes contributing to responses under each condition (see Equations 3 and 4).

(3) $R = P$ (Correct on inclusion task) - P (Incorrect on exclusion task)

(4) $F = P$ (Incorrect on exclusion task)/1-R

The influence of further experimental manipulations on familiarity-based and recollection-based processes can then be assessed separately.

APPLYING PROCESS DISSOCIATION TO EXAMINE MEMORY FOR STEREOTYPE-RELATED BEHAVIORS

The Sherman and Bessenoff (1999) procedure corresponds to the exclusion phase of a process dissociation study. Familiarity and recollection processes were placed in opposition by asking participants to respond "yes" only to behavior items performed by the stereotyped target (i.e., List 2 items) and not to new distractors (List 3 items) or familiar distractors (List 1 items). According to Equation 2, misattributions of List 1 items to List 2 reflect the implicit use of familiarity operating in the wake of the failure of recollection. Systematic variations on this misattribution measure were the central findings of the Sherman and Bessenoff study. When resources were restricted, the proportion of errors in response to List 1 items was higher for stereotype-consistent than stereotype-inconsistent behaviors. However, without a corresponding inclusion task, the finding is ambiguous with respect to process because it is not possible to solve for F and R. Cognitive load may lower the rate of recollection (i.e., raise the value of 1 - R) for stereotype-consistent relative to stereotype-inconsistent items or it may strengthen implicit experiences of the familiarity of stereotype-consistent items relative to stereotype-inconsistent items (i.e., raise the value of F), or some combination of both. Hense, Penner, and Nelson (1995; see also Banaji & Greenwald, 1995) demonstrated greater familiarity for stereotype-consistent than stereotype-inconsistent traits, although they did not examine these implicit effects as a function of processing capacity during recognition, and they did not separate familiarity from response bias. In the present research, we sought to discover whether cognitive load promoted stereotype-consistent misattributions by affecting the deployment of controlled retrieval processes, by distorting more implicit influences of memory, or simply by altering perceivers' response bias.

THE PRESENT RESEARCH

The aim of the present research was to isolate the processes responsible for stereotype-biased false memory under cognitive load by extending Sherman and Bessenoff's (1999) procedure into a full process dissociation design. Whereas the original study asked participants to respond only to the exclusion question "Did Bob do this?" the extended design also included a phase that required participants to consider each item in the context of the inclusion question "Did you see this yesterday?"

As in the original study, participants read two lists of behavior statements. List 1 was ascribed to a neutral target person.[1] List 2 was ascribed to a target person belonging to a stereotypically friendly or unfriendly group (priests or skinheads, respectively). The following day, participants performed two sets of recognition judgments. One set constituted an exclusion task, requiring participants to discriminate List 2 behaviors from items belonging to List 1 or to List 3 (novel distractors). As in the original study, all three lists included behaviors that were consistent and inconsistent with the stereotype associated with the target. The second set of judgments, an inclusion task, required participants to decide simply whether the item was new (List 3) or had appeared in either of the lists shown the previous day (Lists 1 and 2). Half of the participants were placed under a cognitive load throughout each recognition task.

1. In a minor change from the original procedure, List 1 was presented as behaviors describing an information technology worker (a group neutral with respect to kindness), rather than as made-up material. This was intended to make discrimination between Lists 1 and 2 more difficult, and the high proportion of misattributions reported in the Results section testifies to the success of this modification.

USING MULTINOMIAL MODELING TO REFINE PROCESS DISSOCIATION ANALYSES

Without changing the logic of the process dissociation procedure, it is possible to achieve even purer estimates of recollection and familiarity processes by employing a set of multinomial modeling equations offered by Buchner et al. (1995), rather than the original equations used by Jacoby et al. (1993). Buchner et al. pointed out that response bias would be a third process that contributed to judgments, in addition to recollection and familiarity assessment. Response bias would come into play whenever participants lacked memory for an item entirely and were forced to guess. Participants would lack memory whenever they were faced with a novel item, one that was presented but not encoded, or one that was inaccessible in memory.

The multinomial model supplies estimates of all three processes (familiarity, recollection, and bias) by looking not only at responses to List 1 items, but also at responses to List 3 items (novel distractors). By definition, participants will not have any memory or familiarity associated with List 3 items. Therefore, to the extent that participants respond “yes” to these items, then they are likely to also be biased to respond “yes” to List 1 items that are neither recollected nor accepted based on familiarity. The original process dissociation equations (Jacoby et al., 1993) cannot distinguish between List 1 errors that arise from the familiarity of those items and errors that arise from response bias.

At the same time, the Buchner et al. (1995) model does not allow for conscious detection of List 3 distractors as new, a process that is very likely to occur in the present context. As such, we modified the Buchner et al. model according to a standard procedure in multinomial modeling (see Batchelder & Riefer, 1990; Bayen, Murnane, & Erdfelder, 1996; Klauer & Wegener, 1998). This procedure assumes that conscious detection of distractors occurs with the same probability as the conscious recollection of old items. Psychologically, they are highly similar: The latter is based on knowledge about an occurrence, whereas the former is based on knowledge about a nonoccurrence (Strack & Bless, 1994). The modified version of the model differs from the original one only with respect to List 3 items. In the modified model, responses to List 3 items are classified as old if and only if the responses reflect guesses in the absence of conscious novelty detection (see Equations 2a and 4a below).

Figure 1 depicts the logic of the modified model in the form of a multinomial processing tree. The tree is not an information-processing flowchart, but rather is intended to separate logically the components underlying judgment by depicting the conditional probabilities of certain processes occurring given the outcome of other processes. The model does not imply a psychological order of the processes, nor does it imply that processes earlier in the tree are more sensitive to experimental manipulations than are processes later in the tree.

Recollection

The recollection component (r) corresponds to the probability of successfully remembering a List 1 item (including remembering the fact that it came from List 1), as well as consciously rejecting a List 3 item, after response bias is factored out. Reponses based on recollection will by definition be correct (i.e., "yes" under inclusion instructions, "no" under exclusion instructions for List 1 items, and "no" to List 3 items in either condition).

Familiarity

The measure f estimates the probability that responses to List 1 items are based on familiarity, given failed recollection. The estimate is based on comparing the proportion of "yes" responses to List 1 items in the inclusion and exclusion tasks.

Response Bias

Response bias reflects the probability of guessing "yes" in the absence of both recollection and familiarity.[2] The proportion of incorrect "yes" responses to List 3 items constitutes the basis for estimating this bias. Because response bias may differ under exclusion and inclusion

2. The response bias parameter should not be confused with guessing per se. Rather, bias is considered one possible component (along with recollection and familiarity) that may form the basis for decisions (including guesses) in a memory task. That is, participants may guess in a memory task based on recollection, familiarity, or bias. However, guesses that are based on recollection or familiarity are not guesses in the purest sense and would not be considered as response bias in the model, which reserves that term for responses given in the absence of recollection and familiarity.

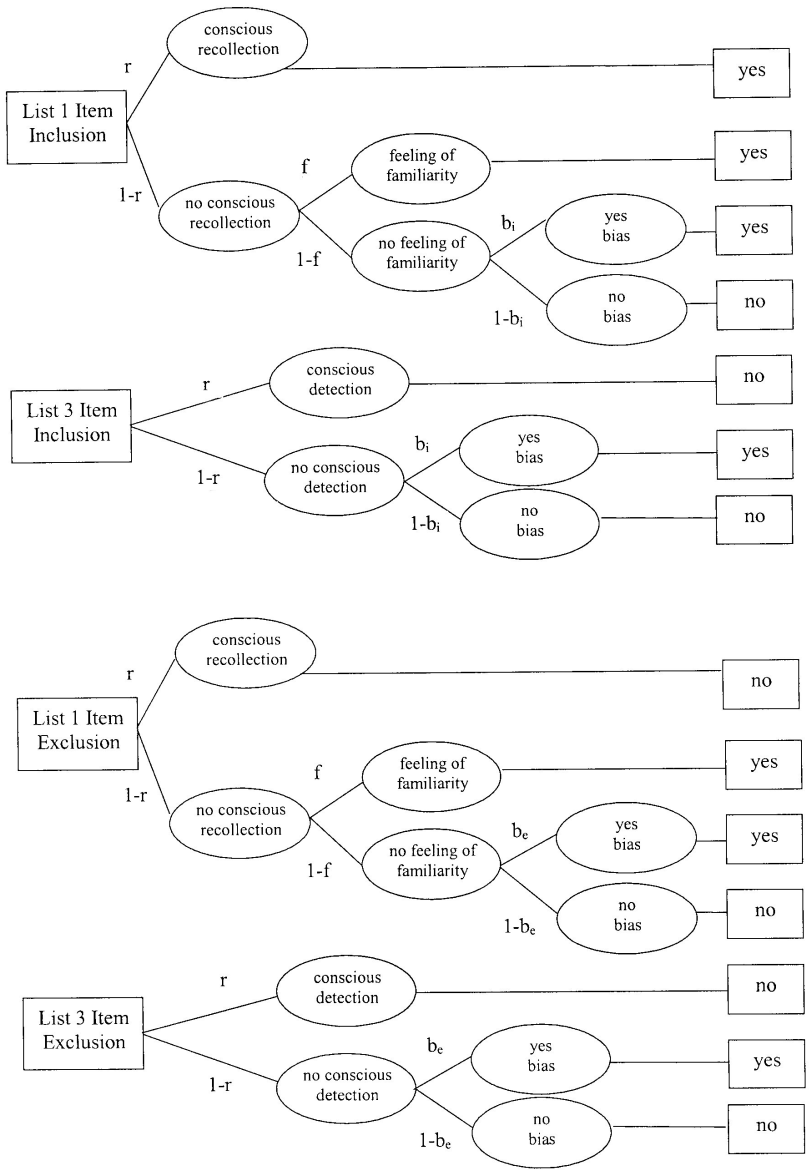

FIGURE 1. The modified multinomial model of process dissociation for judgments of List 1 and List 3 items.

instructions (e.g., because the participant is sensitive to the fact that the ratio of target items to distractors is higher in the inclusion than during the exclusion phase), separate bias parameters are estimated for responses under inclusion instructions (b_i) and under exclusion instructions (b_e).

Model Equations

The processing tree translates directly into four independent equations. Equation 1a represents the probability of correctly responding "yes" to the inclusion question "Did you see this yesterday?" when the item came from List 1. Equation 2a represents the probability of incorrectly responding "yes" under inclusion instructions when the item was a List 3 distractor. The latter two equations represent the probability of incorrectly responding "yes" to the exclusion question "Did Bob do this?" when the item came from List 1 (Equation 3a) or List 3 (Equation 4a).

(1a) p (yes response to List 1 item in inclusion task) =
$r + ((1 - r)^* f) + (1-r)^*(1-f)^*b_i$

(2a) p (yes response to List 3 item in inclusion task) = $(1 - r)^*b_i$

(3a) p (yes response to List 1 item in exclusion task) =
$((1 - r)^*f) + (1 - r)^*(1 - f)^*b_e$

(4a) p (yes response to List 3 item in exclusion task) = $(1 - r)^*b_e$

By separating responses to stereotype-consistent and stereotype-inconsistent items, we were able to estimate the contributions of recollection, familiarity, and response bias, and observe how these contributions were affected by item type and processing capacity.

METHOD

PARTICIPANTS

Forty-eight undergraduates from the Introductory Psychology subject pool at Northwestern University participated for course credit. Participants were run in groups of up to four people.

MATERIALS AND PROCEDURE

The experiment involved two sessions on consecutive days.

Day 1: Impression Formation Tasks

In the first session, participants performed two impression formation tasks in sequence. The two target people were purportedly Chicago-area residents who had supplied information about their activities during earlier interviews. Each task began with a brief biographical sketch of the target person, followed by an instruction to form an impression of the person based on his or her behavior statements. A list of behaviors was then presented via microcomputer at the rate of one statement every 6 seconds.

Participants first learned about "David Johnson," who was described as an information technology worker. Depending on randomly assigned condition, the second target was introduced either as "Bob Hamilton, a skinhead," or as "Father Bob Hamilton, a priest." These stereotypic labels were intended to generate the expectation that the target would be unkind or kind, respectively.

List 1 (David) and List 2 (Bob) each contained 30 behaviors selected on the basis of pretest ratings. One third of the statements in each list were kind (e.g., "took his elderly aunt out for dinner"), one third were unkind (e.g., "swore at the salesgirl") and one third were irrelevant to the dimension of kindness (e.g., "listened to a new CD"). The kind and unkind behaviors were consistent or inconsistent with the stereotype of Bob, depending on the participant's target condition (i.e., skinhead or priest). Individual items appeared in only one list. An additional list was created in the same manner and reserved as a source of novel distractors for the memory tasks on Day 2.

Day 2: Recognition Tasks

In the second session, 24 hours later, participants' recognition memory for the previously learned material was unexpectedly tested twice, once under exclusion instructions and once under inclusion instructions. The order of the tasks was randomly determined. In the exclusion task, participants were asked, "Did Bob do this?" Thus, "yes" was the correct response only for List 2 items. For List 1 items and the new distractor items from List 3, the correct response was "no." In the inclusion task, participants were asked, "Did you see this yesterday?" A "yes" response was appropriate this time for both List 1 and List 2 items, but not for List 3 items. During both the inclusion and exclusion tasks, participants made yes/no judgments for all 90 behaviors at their own pace, but were encouraged to respond as quickly as possible without sacrificing accuracy.

Half of the participants were placed under conditions of low cognitive capacity by requiring that they rehearse 8-digit numbers while performing the memory tasks, a manipulation that has successfully served as a cognitive load in previous research (e.g., Sherman & Bessenoff, 1999; Sherman et al., 1998). To avoid placing participants under load during the instructions for each task, two 8-digit numbers were used, one for each task. One number was provided immediately following instructions for the exclusion task, and a different number was provided following instructions for the inclusion task. As compliance checks, participants wrote each number down as soon as the task was complete.

RESULTS

MISATTRIBUTIONS IN THE EXCLUSION TASK

Using the same criteria as Sherman and Bessenoff (1999), a misattribution was defined as a "yes" response to a List 1 or List 3 item, when instructed to respond "yes" only to Bob's behaviors (i.e., during the exclusion task). Misattributions as a proportion of all responses to distractors were subjected to a 2 (target type: skinhead vs. priest) × 2 (cognitive capacity: high vs. low) × 2 (order: inclusion task first vs. exclusion task first) × 2 (stimulus type: stereotype-consistent vs. stereotype-inconsistent behavior) × 2 (distractor novelty: List 1 vs. List 3) mixed ANOVA. The final two factors were entered as within-subjects variables. Target type did not enter into any main effects or interactions, and the remaining analyses collapsed across this factor.

As found by Sherman and Bessenoff (1999), there was a main effect of distractor novelty, $F(1, 40) = 81.12, p < .001$. List 1 behaviors ($M = .42$) were more likely to be misattributed to Bob than were novel distractors ($M = .12$). This effect was qualified by the predicted three-way interaction among novelty, capacity, and stimulus type, F

(1, 40) = 7.01, $p = .01$. For List 1 behaviors consistent with Bob's stereotype, participants with low capacity were more likely to make misattributions than were participants with high capacity (low capacity: $M = .49$, high capacity: $M = .33$), $F\ (1, 40) = 6.16$, $p < .02$. Capacity had no effect on misattributions of List 1 behaviors that were inconsistent with Bob's stereotype (low capacity: $M = .44$, high capacity: $M = .42$), $F < 1$. For List 3 behaviors consistent with the stereotype, misattributions were actually slightly more likely when capacity was high ($M = .14$) than when it was low ($M = .10$), $F < 1$. Again, capacity had little effect on misattributions of stereotype-inconsistent behaviors (low capacity: $M = .13$, high capacity: $M = .11$), $F < 1$.

However, the three-way interaction was itself moderated by a four-way interaction with task order, $F\ (1, 40) = 4.39$, $p < .05$. Separate analyses of the data from the two task orders revealed the nature of this moderating order effect. Participants who performed the inclusion task first did not show the expected and previouslyreported three-way interaction among capacity, item novelty, and stimulus consistency, $F = .16$. In contrast, participants who performed the exclusion task first showed a highly reliable three-way interaction among novelty, capacity, and stimulus type, $F\ (1, 20) = 10.68$, p < .01 (see Figure 2). The pattern for exclusion-first participants was similar to that reported for the data set as a whole. For misattributions of List 1 behaviors, there was a significant two-way interaction between capacity and stimulus type, $F\ (1, 20) = 6.70$, $p < .02$ (see left panel, Figure 2). When judging a List 1 behavior that was consistent with Bob's stereotype, participants with low cognitive capacity were more likely than were participants with high capacity to make a misattribution (low capacity: $M = .58$, high capacity: $M = .24$). Yet when a List 1 behavior was inconsistent with Bob's stereotype, misattributions occurred at equivalent rates across the capacity conditions (low capacity: $M = .49$, high capacity: $M = .38$). The simple contrast between high- and low-capacity conditions was significant for consistent items, $F\ (1, 20) = 15.52$, $p < .01$, but not for inconsistent items, $F < 1$. Conversely, this pattern did not hold when the distractor was a List 3 item, $F\ (1, 20) = 2.58$, $p = .12$ (see right panel, Figure 2). Indeed, the trend was in the opposite direction. Compared to participants with high capacity, participants in the low-capacity condition were slightly less likely to make misattributions for stereotype-consistent

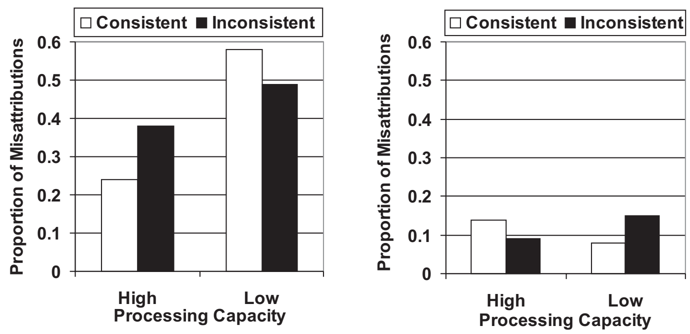

FIGURE 2. Source misattributions as a function of processing capacity and stimulus type: Left panel, List 1 distractors; Right panel, List 3 distractors.

behaviors (low capacity: $M = .08$, high capacity: $M = .14$), $F(1, 20) = 1.11$, $p = .31$, but more likely to make misattributions for stereotype-inconsistent behaviors (low capacity: $M = .15$, high capacity: $M = .09$), $F < 1$. Thus, the results from participants performing the exclusion task first replicated the results from Sherman and Bessenoff's (1999) original experiment (in which participants performed only an exclusion task): Participants relied more heavily on their stereotypes when resources were depleted, but only when source confusion was high (i.e., when the distractor was familiar).

It is not clear why the results from participants performing the inclusion task first did not replicate those of Sherman and Bessenoff (1999). One possibility is that the exclusion task is easier to perform when the inclusion task is performed first. When the exclusion task is performed first, participants must assess both whether or not an item is old and whether or not it was performed by David or Bob (if it is deemed old). In contrast, when the exclusion task is performed second, participants need not decide whether or not the item is old because they have presumably already made this judgment during the inclusion task. As such, the exclusion task is less demanding when it is performed second, obviating the need to rely on stereotypes to resolve the discrepancy between familiarity and recollection in performing the task on List 1 items, even when capacity is low. Whatever the cause, these data make clear that stereotypic memory errors are not inevitable. We next try to understand the factors that contribute to such errors when they occur.

MULTINOMIAL MODEL ESTIMATES OF PROCESSES UNDERLYING RESPONSES

Estimates of response bias for both the inclusion and exclusion task phases (b_i and b_e) together with estimates of recollection and familiarity were calculated using the modified version of Buchner et al.'s multinomial model. Participants received separate scores for stereotype-consistent and stereotype-inconsistent items. The estimates reported exclude participants who performed the inclusion task first, because it makes little sense to decompose an effect using data that do not show that effect. The goal is to understand the nature of the effect when it occurs. Data from one participant were excluded from parameter estimation procedures because he or she responded yes to all of the List 3 items in one stimulus category.[3]

Overall Model Fit

The first question in applying a multinomial model is whether or not the model provides a good overall fit of the data. Goodness of fit is assessed

3. We report within-subjects analyses. For these analyses, parameter estimates for each participant were obtained using their responses in the exclusion and inclusion task phases. Alhough it is possible to obtain these individual estimates algebraically, the solutions result in very complex equations. As such, we derived the parameter estimates iteratively via Maximum Likelihood Estimation (MLE) with standard software. The two methods produce identical estimates. These estimates were then analyzed with standard ANOVA procedures. We found the same pattern of results when we used multinomial model-fitting procedures in which parameters were derived via MLE from cell frequencies that were aggregated across participants. In the interest of presentation clarity, we do not present these analyses, but they are available on request from the authors. All analyses reported here also were run on the basis of the original Buchner et al. (1995) model and yielded essentially the same results. These analyses also are available upon request.

by evaluating differences between the observed and the estimated response frequencies of the model. This assessment yields a likelihood ratio statistic, G^2, which equals asymptotically χ^2 distributed. For saturated models like the one we used (i.e., models in which the number of parameters equals the number of response categories), the value of G^2 should equal 0, and that was, in fact, the case for our model. This demonstrates that our model is capable of accounting for the observed data.[4]

Bias Under Exclusion Instructions (b_e)

Across all conditions, participants were relatively reluctant to ascribe novel distractors to Bob, tending to produce false alarms at a rate below 25%. Individual estimates of this bias were analyzed in a 2 (target type: skinhead vs. priest) × 2 (cognitive capacity: high vs. low) × 2 (stimulus type: stereotype-consistent vs. stereotype-inconsistent) mixed-design ANOVA. This analysis produced a reliable two-way interaction involving cognitive capacity and stimulus type F (1, 19) = 4.62, $p < .05$ (see Figure 3). This effect reflects the fact that bias for inconsistent items was slightly greater in the low- (M = .25) than high- (M = .19) capacity condition, F (1, 19) = 1.46, p = .24, whereas bias for consistent items was greater in the high- (M = .29) than low- (M = .17) capacity condition, F (1, 19) = 2.73, p = .12. Not surprisingly, this effect mimics the effect reported above concerning the simple number of false alarms for List 3 items in the exclusion condition. As above, this effect is clearly inconsistent with the hypothesis that simple bias accounts for the misattribution results reported here and by Sherman and Bessenoff (1999).

Bias Under Inclusion Instructions (b_i)

Individual estimates of this bias were analyzed in a 2 (target type: skinhead vs. priest) × 2 (cognitive capacity: high vs. low) × 2 (stimulus type: stereotype-consistent vs. stereotype-inconsistent) mixed-design

4. Statistical tests of model fit are not possible with saturated models. The original Buchner et al. (1995) model also provided a good fit of the data, G^2, = 0. However, the "two-high" threshold source monitoring model developed by Bayen and her colleagues (Bayen et al., 1996; Bayen et al., 2000; Spaniol & Bayen, 2002) did not fit the data. That model produces separate recollection parameters for item (old vs. new) and for source (List 1 vs. List 2). The model also estimates two guessing parameters, one for item and one for source. The model's lack of fit appears to be due to a lower frequency of correct responses on List 1 items than the model predicts. These analyses also are available upon request.

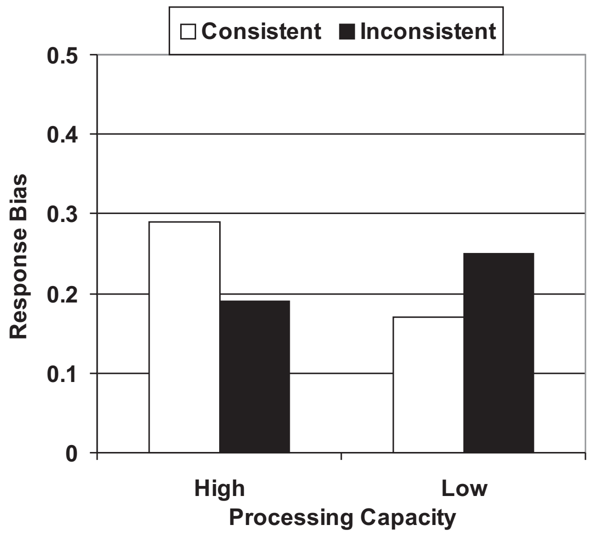

FIGURE 3. Response bias as a function of processing capacity and stimulus type: Exclusion instructions only.

ANOVA. This analysis demonstrated a main effect of stimulus type, such that bias was stronger for consistent (*M* = .34) than inconsistent (*M* = .09) items, $F\ (1, 19) = 27.09$, $p < .001$. However, this main effect was qualified by a reliable two-way interaction involving cognitive capacity and stimulus type, $F\ (1, 19) = 4.78$, $p < .05$ (see Figure 4). In this case, bias for inconsistent items was equal in the high- (*M* = .09) and low- (*M* = .09) capacity conditions, $F < 1$, whereas bias for consistent items was greater in the high (*M* = .44) than low-capacity (*M* = .25) condition, $F\ (1, 19) = 4.24$, $p = .05$. Together with the b_e results, these results argue strongly against the idea that stereotypic misattributions increase when capacity is restricted because participants resort to a simple response bias in the direction of one's stereotypes under such conditions.

Recollection

Recollection, *r*, reflects the likelihood of successfully remembering List 1 items and consciously rejecting List 3 items, after response bias is factored out. Individual estimates of this parameter were analyzed in a 2 (target type: skinhead vs. priest) × 2 (cognitive capacity: high vs. low) ×

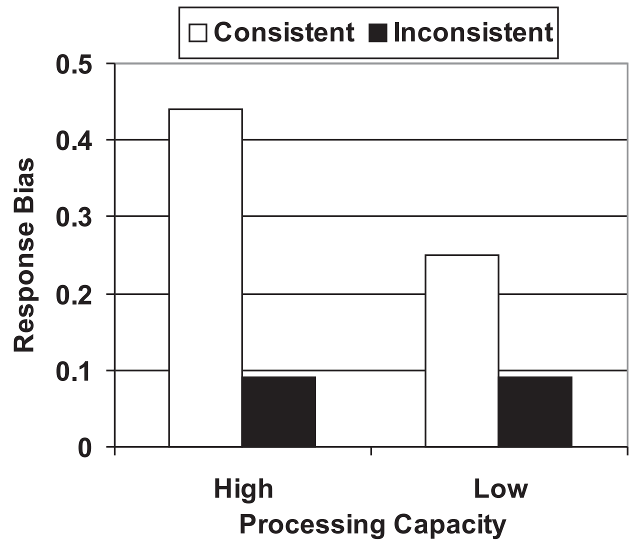

FIGURE 4. Response bias as a function of processing capacity and stimulus type: Inclusion instructions only.

2 (stimulus type: stereotype-consistent vs. stereotype-inconsistent) mixed-design ANOVA. This analysis produced a reliable two-way interaction involving cognitive capacity and stimulus type, F (1, 19) = 13.25, $p < .01$ (see Figure 5). For stereotype-consistent items, assignments were less likely to reflect recollection when capacity was low (M = .18) than when it was high (M = .37). Conversely, for stereotype-inconsistent items, recollection was equal in the two capacity conditions (low capacity: M = .24, high capacity: M = .21). The simple contrast between high- and low-capacity conditions was significant for consistent items, F (1, 19) = 11.52, $p < .01$, but not for inconsistent items, F (1, 19) = .15, *ns*.

Thus, participants under cognitive load showed a reduced probability of consciously recollecting stereotypic items together with a tendency to misattribute List 1 stereotypic items in the same condition. It appears that the misattribution data may be explained, at least in part, in terms of a reduction in the conscious recollection of familiar stereotypic items when capacity is low.

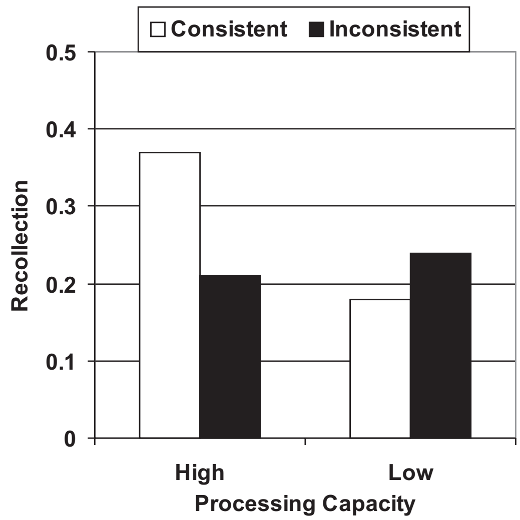

FIGURE 5. Recollection as a function of processing capacity and stimulus type.

Familiarity

The familiarity parameter, *f*, estimates the probability that responses to List 1 items are based on familiarity, given failed recollection. Individual estimates of this parameter were analyzed in a 2 (target type: skinhead vs. priest) × 2 (cognitive capacity: high vs. low) × 2 (stimulus type: stereotype-consistent vs. stereotype-inconsistent) mixed-design ANOVA. This analysis produced a reliable two-way interaction involving cognitive capacity and stimulus type, F (1, 19) = 9.80, $p < .01$ (see Figure 6). Once again, the effect of capacity was restricted to the stereotype-consistent items. Simple contrasts confirmed that depleting resources increased familiarity-based responses for consistent items (low capacity: $M = .64$, high capacity: $M = .32$), F (1, 19) = 9.09, $p < .01$. In contrast, capacity had no effect on familiarity for stereotype-inconsistent items (low capacity: $M = .57$, high capacity: $M = .55$), F (1, 19) = 0.75, $p = .40$.[5]

5. This two-way interaction was qualified by a three-way interaction with target type, F (1, 19) = 5.14, $p = .04$. This three-way interaction reflects the fact that the two-way interaction was stronger among participants told that List 2 described a skinhead rather than a priest, although the pattern of means was in the same direction in both cases.

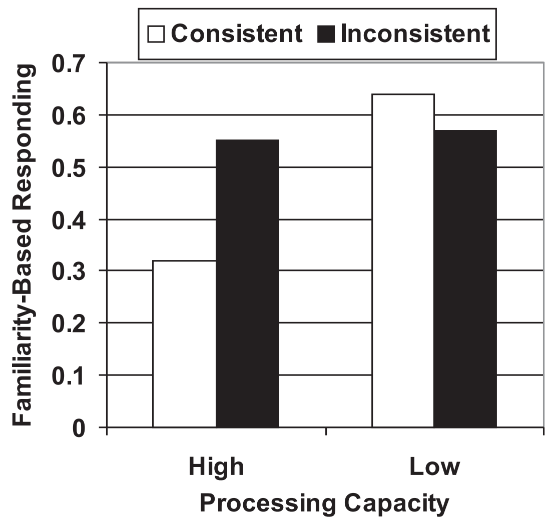

FIGURE 6. Familiarity-based responding as a function of processing capacity and stimulus type.

Thus, participants under cognitive load showed an increased influence of familiarity in responding to stereotypic items together with a tendency to misattribute stereotypic List 1 items in the same condition. At least part of the basic misattribution effect may therefore be due to a load-induced shift in favor of responding to those items based on sheer item familiarity whenever recollection fails.

DISCUSSION

Memory often fails. Even if one is able to accurately remember that an event did, in fact, occur, attempts to ascribe the event to the appropriate context and/or actor may fail (Johnson, Hashtroudi, & Lindsay, 1993; Koriat et al., 2000). The likelihood of such memory distortion is affected by a variety of factors, including properties of the event in question. Of interest in the present research, the extent to which an event is consistent with a stereotype-based expectancy affects the likelihood that the event will be misattributed to an incorrect context or actor (Bayen et al., 2000; Mather et al., 1999; Sherman

& Bessenoff, 1999; Spaniol & Bayen, 2002). Factors affecting the perceiver's mental state also may influence memory distortions. We have shown that stereotype-based misattributions are particularly likely to occur when processing capacity is diminished during attempted recollection. Such states frequently occur outside the laboratory, where levels of stress and distraction are higher than inside the laboratory.

DECOMPOSING STEREOTYPIC MISATTRIBUTIONS

Certainly, stereotypic memory errors are not inevitable (as demonstrated by the four-way interaction for misattributions). However, in this article, our goal was to achieve a more fine-grained analysis of the factors that underlie such errors when they occur. We hoped to move beyond characterizing the nature of these distortions and examine the cognitive processes that produce them. Relying on the logic of Jacoby's (1991) process dissociation paradigm and multinomial models of that logic (e.g., Buchner et al., 1995), we sought to identify the contributions of conscious recollection, feelings of familiarity, and response bias to stereotypic misattributions. Results showed that these processes contribute differentially, depending on the item to be remembered and cognitive load. Cognitive load did not affect the manner in which stereotype-inconsistent items were remembered. For these items, neither recollection nor familiarity varied as a function of cognitive load. In contrast, for stereotype-consistent items, load did affect conscious recollection and familiarity. When capacity was depleted, the contribution of conscious recollection was diminished and the contribution of familiarity was enhanced.[6] Thus, the tendency to misattribute stereotypic events when capacity is low appears to be due, in part, to a diminished role for conscious recollection coupled with an increased role

6. These results may also be described as showing relatively enhanced recollection and diminished familiarity-based responding for inconsistent compared to consistent items in the low capacity (vs. high capacity) condition. These results are similar to the effects of cognitive load on the encoding of stereotype-consistent and stereotype-inconsistent information (e.g., Sherman & Frost, 2000; Sherman et al., 1998).

for familiarity. Attribution performance suffers because feelings of familiarity do not discriminate between List 1 and List 2 items. Discriminating between these two equally familiar sources requires conscious recollection. These findings add to the litany of mechanisms through which stereotypes perpetuate themselves (for a review, see Hamilton & Sherman, 1994).

THE BASES OF SHIFTS IN PARAMETER INFLUENCE

Although this study allowed us to examine the independent contributions of recollection and familiarity processes to stereotype-based memory errors, it did not provide definitive evidence of the bases of the effects in these parameters. In the case of recollection, we do not know whether the diminished influence on stereotypic items under load reflects changes in effort or ability to recollect those items. There are good reasons to expect both effects. Participants may exert less effort to recollect the details of stereotype-consistent items because they match expectations. The burden of proof is much higher for expectancy violations (see Hamilton & Sherman, 1994, for a review). Ability to recollect consistent behaviors also may suffer when capacity is depleted. Stereotype-consistent and stereotype-inconsistent events are encoded differently when they are encountered. Perceivers may extract only the basic gist of consistent items because they fit expectations. In contrast, perceivers are more likely to encode the specific details of inconsistent behaviors (Sherman, 2001; Sherman et al., 1998). When these perceivers are subsequently asked to recollect the details, it would be more difficult to produce them for the consistent items, particularly if resources are limited. Thus, it is likely that changes in recollective influence on stereotypic items reflect both motivation and ability factors.

There also are two good explanations for the enhanced influence of familiarity on judgments of stereotype-consistent items when resources are low. One possibility is that consistent items feel more familiar in the high-load condition. Because stereotypes are more active under such conditions (Sherman et al., 2000), consistent items may feel more stereotypic, and therefore more familiar or fluent (Banaji & Greenwald, 1995; Whittlesea & Leboe, 2000). Another possibility is that the items feel equally familiar in the two conditions, but that participants are more willing to accept that level of familiarity as a memory cue for consistent items, particularly in the high-load condition. Memory cues based on feelings of familiarity are easily applied when capacity is low and recollection is difficult. Thus, their enhanced influence may result from their ease of application in combination with a reduced standard of evidence for stereotypical information when resources are low. Of course, it is entirely possible that both factors contribute to the familiarity results demonstrated in this study. Distinguishing among the different explanations for the fluctuating influence of recollection and familiarity will be an important avenue for future research.

IMPLICATIONS

These results suggest a number of important practical implications for reducing stereotypic memory distortions. Because the effect appears to be due both to diminished recollection of stereotypic behaviors and to increased reliance on familiarity in judging those behaviors, a dual plan of attack is warranted to reduce such errors. First, enhancing perceivers' accuracy motivation, both when they initially observe potentially stereotype-relevant behaviors and when they subsequently try to remember them, should enhance accurate recollection and diminish errors. However, when recollection fails, the present results indicate that implicit feelings of familiarity may also produce stereotypic misattributions. As suggested in the introduction, increased motivation would be unlikely to diminish the extent to which people rely on familiarity, and may well enhance the use of such cues. Instead, the best strategy to reduce such effects may be to minimize the subjective familiarity of stereotypic behaviors. Perhaps the best way to accomplish this would be to weaken or change people's stereotypes. So it seems that reducing stereotypic memory errors may require changes both in people's motivations and in their group knowledge.

Our results suggest two other basic remedies to the problem of stereotypic memory distortions. First, given the fact that the distortions appear to occur primarily when perceivers are under cognitive load, one obvious solution is to try to ensure that people are relaxed and fully capacitated when they are trying to remember stereotype-relevant events. This suggestion would seem particularly important when the stakes are high, as in cases of eyewitness testimony.

Another remedy is suggested by the fact that stereotypic distortions were absent when perceivers performed an inclusion task before performing an exclusion task. In particular, this indicates that it may be prudent to test memory in two separate phases. In the first phase, perceivers would simply be asked if they did or did not witness a particular action. Only after that judgment would they be asked to indicate the specific person who performed the action.

FAMILIARITY AND AUTOMATICITY

To the extent that the influence of familiarity is presumed to be automatic, it may seem surprising that the use of familiarity-based responding was influenced by processing capacity in the present research. However, we view the status of familiarity processes as automatic to be an empirical question. The familiarity parameter in our research reflects both the experience of familiarity and the use of familiarity in making judgments. As described earlier, there are reasons to expect that stereotypic information may feel more or less familiar, depending on processing resources. There also are reasons to expect that variations in processing capacity may lead perceivers to rely more or less on familiarity cues in making memory judgments about stereotypic events. More generally, we would argue that memory cues based on familiarity operate in much the same way as do other kinds of memory cues (e.g., vividness, amount of contextual detail, schema-consistency), in that their use may be subject to many different moderating influences (e.g., Jacoby, Marsh, & Dolan, 2001; Johnson et al., 1993; Leboe & Whittlesea, 2002; Whittlesea & Williams, 2001).

We also would note that "automaticity" is not a monolithic concept. Bargh (1994) described four different ways in which a process may be said to be automatic: if it occurs without the person's awareness, if it occurs without intention, if it is relatively independent of processing resources, and if it cannot be controlled or inhibited. By Bargh's account, the use of familiarity may be seen as automatic in that it is used without much (if any) awareness or intent. But familiarity is only relatively automatic in terms of its resource-dependence, at best. It is certainly resource-independent in comparison to recollection, but it also can be influenced by variations in processing capacity, as in the present study (see also Whittlesea & Williams, 2001). We have suggested that it is unlikely that people can control the use of familiarity; however, we know of no direct empirical evidence pertaining to this question.

ON RESPONSE BIASES

One factor that did not influence the misattribution results was response bias. In contrast to the List 1 items, misattributions of List 3 items were less stereotypical when capacity was diminished, suggesting that a pure response bias cannot account for the misattribution data. The bias parameter under exclusion instructions showed the same counterstereotypical bias when capacity was restricted. Under inclusion instructions, there was an overall bias toward stereotypic responding. However, this bias was smaller in the low- than high-capacity condition. Again, this shows that response bias did not produce the overall effect of increasing stereotypic misattributions under cognitive load. Indeed, the misattribution results are even stronger in light of the direction of the bias. That is, despite a shift in bias away from calling consistent items old when capacity was depleted, the tendency to misattribute List 1 consistent items compared to inconsistent items was greatly increased under a cognitive load.

These conclusions may be contrasted with those of other researchers (Banaji & Greenwald, 1995; Bayen et al., 2000; Mather et al., 1999; Spaniol & Bayen, 2002), who proposed a significant role for response bias in producing stereotypic memory distortions. However, none of those studies attempted to separately assess response bias and familiarity, and therefore they could not distinguish between the two influences. Hense et al. (1995) concluded that such stereotypic errors were based on the familiarity of stereotype-consistent information. However, Hense et al. also did not separate familiarity and bias, and therefore they could not rule out bias as a potential explanation for their effects. One significant contribution of the present research is the separate assessment of bias and familiarity. Our results showed a clear dissociation between the two components. Whereas familiarity-based responding increased for consistent items when capacity was diminished, responses based on bias in favor of consistent items decreased in the same condition.

We do not wish to suggest that response biases never influence memory for stereotype-relevant material. Indeed, we cannot conclude from the present data that our participants did not have such a bias. It is entirely possible that our participants had a bias to respond positively to stereotypic behaviors that was counteracted by attempts to correct for the bias (Wegener & Petty, 1997). In the introduction of the article, we suggested that one possible solution for a response bias would be to correct in the opposite direction. Given the clear relevance of stereotypes about skinheads and priests to participants' memory judgments, it would not be surprising if participants recognized the possibility for bias and corrected for it. Perhaps this correction was deemed more necessary under low capacity conditions, when participants knew they lacked the resources to engage in careful recollection. In any case, we are not in a position to claim that our participants lacked bias. However, we can state confidently that such a bias cannot account for the pattern of misattributions we observed. Clearly, there is much to do in future research to specify the conditions under which bias and familiarity influence memory, and the nature of those influences. To do so will require techniques that permit the separation of the two components.

ON THE JOINT CONTRIBUTIONS OF IMPLICIT AND EXPLICIT PROCESSES

Recently, the distinction between implicit and explicit processes has taken a central role in many areas of social-psychological research, including research on stereotyping and prejudice. Most often, social-psychological processes have been characterized as either implicit or explicit by nature. The extent to which a process is implicit or explicit is frequently determined by performance on two separate tasks, one that is presumed to measure implicit processes, and one that is presumed to measure explicit processes (e.g., Devine, 1989; Dovidio, Kawakami, Johnson, Johnson, & Howard, 1997; Fazio, Jackson, Dunton, & Williams, 1995; Greenwald, McGhee, & Schwartz, 1998; Wittenbrink, Judd, & Park, 1997; for exceptions see Payne, 2001; Payne, Lambert, & Jacoby 2002). One problem with this approach is that it confounds task with processing type. There may be a number of important differences between the two tasks that affect performance, beyond their characterization as "implicit" or "explicit." Moreover, few, if any, measures are "process pure." That is, performance on very few tasks is determined wholly by either implicit or explicit processes. Rather, most tasks include both implicit and explicit components. It was this realization that led to the development of the process dissociation paradigm (Jacoby, 1991), and the attempt to measure the joint contributions of implicit and explicit processes to performance on a single memory task.

We believe there are great advantages to decomposing a single task into multiple components in social psychological research as well. Most obviously, it presents a more complex and detailed profile of task performance. For example, in the present research, we were able to separately assess the influence of response bias and familiarity on memory errors. Doing so permitted us to draw conclusions that could not be drawn using other techniques (e.g., Banaji & Greenwald, 1995; Hense et al., 1995; Mather et al., 1995). We also demonstrated that the processing of stereotypic information is affected in two distinct ways when capacity is diminished: recollection is diminished, and the role of familiarity is enhanced. Note that there is no mathematical necessity in the model we applied that the two processes demonstrate this reciprocal relationship. The two parameters may vary independently, and have been shown to do so (Buchner et al., 1995). In fact, the within-subject correlations between recollection and familiarity in the present research were weak: -.24 ($p = .26$) for consistent items and -.20 ($p = .34$) for inconsistent items. Thus, within the same task, the results demonstrate two important independent effects on memory judgments for consistent items. This is much richer data than can be gathered with a task that is presumed to reflect only one or the other process. It also permits much stronger conclusions than results obtained on two different measures that may be confounded in a variety of different ways. Finally, we believe that this approach presents a more accurate portrayal of social cognitive processes, few of which are entirely implicit or explicit.

CONCLUSION

Stereotype-driven errors in source memory may lead you to falsely remember that the Sharks rather than the Jets committed homicide in *West Side Story.* But identifying the bases of these errors and possible remedies is no trivial pursuit. In the real world, the consequences of such errors may be great. Violent crimes may be misattributed to innocent bystanders or “usual suspects” who happen to belong to minority groups stereotyped as violent, leading to the false conviction and punishment of innocent individuals. Such errors may be particularly likely when a witness is pushed to make a quick identification or is made nervous by the presence of armed police officers. In this context, it is sobering to consider the extent to which Blacks in the United States are overrepresented on death row, in many cases on the basis of eyewitness testimony. The goal of the present research was to begin to identify the specific cognitive processes that underlie stereotypic memory errors. Not surprisingly, the cognitive bases of these errors are complex and subtle. When resources are depleted, the very strategies and/or memory cues that perceivers use to make source attributions are altered for stereotypic events. Overcoming such effects may be extremely difficult and likely requires correction strategies that are equally complex and subtle, and that target multiple cognitive processes and motivations.

REFERENCES

Banaji, M. R., & Greenwald, A. G. (1995). Implicit gender stereotyping in judgments of fame. *Journal of Personality and Social Psychology, 68,* 181-198.

Bargh, J. A. (1994). The four horsemen of automaticity: Awareness, intention, efficiency, andcontrol in social cognition. In R. S. Wyer, Jr., & T. K. Srull (Eds.), *Handbook of social cognition* (2nd ed., Vol. 1, pp. 1-40). Hillsdale, NJ: Erlbaum.

Batchelder, W. H., & Riefer, D. M. (1990). Multinomial processing models of source monitoring. *Psychological Review, 97,* 548-564.

Bayen, U. J., Murnane, K., & Erdfelder, E. (1996). Source discrimination, item detection, and multinomial models of source monitoring. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22,* 197-215.

Bayen, U. J., Nakamura, G. V., Dupuis, S. E., & Yang, C.-L. (2000). The use of schematic knowledge about sources in source monitoring. *Memory and Cognition, 28*, 480-500.

Brewin, C. R., & Andrews, B. (1998). Recovered memories of trauma: Phenomenology and cognitive mechanisms. *Clinical Psychology Review, 18,* 949-970.

Buchner, A., Erdfelder, E., & Vaterrodt-Plünnecke, B. (1995). Toward unbiased measurement of conscious and unconscious memory processes within the process dissociation framework. *Journal of Experimental Psychology: General, 124,* 137-160.

Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology*, *56*, 5-18.

Dovidio, J. F., Kawakami, K., Johnson, C., Johnson, B., & Howard, A. (1997). On the nature of prejudice: Automatic and controlled processes. *Journal of Experimental Social Psychology*, *33*, 510-540.

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

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

Hamilton, D. L., & Sherman, J. W. (1994). Stereotypes. In R. S. Wyer, Jr., & T. K. Srull (Eds.), *Handbook of social cognition* (2nd ed., Vol. 2, pp. 1-68). Hillsdale, NJ: Erlbaum.

Hense, R. L., Penner, L. A., & Nelson, D. L. (1995). Implicit memory for age stereotypes. *Social Cognition*, *13*, 399-415.

Jacoby, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language, 30*, 513-541.

Jacoby, L. L.., Marsh, E. L., & Dolan, P. O. (2001). Forms of bias: Age related differences in memory. In M. Naveh-Benjamin, M. Moscovitch, & H. Roediger (Eds.), *Perspectives of human memory and cognitive aging: Essays in honor of Fergus Craik* (pp. 240-252). New York: Routledge.

Jacoby, L. L., Toth, J. P., & Yonelinas, A. P. (1993). Separating conscious and unconscious influences of memory: Measuring recollection. *Journal of Experimental Psychology: General, 122*, 139-154.

Johnson, M. K., Hashtroudi, S., & Lindsay, D. S. (1993). Source monitoring. *Psychological Bulletin, 114*, 3-28.

Klauer, K. C., & Wegener, I. (1998). Unraveling social categorization in the "Who said what?" paradigm. *Journal of Personality and Social Psychology, 75,* 1155-1178.

Koriat, A., Goldsmith, M., & Pansky, A. (2000). Toward a psychology of memory accuracy. *Annual Review of Psychology, 51,* 481-537.

Leboe, J. P., & Whittlesea, B. W. A. (2002). The inferential basis of familiarity and recall: Evidence for a common underlying process. *Journal of Memory and Language*, *46*, 804-829.

Lenton, A. P., Blair, I. V., & Hastie, R. (2001). Illusions of gender: Stereotypes evoke false memories. *Journal of Experimental Social Psychology, 37,* 3-14.

Mather, M., Johnson, M. K., & De Leonardis, D. M. (1999). Stereotype reliance in source monitoring: Age differences and neuropsychological test correlates. *Cognitive Neuropsychology, 16,* 437-458.

Memon, A., & Young, M. (1997). Desperately seeking evidence: The recovered memory debate. *Legal and Criminological Psychology, 2,* 131-154.

Payne, B. K. (2001). Prejudice and perception: The role of automatic and controlled processes in misperceiving a weapon. *Journal of Personality and Social Psychology*, *81*, 181-192.

Payne, B. K., Lambert, A. J., & Jacoby, L. L. (2002). Best laid plans: Effects of goals on accessibility bias and cognitive control in race-based misperceptions of weapons. *Journal of Experimental Social Psychology*, *38*, 384-396.

Sherman, J. W. (2001). The dynamic relationship between stereotype efficiency and mental representation. In G. Moskowitz (Ed.), *Cognitive social psychology: The Princeton Symposium on the Legacy and Future of Social Cognition* (pp. 177-190). Mahwah, NJ: Erlbaum.

Sherman, J. W., & Bessenoff, G. R. (1999). Stereotypes as source monitoring cues: On the interaction between episodic and semantic memory. *Psychological Science, 10*, 106-110.

Sherman, J. W., & Frost, L. A. (2000). On the encoding of stereotype-relevant information under cognitive load. *Personality and Social Psychology Bulletin*, *26*, 26-34.

Sherman, J. W., Lee, A. Y., Bessenoff, G. R., & Frost, L. A. (1998). Stereotype efficiency reconsidered: Encoding flexibility under cognitive load. *Journal of Personality and Social Psychology. 75,* 589-606.

Sherman, J. W., Macrae, C. N., & Bodenhausen, G. V. (2000). Attention and stereotyping: Cognitive constraints on the construction of meaningful social impressions. *European Review of Social Psychology, 11*, 145-175.

Slusher, M. P., & Anderson, C. A. (1987). When reality monitoring fails: The role of imagination in stereotype maintenance. *Journal of Personality and Social Psychology, 52*, 653-662.

Spaniol, J., & Bayen, U. T. (2002). When is schematic knowledge used in source monitoring? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28*, 631-651.

Stangor, C., & McMillan, D. (1992). Memory for expectancy-congruent and expectancy-incongruent information: A review of the social and social developmental literatures. *Psychological Bulletin, 111*, 42-61.

Strack, F., & Bless, H. (1994). Memory for nonoccurrences: Metacognitive and presuppositional strategies. *Journal of Memory and Language, 33,* 203-217.

Wegener, D. T., & Petty, R. E. (1997). The flexible correction model: The role of naive theories of bias in bias correction. *Advances in Experimental Social Psychology, 29,* 141-208.

Whittlesea, B. W. A., & Leboe, J. P. (2000). The heuristic basis of remembering and classification: Fluency, generation, and resemblance. *Journal of Experimental Psychology: General, 129,* 84-106.

Whittlesea, B. W. A., & Williams, L. D. (2001). The discrepancy-attribution hypothesis: I. The heuristic basis of feelings and familiarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*, 3-13.

Wittenbrink, B., Judd, C., & Park, B. (1997). Evidence for racial prejudice at the implicit level and its relationship with questionnaire measures. *Journal of Personality and Social Psychology*, *72*, 262-274.