

Forgetting by remembering: Stereotype inhibition through rehearsal of alternative aspects of identity[☆]

Elizabeth W. Dunn* and Barbara A. Spellman¹

Department of Psychology, University of Virginia, 102 Gilmer Hall, P.O. Box 400400, Charlottesville, VA 22904-4400, USA

Received 27 December 2001; revised 16 December 2002

Abstract

We applied previous research on retrieval-induced forgetting to the issue of stereotype inhibition. All participants learned about a target person who belonged to a stereotyped group, and then practiced retrieving a subset of the target's characteristics. When participants practiced individuating information about the target, they showed inhibited memory for the target's stereotypic traits. When participants practiced stereotypic information about the target, they showed inhibited memory for: (a) traits associated with another stereotyped aspect of the target's identity; (b) individuating traits of the target; and (c) other, unpracticed traits of the target associated with the same stereotype. Stereotype belief moderated these inhibition effects; the more strongly participants believed in the stereotype, the less inhibition of stereotype-relevant traits they showed.

© 2003 Elsevier Science (USA). All rights reserved.

Keywords: Stereotype; Inhibition; Cross-categorization; Retrieval-induced forgetting

Introduction

Imagine a White female college student interacting with her Black male professor. Although the student may sincerely wish to avoid viewing the professor through the lens of racial stereotypes, her conscious attempts to avoid stereotypical thinking may backfire, particularly if her mental resources are partially occupied by other tasks (e.g., Wyer, Sherman, & Stroessner, 2000). Is the student doomed, then, to the inescapable

influence of racial stereotypes in perceiving and interacting with the professor? Perhaps not. After all, the professor is not only an African-American, but also a *professor*—not to mention, a complex person with a range of individuating personality traits (e.g., an exceptional sense of humor). By focusing on the teacher's professorial or individuating characteristics, then, the student may be able to successfully inhibit stereotypes about African-Americans.

A handful of recent studies provide initial evidence that such inhibitory processes may play an important role in stereotyping and person perception (Dijksterhuis & van Knippenberg, 1995, 1996, 1998; Hugenberg & Bodenhausen, 2001; Macrae, Bodenhausen, & Milne, 1995; Rudman & Borgida, 1995; Stroessner, 1996). For example, Macrae et al. (1995) showed participants a video of a Chinese woman either putting on makeup (priming the female stereotype) or eating noodles with chopsticks (priming the Chinese stereotype). In the lexical decision task that followed, participants who had seen the makeup version were faster to identify female-relevant words (e.g., *emotional*), but slower to identify Chinese-relevant words (e.g., *calm*), compared to unprimed control participants; the reverse pattern emerged for participants who had seen the chopsticks version.

[☆] This research was supported in part by a National Science Foundation Graduate Research Fellowship to the first author and an NIMH Grant to the second author. We thank Jason Mitchell for helpful discussions, Joanna Hale and Carmina Liuksila for data collection, and Bella DePaulo for statistical advice. We are grateful to Timothy Wilson, John Skowronski, Stacey Sinclair, Galen Bodenhausen, Neil Macrae, and two anonymous reviewers for helpful comments on earlier drafts. Portions of these results were presented at the 2000 meeting of the American Psychological Society, Miami Beach, FL; the 2001 Meeting of the Society for Personality and Social Psychology, San Antonio, TX; and the 2001 International Conference on Memory, Valencia, Spain.

* Corresponding author. Fax: +1-434-982-4694.

E-mail addresses: edunn@post.harvard.edu (E.W. Dunn), spellman@virginia.edu (B.A. Spellman).

¹ Also corresponding author.

Macrae et al. assert that “both applicable categories (i.e., Chinese and woman) are activated in parallel and that a competition for mental dominance ensues,” leading to the excitation of the winner and the inhibition of the loser (p. 398).²

This research suggests that inhibition may be involved in the category-selection process that occurs when we initially encounter multiply-categorizable targets. However, what about the later stages of person perception when we have already learned about a person and need to remember information about the person’s characteristics—might similar inhibitory processes come into play then as well? If focusing on one stereotyped aspect of a person inhibits another stereotyped aspect, might focusing on one stereotyped aspect inhibit memory for individuating information about the person? More optimistically, might focusing on individuating information inhibit the stereotype? And what role would belief in such stereotypes play?

To examine these questions we borrowed a technique from the cognitive literature called the retrieval-induced forgetting paradigm (see Levy & Anderson, 2002, for a review). This technique provides a method for creating and examining inhibition in long-term memory. First, we describe the retrieval-induced forgetting paradigm and how it has been applied to person perception. Next, we describe important factors to consider in applying this paradigm to the stereotyping domain, including the potential role of stereotype belief. Then, we discuss mechanisms through which belief in the stereotype might moderate inhibition. Finally, we present three experiments that address the questions posed above.

Social and non-social demonstrations of retrieval-induced forgetting

The phenomenon of retrieval-induced forgetting (RIF) was first demonstrated by Anderson, Bjork, and Bjork (1994) and has been replicated many times using various kinds of stimuli (see Levy & Anderson, 2002, for a list; see also Bäuml, 2002; Bäuml & Hartinger, 2002; Macrae & Rosevare, 2002). The standard version of the retrieval-induced forgetting paradigm has four phases: a learning phase, a retrieval-practice phase, a distracter phase, and a test phase. For example, in Anderson, Bjork, and Bjork (1994), participants studied category–exemplar pairs from eight target categories, each of which had six exemplars (e.g., FRUIT–ORANGE; HOB-

BIES–COINS; FRUIT–BANANA). Participants were told to study each pair by relating the exemplar to the category. Next, participants practiced retrieving half of the items from half of the target categories; they were shown a category name and the first two letters of an exemplar (e.g., FRUIT–OR__) and told to write down an item that they had studied earlier that would fit the word stem. After a distracter task, participants were shown the name of each previously studied category in turn and asked to write down all of the words they had studied from that category.

These manipulations create three kinds of items: practiced items from practiced categories (e.g., ORANGE); unpracticed items from practiced categories (e.g., BANANA); and baseline items—that is, items from categories in which no items were practiced (e.g., COINS). Not surprisingly, Anderson et al. (1994) found that practiced items were recalled better than baseline items. The interesting question, however, is what would happen to the unpracticed items like BANANA? Would these words be remembered better than the baseline items because of their association with the strengthened items from their category? Or would memory for them be impaired? In fact, memory for the unpracticed items was impaired relative to the baseline items. This phenomenon was called “retrieval-induced forgetting” because it results from the retrieval of competing items from memory.³

What is the mechanism behind this forgetting? Anderson and Spellman (1995) pointed out that although the previous results might be due to inhibition of the concept BANANA, the evidence is ambiguous. Because the impaired items (e.g., BANANA) were both studied and tested under the same cue (e.g., FRUIT) as the strengthened items (e.g., ORANGE), the impairment could be due to other, non-inhibitory, mechanisms such as blocking. According to the blocking explanation, the strengthened item ORANGE repeatedly comes to mind when the test probe FRUIT is presented. The representation of BANANA itself might not be inhibited; rather, access to it through the cue FRUIT might be blocked by ORANGE. (See Anderson & Neely, 1996, for a review of the possible non-inhibitory mechanisms.) Using a more complicated design, Anderson and Spellman demonstrated that the representations of the impaired items were, in fact, inhibited. (See also Anderson & Bell, 2001.)

According to Anderson and his colleagues, this form of inhibition occurs because when given the cue FRUIT-

² Note that these results do not demonstrate inhibition that is specific to the un-primed category. Participants viewing the makeup version are faster on the female-relevant words than on the Chinese-relevant words but we do not know what happens to other neutral words. Perhaps neutral words would also be slower, indicating a general criterion shift, not inhibition that is specific to the competing Chinese stereotype.

³ Retrieval-induced forgetting does not result merely from the strengthening of competing items. See Anderson, Bjork, and Bjork (2000), in which additional presentations of a category–exemplar pair increases recall for that pair but does not impair recall for non-represented pairs. See also Bäuml and Hartinger (2002), and MacLeod and Macrae (2001).

OR__ both ORANGE and BANANA (plus all the other fruits studied) are likely to be initially activated. In order to retrieve the target ORANGE from out of the activated noise, the other items that were activated need to be suppressed. As a result of this suppression, non-practiced FRUIT words become difficult to remember when given the original FRUIT cue or an alternative cue.

Social stimuli in retrieval-induced forgetting. Recent work by Macrae and MacLeod (1999) and MacLeod and Macrae (2001) illustrates that the phenomenon of retrieval-induced forgetting extends to some social stimuli. Participants were asked to form impressions of two men by studying a series of index cards, each of which listed one man's name along with a single personality trait. Afterward, participants rehearsed half of the traits associated with one of the men. Following a distracter period, participants were asked to write down all of the traits that had originally been associated with each of the men. Consistent with previous research on RIF, participants: (a) exhibited enhanced memory for the traits they had practiced while (b) displaying inhibited memory for the unpracticed traits associated with the man whose other traits had been rehearsed.

In these experiments, each man was described by 10 traits that were unrelated to each other. For example, one man was described by the traits "sensible," "romantic," and "vigorous," a combination that almost precludes forming a clear impression of him. In utilizing traits that do not hang together, the experimental design blunts the perceiver's usual role in organizing the information based on pre-existing knowledge about the social world. While a perfectly appropriate strategy for a first investigation of RIF in the social domain, this work does not attempt to address how perceivers' social knowledge interacts with the general phenomenon of retrieval-induced forgetting. For example, if one learns that a man is sensible and vigorous and then rehearses "sensible," one may easily forget that an unrelated trait (i.e., "vigorous") also describes the man. But would it be so easy to forget a trait such as "rational," given that sensible people also tend to be rational? We would argue that such pre-existing knowledge about the social world may influence the extent to which people exhibit RIF in forming impressions of others.

Indeed, rather than defining others by a string of unconnected traits, people are likely to organize information about other individuals around important aspects of their identity. Because group membership often represents a central aspect of people's identity (e.g., Tajfel & Turner, 1986), information about others may frequently coalesce around their membership in various social groups, information that, for better or worse, often takes the form of stereotypes. Thus, in applying RIF to person perception, it is useful to consider the special case of stereotypes. In the present series of studies, we examine the role of retrieval-induced for-

getting in memory for stereotypical traits associated with a social group or group member, compelling a careful consideration of the special nature of social stereotypes.

Structural differences in stimuli. Importing the retrieval practice paradigm into the stereotype arena is not as straightforward as it may initially seem. Typical RIF experiments use category-exemplar pairs (e.g., FRUIT-ORANGE); the kind of pairs we are interested in using are category-feature pairs (e.g., ARTIST-CREATIVE). These pairings are not analogous: oranges are a type of fruit; creativity is a common trait (i.e., feature) among artists or belonging to particular artists. We know from previous research that category-exemplar pairs can demonstrate inhibition (e.g., Anderson et al., 1994); we also know that using social stimuli, individual-feature pairs can demonstrate inhibition (e.g., Macrae & MacLeod, 1999). We do not know, whether for social stimuli, category-feature pairs can demonstrate inhibition.

How stereotype belief might moderate inhibition

In applying previous RIF research to the issue of stereotype inhibition, it is crucial to consider how social knowledge interacts with the cognitive process of retrieval-induced forgetting. There are two findings in the cognitive RIF literature that suggest ways in which stereotype belief might moderate inhibition. These findings have to do with: (a) the strength of the links between the category and the exemplars and (b) the coherence or "integration" of a category's exemplars. Extending the analogy from those two findings to research on stereotype inhibition leads to competing predictions about how belief will affect inhibition.

The strength of the links between the category and the exemplars. Clever cognitive experiments notwithstanding, for most standard categories, most of the time, most people agree that membership is clear: oranges and bananas are fruits. But for stereotypes it is not so clear that most of the time most people agree that the stereotypic traits actually apply, for example, that Asian-Americans are intelligent and reserved. Although most people may be aware of common social stereotypes, belief in such stereotypes varies a great deal from individual to individual (Devine, 1989). The belief that a banana is a fruit is likely to vary little across individuals, but some fruits are fruitier than others; when asked to generate exemplars of the category fruits, participants are more likely to come up with apple, orange, and pear, than pomegranate, blueberry, and mango (exemplars 1, 2, and 3, and 21, 22, 23 from Battig & Montague, 1969). To the extent that people who strongly believe in a given stereotype tend to think of stereotype-relevant traits when faced with a member of that category, stereotype belief could be conceptualized as a set of strong links between the category and its stereotypic traits—just as

the category FRUIT is strongly linked to apples and other exemplars that come to mind when presented with this category.

The cognitive literature shows that exemplars that are strongly associated with their category suffer *more* RIF—both absolute and proportional—than weakly associated items; exemplars that are strongly associated with a category are likely to come to mind upon seeing the category and therefore must be strongly inhibited in order to promote successful recall of other items from the same category (Anderson et al., 1994; Bäuml, 1998). Thus, if belief in a stereotype strengthens the association between the relevant social category and its stereotypic traits, then higher levels of stereotype belief should be associated with greater retrieval-induced forgetting of these traits.

Coherence or “integration” of exemplars. Apart from strengthening the links between a category and its associated traits, stereotype belief may strengthen the links between the traits themselves via cross-item integration. Cross-item “integration” has to do with the extent to which each item is associated with and evokes thoughts of the others. For example, oranges and bananas have associations to each other beyond the fact that they are similar because they are both fruits: they are both breakfast foods, they might both be in your smoothie, and they both make good afternoon snacks. But traits like “intelligent” and “reserved” would not have those sorts of independent associations—until the label is provided. Indeed, stereotypic traits often seem to cohere only through their common association with the stereotype label.

Integrating items within a category has been shown to reduce RIF; if participants think back to other items associated with the same category whenever they see a new item, the unpracticed items from the category will show resistance to inhibition (Anderson & McCulloch, 1999). The reduction in RIF is shown both if the participants spontaneously choose to integrate items while studying or if the experimenter instructs participants to do so. Why does it happen? One possible explanation is that of “mediated retrieval.” When thinking back to previous items, connections are being created between the to-be-practiced and to-be-non-practiced items. Later, during the final recall test, such cross-item integration provides additional retrieval routes from the strengthened practiced items (which are likely to be recalled) to the unpracticed items from the same category, benefiting the latter relative to baseline items. If one embraces a stereotype, it may provide a powerful integrative framework. For example, if one really believes that Asian-Americans are studious, intelligent, reserved, diligent, educated, and precise, then remembering any of these words may bring the others back to mind because the traits are closely linked to each other by way of their common association with the stereotype. If stereotype belief strengthens the links between traits related to the

stereotype in this manner, then people who believe a given stereotype should show less retrieval-induced forgetting of stereotypic traits.

The present studies: Overview of methods and predictions

In the studies that follow, we examined how RIF applies to social stereotypes. All of the present experiments use a similar four-phase procedure based on the methods of Anderson and colleagues described above. In the initial study phase, participants learn to associate trait words with several different people or groups. In the retrieval-practice phase, participants’ attention is focused on a subset of the information they have just learned. They are shown the name or description of one of the people or groups (“target”) and are asked to retrieve half of the associated traits from memory. After a distracter period, participants are asked to recall all of the trait words they saw during the initial study phase. Finally, participants fill out a questionnaire regarding their belief in the stereotypes addressed in the experiment.

In Experiments 1 and 2, we predicted that studying a compound category (e.g., ASIAN-AMERICAN ARTIST) and rehearsing stereotypic traits associated with one of the sub-categories (e.g., ARTIST) would inhibit memory for traits associated with the other sub-category (e.g., ASIAN-AMERICAN). Further, in Experiment 3, we predicted that practicing individuating traits would inhibit stereotypic traits, and vice versa. In this experiment, we also investigated whether practicing a sub-set of the characteristic traits of a stereotype would lead to inhibited memory for other unpracticed traits related to the same stereotype.

Throughout these studies, we examined the moderating role of stereotype belief in retrieval-induced forgetting of stereotype-relevant traits. If believing in a stereotype serves merely to strengthen the links between a group label and relevant traits, then strongly held stereotypes would be expected to suffer greater inhibition than weakly held stereotypes. If, however, believing in a stereotype serves to integrate stereotype-relevant traits, then strongly held stereotypes should prove relatively resistant to inhibition.

Experiment 1: Alternative stereotyped identities

Experiment 1 addresses the question of whether focusing on one (stereotyped) aspect of a person’s identity will inhibit memory for another (stereotyped) aspect of a person’s identity.

Method

Pre-testing. We asked 69 undergraduates to write down six words that were stereotypically associated with

each of seven social groups. We obtained data for a total of 14 groups; the data revealed that our participants shared relatively clear stereotypes about a subset of those groups, including Asian-Americans, artists, feminists, and mothers. We then created a list of 84 trait words that participants frequently associated with the groups in this subset.

Next, we asked 14 participants to rate on a 5-point scale how stereotypical each of these traits was of each social group, on a scale ranging from “Very Atypical (1)” to “Unrelated (3)” to “Very Typical (5).” For the groups Asian-Americans, artists, feminists, and mothers, we selected three words that were rated as very stereotypical ($M = 4.7$) of each group, but unrelated to each of the other groups ($M = 3.1$); see Table 1. We combined these groups to create two compound categories: Asian-American artists and feminist mothers.⁴

Participants. A new sample of 55 undergraduates, including 10 males, 43 females, and 2 who did not list their gender, participated in partial fulfillment of a course requirement. Experimental sessions were run in groups of 5–10 participants.

Initial study phase. After signing consent forms, all participants received a study booklet that displayed the name of one compound category and a single associated trait word on each page. The categories included the two target categories, “Asian-American artist” and “feminist mother,” as well as a filler category, “red food.” Each category had a total of six trait words associated with it. The words were displayed with the compound category to the left and the trait word to the right (e.g., ASIAN-AMERICAN ARTIST: INTELLIGENT).

Participants studied each page for 5 s and were instructed to spend all that time relating the trait word to the compound category. The booklet contained 18 pages. It was ordered such that for the target categories: (a) a compound category never appeared twice in a row; (b) successive trait words belonging to the same compound category alternated between sub-categories; (c) the first and last pages of the booklet displayed the filler category thereby reducing primacy and recency effects on later recall of the target categories. The filler category was used flexibly to achieve the above constraints.

Retrieval-practice phase. Participants then were asked to retrieve a subset of the words they had just studied.

⁴ We asked an additional set of 36 participants to rate how surprised they would be to encounter someone who belonged to these compound categories. On a scale ranging from “not at all surprised (1)” to “very surprised (7),” participants reported that they would not be very surprised to meet an Asian-American artist ($M = 2.1$) or a feminist mother ($M = 2.1$). By selecting compound categories that our participants found relatively unsurprising, we hoped to minimize emergent attributes, traits that people sometimes generate to explain why someone belongs to two categories that are not usually combined (e.g., Harvard-educated carpenters are seen as non-materialistic; Kunda, Miller, & Claire, 1990).

Table 1
Stereotypical words associated with each group (Experiments 1 and 2)

Group	Stereotypical words		
Mothers	caring	supportive	worried
Feminists	strong	loud	ugly (bitter)
Asian-Americans	intelligent	diligent	short (reserved)
Artists	creative	talented	peculiar

Note: Words in parentheses were used in Experiment 2 to replace the adjacent words from Experiment 1.

Each page of the retrieval-practice booklet displayed one of the previously studied compound categories along with a word stem that prompted participants to recall one of the associated words (e.g., ASIAN-AMERICAN ARTIST: IN__). Participants were given 8 s to complete the stem by writing out the appropriate word beneath the blank (e.g., intelligent).

Each participant was randomly assigned to practice the words related to only one sub-category, that is, either Asian-American, artist, feminist, or mother. For example, in the Asian-American condition, the full compound category (i.e., ASIAN-AMERICAN ARTIST) was provided as a cue but the word stems led participants to retrieve only the words associated with the Asian-American stereotype (i.e., intelligent, diligent, short). Participants were asked to retrieve each of the three words included in their condition three times. To reduce cross-item integration, the same target never appeared twice in a row, and to increase item strengthening, the target items appeared at expanding intervals. In addition, the first and last pages of the booklet displayed the filler category, thereby reducing primacy and recency effects on later recall of the target categories. The filler category was used flexibly to achieve these constraints.

Distracter phase. Following the retrieval practice phase, participants spent 10 min completing a logical reasoning task. This exercise was chosen because it contained no information relevant to social stereotypes and required no memorization or retrieval.

Final test phase. Participants were then given a surprise cued recall test. Each page of the test booklet had the name of a compound category (either ASIAN-AMERICAN ARTIST or FEMINIST MOTHER) at the top. Participants were given 30 s to list all the words they had seen associated with each compound category during the initial study phase.⁵ The target categories

⁵ Final tests were limited to 30 s in order to prevent participants from using a generate-and-test recall strategy (e.g., using the cue FEMINIST MOTHER to generate feminist and mother stereotype words and then judging whether those words had been presented in the study phase). This time limit has been used in other RIF experiments in which participants' pre-experimental knowledge may simulate actual recall.

were counterbalanced for order across booklets, with the filler category always appearing last.

Questionnaire. Finally, participants completed a questionnaire that asked them to report the extent to which they believed in the various social stereotypes tapped during the experiment. For each trait word, participants were asked to rate on a 7-point scale how much they agreed that the trait accurately described the relevant group. For example, they were asked to rate how much they agreed with the statement that Asian-Americans tend to be diligent. The participants were then provided with a full explanation of the study and dismissed.

Results and discussion

Specific predictions. According to our hypotheses, focusing on information relevant to one aspect of identity should inhibit memory for information related to another aspect of identity. For example, relative to their memory for words associated with the feminist and mother stereotypes, participants who rehearsed words associated with the artist stereotype should show reduced memory for words associated with the Asian-American stereotype. For participants who rehearsed artist words, artist words represent the *focused-on set* and should be facilitated, Asian-American words represent the *alternative set* and should be inhibited, and feminist and mother words represent the *irrelevant set* and should provide a baseline measure of memory ability.⁶ Across the four retrieval-practice conditions, we expected participants to exhibit better memory for the focused-on set compared to the irrelevant set (facilitation) and to exhibit weaker memory for the alternative set compared to the irrelevant set (inhibition).

Retrieval practice success rate. Retrieval practice success (i.e., the percentage of times participants retrieved the appropriate word during retrieval practice) was approximately 83% across conditions.⁷

Transforming recall data to z-scores. Before conducting our analyses, we transformed the raw number of words participants recalled from each sub-category into z-scores. Because our choice of trait words was highly constrained by the requirement that each word be strongly related to one group and irrelevant to all others, we were unable to control for word length and frequency. As a result, words from some sub-categories were easier to remember than words from other sub-

categories. Using z-scores compensates for this problem, in that each participant's z-score reflects how well he or she recalled a given sub-category, relative to how well all other participants recalled words from that same sub-category. For example, a participant's z-score for the stereotypically maternal words reflects how well he or she remembered these words relative to how well the rest of the participants remembered the maternal words. The z-scores thus normalized each sub-category of words across conditions. It allows us to look within subject for facilitation and inhibition without worrying about which of the sub-categories were harder or easier to remember overall.

Recall data. Our primary hypotheses received strong support. Not surprisingly, a planned comparison revealed that participants exhibited better memory for the focused-on set than for the irrelevant set, $F(1, 54) = 42.76$, $p < .001$, $\eta^2 = .442$. More interestingly, participants showed inhibited memory for the alternative set compared to the irrelevant set, $F(1, 54) = 20.59$, $p < .001$, $\eta^2 = .28$ (see Table 2). Thus, studying one set of trait words not only strengthened subsequent memory for those traits, but also reduced memory for unpracticed traits associated with the target person. This result suggests that rehearsing traits related to one aspect of an individual's identity, such as profession, can weaken memory for traits related to another aspect of identity, such as race.

Effects of stereotype belief. Inhibition was moderated by participants' belief in the stereotype about the alternative set. We performed a regression using level of belief in the alternative and focused-on stereotypes to predict magnitude of RIF, an index created by subtracting z-scored recall for the alternative set from z-scored recall for the irrelevant set. The regression revealed that the *more* a participant believed in the stereotype about the alternative set the *less* retrieval-induced forgetting that participant would show for that set, $F(1, 53) = 4.87$, $p = .03$, $\beta = -.29$. For example, participants who believed that Asian-Americans were intelligent, diligent, and short demonstrated relatively little retrieval-induced forgetting for these words after rehearsing words from the artist category.

Even participants with high belief in the alternative stereotype showed some RIF, however, as illustrated by a subsidiary analysis in which a median split was used to place participants into either the high stereotype endorsement or the low stereotype endorsement group. Participants in the high stereotype endorsement group exhibited significant RIF, $M = .38$, $F(1, 26) = 4.24$, $p = .05$, $\eta^2 = .14$, albeit somewhat less than their counterparts in the low stereotype endorsement group, $M = .73$, $F(1, 27) = 20.96$, $p < .001$, $\eta^2 = .44$ (see Table 2). Inhibition of the alternative set was not moderated by level of belief in the focused-on stereotype ($\beta = -.04$, $F < 1$). The facilitation effect was not

⁶ In previous work on retrieval-induced forgetting (e.g., Anderson & Spellman, 1995), the focused-on set is called the Rp+ category, the alternative set is called the Rp- category, and the irrelevant sets are called the Nrp categories.

⁷ This figure is based on 30 out of 55 participants (or 54%) because some of the raw retrieval-practice data were lost. Data were available for at least 5 participants in each condition.

Table 2

Mean *z*-scores and raw recall scores for the focused-on, irrelevant, and alternative categories for Experiment 1 (*N* = 55)

	Overall recall			Recall by belief in stereotype about alternative category (<i>z</i> -scores)	
	Percent	Raw scores (out of 3)	<i>z</i> -Scores	Low belief	High belief
Type of word					
Focused-on	84	2.53	.82	.77	.87
Irrelevant	58	1.74	-.09	-.13	-.04
Alternative	41	1.22	-.64	-.86	-.42
Effect					
Facilitation	26	0.79	.91	.90	.91
Inhibition	17	0.52	.55	.73	.38

Note. Facilitation of focused-on set is focused-on minus irrelevant. Inhibition of alternative set is irrelevant minus alternative.

moderated by level of belief in either the focused-on or alternative stereotypes (both F 's < 1). Memory for the irrelevant sets was unrelated to belief in the associated stereotypes, average $r = .08$, all p 's > .30. Therefore, believing in a stereotype does not make stereotype-relevant traits easier to remember; however, it does make them harder to forget. This result suggests that if one strongly endorses a stereotype, then stereotype-relevant traits may surface in memory even if one tries to focus on another aspect of a target person's identity.

Measuring belief post-experimentally. Of course, because participants completed the stereotype belief questionnaire after the final recall test, it could be argued that recall increased stereotype belief rather than vice versa. The data do not support this interpretation. First, we need to reiterate that the relation between belief and recall emerged only for the alternative set. These words were not rehearsed during the retrieval-practice phase; therefore, any effect of recall on belief must have occurred at the final-test stage. If test-phase recall affected belief, then recall for the irrelevant sets should have been correlated with belief in the irrelevant stereotypes, a relationship that was notably absent, as discussed above. In addition, if recalling stereotype-relevant traits makes the stereotype seem more believable, then we would expect greater belief in sets when they are focused-on than when they are irrelevant because focused-on sets are better recalled during the test phase. In fact, however, participants reported marginally lower belief in the stereotype related to the focused-on set than the irrelevant set, $M_{diff} = -.21$, $t(54) = -1.75$, $p = .09$. Belief in the alternative stereotype did not differ from belief in the irrelevant stereotype, $t < 1$. Thus, it seems very likely that believing in the alternative set stereotype led to better recall (i.e., less RIF) of words from the alternative set, rather than vice versa.

Summary. Focusing on one stereotypic aspect of identity can inhibit access to another stereotypic aspect of identity. In the general discussion, we describe how this finding is consistent with other findings in the stereotype literature. In addition, we found that more strongly held stereotypes suffer less inhibition, consistent

with the notion that stereotypes provide cross-item integration, thereby supplying protection against RIF.

Experiment 2: Limited cues

The results of Experiment 1 suggest that when participants tried to recall words related to one aspect of a target's identity, words related to an alternative aspect of identity came to mind, and therefore required active inhibition, resulting in reduced memory for these competing traits at the final test stage. However, the design of our study may have artificially fostered the competition that occurred at retrieval between the focused-on set and the alternative set because we provided the compound-category cue at retrieval. We wanted participants to focus on a single aspect of the target's identity during the retrieval phase, so, for example, in the mother condition they were cued to retrieve all and only the mother stereotype words (along with filler words). Yet, when we cued them to retrieve, for example, supportive, they were given the cue FEMINIST-MOTHER SU__. Thus, feminist words may have competed with mother words at the retrieval stage because the word "feminist" was provided as part of the retrieval cue. If we had not reminded participants of the target's dual category membership during the retrieval practice phase, perhaps words related to the non-practiced (i.e., alternative) aspect of the target's identity would not have come to mind, and therefore would not have required inhibition.

In Experiment 2, we conducted a stricter test of our hypothesis by supplying participants with only the relevant sub-category (e.g., MOTHER) as the cue in the retrieval-practice phase. It is conceivable that participants would not think of words related to the alternative aspect of the target's identity if they were forced to rely on a single retrieval cue, thereby eliminating the need for inhibition. If, however, we found RIF using this design, we would have stronger evidence that trying to retrieve traits related to one aspect of identity requires the active inhibition of traits related to alternative aspects of identity. Therefore, participants in Experiment 2 were

given the full compound-cue during study (e.g., FEMINIST-MOTHER supportive) but only the relevant half of the cue during retrieval practice (e.g., MOTHER SU__). In order to provide continuity to the task and to be sure the participants would not be confused by the cue change (without having to explicitly forewarn them about it), we presented pictures of real people along with the cues: a picture of an Asian-American woman was accompanied by the label ASIAN-AMERICAN ARTIST and a picture of a woman nursing a baby was accompanied by the label FEMINIST MOTHER.

These particular compound categories illustrate an interesting property of real-life stereotype categories: some category memberships (e.g., race) are constantly visible while others (e.g., profession) may easily remain invisible. For each target, one sub-category membership was immediately visible (Asian-American, mother), whereas the other sub-category membership was not immediately visible (artist, feminist). In this way, we were able to examine whether retrieving words related to one aspect of identity inhibited words related to both visible and invisible alternative aspects of identity.

Method

Participants. Participants were 89 undergraduates, including 34 males and 55 females, who participated in partial fulfillment of a course requirement. Experimental sessions were run in groups of one to four participants.

Initial study phase. After signing a consent form, each participant was seated in front of a Macintosh 7300 computer. Using Psyscope software (Cohen, MacWhinney, Flatt, & Provost, 1993), the computer displayed a series of pictures of an Asian-American woman and a White mother, who was depicted holding a baby. During the study phase, each picture appeared six times along with an appropriate compound category description (ASIAN-AMERICAN ARTIST or FEMINIST MOTHER) and a single trait word, with a different trait shown each time. For example, participants saw a picture of an Asian-American woman, with the words “ASIAN-AMERICAN ARTIST: diligent” below it. We used the trait words from Experiment 1, except that we replaced *short* with *reserved* in the Asian-American category and *ugly* with *bitter* in the feminist category, in order to avoid combining physical trait descriptions with the pictures. The words stayed on the screen for 5 s; participants were instructed to spend this time relating the trait word to the target person.

We created fillers by displaying pictures of two chimpanzees (labeled WILD CHIMP and DOMESTICATED CHIMP) along with associated words (e.g., overeater). The chimpanzees appeared as the first three and last three targets and were interspersed throughout, allowing us to include at least two intervening items between repetitions of the same target person.

Retrieval-practice phase. The computer displayed the target’s picture and a sub-category label along with a word stem related to that sub-category (e.g., ASIAN-AMERICAN: DI__). Thus, in contrast to Experiment 1 in which participants saw the full compound category (e.g., ASIAN-AMERICAN ARTIST), participants were shown only the sub-category relevant to the traits they were asked to retrieve. In each condition, participants retrieved each of the three words three times. Participants were given 8 s to type out the appropriate trait word. The order of presentation met the same constraints as in Experiment 1.

Distracter phase. Same as Experiment 1.

Final test phase. The computer displayed each target’s picture and full compound category description, giving participants 30 s to list all the words originally studied with each target. Target order was counterbalanced, with the chimpanzees always appearing first and last.

Final questionnaire. Analogous to Experiment 1.

Results and discussion

Analyzing the recall data using *z*-scores as in Experiment 1, we found that participants exhibited better memory for words from the focused-on set ($M = .63$) than the irrelevant set ($M = -.11$), demonstrating facilitation, $F(1, 88) = 54.12$, $p < .001$, $\eta^2 = .38$. More importantly, participants exhibited worse memory for words from the alternative set ($M = -.41$) than the irrelevant set ($M = -.11$), demonstrating inhibition, $F(1, 88) = 10.34$, $p < .002$, $\eta^2 = .11$ (see Table 3). The inhibition effect was significant both when the alternative category was visible, $F(1, 43) = 6.36$, $p < .02$, $\eta^2 = .13$, and when the alternative category was invisible, $F(1, 44) = 3.93$, $p = .05$, $\eta^2 = .08$. Although the inhibition effect appeared slightly stronger when the alternative category was visible, the difference between the visible and invisible conditions was not significant, $F < 1$.

The inhibition effect was not significantly moderated by participants’ belief in the stereotype about the alternative or focused-on set, F ’s < 1 . Participants correctly retrieved 74% of words during the retrieval-practice phase, and participants retrieved a slightly higher percentage when they were asked to retrieve words related to a visible category (79%) than an invisible category (70%), $F(1, 87) = 2.64$, $p = .11$.

Summary. This experiment provided stronger evidence that remembering traits related to one aspect of a person’s identity impairs memory for traits related to another aspect of the person’s identity. Although participants were given only the relevant sub-category (e.g. ASIAN-AMERICAN) during the retrieval-practice phase, words from the alternative category (e.g., ARTIST) apparently came to mind, thereby requiring inhibition. Notably, this pattern of competition and inhibition

Table 3
Mean *z*-scores and raw recall scores for the focused-on, irrelevant, and alternative categories for Experiment 2 (*N* = 89)

	Overall recall			Recall by visibility of alternative category (<i>z</i> -scores)	
	Percent	Raw scores (out of 3)	<i>z</i> -Scores	Visible (<i>N</i> = 44)	Invisible (<i>N</i> = 45)
Type of word					
Focused-on	63	1.89	.63	.41	.85
Irrelevant	43	1.28	-.11	-.16	-.07
Alternative	34	1.03	-.41	-.53	-.30
Effect					
Facilitation	20	.61	.74	.57	.92
Inhibition	9	.25	.30	.37	.23

Note. Facilitation of focused-on set is focused-on minus irrelevant. Inhibition of alternative set is irrelevant minus alternative.

occurred even when the alternative category was not apparent from the target's picture.

Experiment 3: Individuating vs. stereotypic identities

In Experiments 1 and 2, both aspects of identity for each target were centered around a stereotype, whether racial, professional, political, or parental. As such, these experiments tell us that focusing on traits related to one stereotype can inhibit memory for traits related to another stereotype. But, of course, people are more than the sum of their stereotypes. As we get to know people we learn about their individuating characteristics. Would focusing on individuating information inhibit access to stereotypic information about that person? Would focusing on stereotypic information inhibit individuating information? Finally, would focusing on stereotypic information inhibit other, unpracticed information related to the same stereotype?

To address these questions we conducted a third experiment in which participants were randomly assigned to one of two study phase conditions. In the all-stereotype condition, participants studied traits about a target person that were all related to the same stereotype and then practiced retrieving half of these traits. In the mixed condition, participants studied both stereotypic traits ("S-traits") and individuating traits ("I-traits") about a target person and then practiced retrieving either the S-traits or I-traits.

The all-stereotype condition may be viewed as the condition most similar to previous experiments using typical cognitive categories because all the trait words are related to the target person through the same stereotype. As discussed in the introduction, stereotype belief may play one of two roles. If believing in a stereotype promotes stronger links between the group label and the associated traits, then strong belief should lead to *more* inhibition. If, however, belief in a stereotype promotes a more powerful integrative framework across traits, then strong belief should lead to *less* inhibition.

The all-stereotype condition allowed us to test these competing predictions.

In the same way that stereotype belief may affect how strongly associated a set of traits words are to each other (or to the group label), the nature of the words themselves may influence associative links. The mixed condition allowed us to test the prediction that studying individuating words would inhibit memory for stereotypic words, and vice versa—and allows us to compare the size of the inhibition effect. Because stereotypic words are more associated to the group label than individuating words, one might expect them to suffer more inhibition; however, because stereotypic words are more highly integrated with each other, one might expect them to suffer less inhibition. Again, our results should tell us something about the nature of the information that stereotypes provide.

Method

Pre-testing. Because this experiment required the use of more trait words than we had used in the previous experiments, we asked another 12 participants to rate the extent to which 72 words were stereotypically associated with Asian-Americans and mothers, using the 5-point scales employed in Experiment 1. Based on this data and the pre-testing data from Experiment 1, we selected six words that participants viewed as stereotypical of Asian-Americans ($M = 4.5$), six words viewed as stereotypical of mothers ($M = 4.8$), and 12 words that were unrelated to either stereotype ($M = 2.9$). See Table 4.

Participants. Participants were 149 undergraduates, including 50 males, 90 females, and 9 who did not list their gender. Experimental sessions were run in groups of one to four participants.

Initial study phase. The study phase was identical to Experiment 2, except that we used different trait words and did not provide explicit category labels; the picture of the Asian-American woman was simply labeled "JUNE," the picture of the mother was labeled "CHERYL," and the chimps were labeled "MIKEY" and "SCOUT." Participants were randomly assigned to

Table 4
Stereotypical and individuating words associated with each target person (Experiment 3)

Group	Type of word	
	Stereotypical	Individuating
Mothers (Cheryl)	patient	musical
	supportive	quirky
	worried	dreamy
	emotional	perky
	helpful	exciting
	caring	funny
Asian-Americans (June)	studious	wealthy
	intelligent	elegant
	reserved	mysterious
	diligent	liberal
	educated	witty
	precise	moody

one of two study conditions. In the all-stereotype condition, all six traits were related to the stereotype about each woman’s group. In the mixed condition, half of the traits were related to the group stereotype, whereas half were unrelated to the stereotype. Stereotypic and individuating traits alternated for each person.

To counterbalance for specific words, we created two versions of the mixed condition. Both versions used three S-traits and three I-traits to describe June and Cheryl, but the words themselves differed. Each version of the mixed condition used three of the words from the all-stereotype condition to describe each person, such that the same stereotypic words were represented in the mixed and all-stereotype conditions. For the mixed condition, we selected individuating words that were unrelated to the stereotype, but would allow participants to form a coherent impression of each target.

Retrieval-practice phase. For the all-stereotype study condition, we divided the stereotypic traits into two sets, which corresponded to the sets of stereotypic traits used

in the two versions of the mixed condition. Participants in the all-stereotype condition were randomly assigned to practice recalling one of these two sets of stereotypic traits associated with June or Cheryl. Similarly, participants in the mixed study condition were randomly assigned to rehearse either the S-traits or the I-traits associated with either June or Cheryl. The order, timing, and presentation used in the retrieval-practice phase were equivalent to Experiment 2.

Distraction phase. Same as Experiment 1.

Final test phase. Identical to Experiment 2, except that the targets’ names replaced their category descriptions.

Final questionnaire. Analogous to Experiment 1.

Results and discussion

Recall data. Overall, collapsing across study conditions, planned comparisons revealed that participants remembered more words from the focused-on set than the irrelevant set, $F(1, 148) = 74.32, p < .001, \eta^2 = .33$, and remembered fewer words from the alternative set than the irrelevant set, $F(1, 148) = 63.26, p < .001, \eta^2 = .30$, demonstrating facilitation and inhibition, respectively.

Participants in the all-stereotype condition demonstrated substantial RIF, $F(1, 72) = 34.81, p < .001, \eta^2 = .33$. (See Table 5.) In the mixed study condition, we found significant RIF regardless of which traits were rehearsed; when S-traits were rehearsed I-traits were inhibited, $F(1, 37) = 27.06, p < .001, \eta^2 = .42$, and when I-traits were rehearsed S-traits were inhibited, $F(1, 37) = 7.0, p < .01, \eta^2 = .16$. The RIF effect was not significantly moderated by the type of words participants rehearsed, although there was a slight trend suggesting that rehearsing S-traits inhibited I-traits more than vice versa, $F(1, 74) = 1.66, p = .20, \eta^2 = .022$.

Effects of stereotype belief. In the all-stereotype condition, to examine whether stereotype belief was

Table 5
Mean z-scores and raw recall scores for the focused-on, irrelevant, and alternative categories for Experiment 3 (N = 149)

Type of word	Overall recall			Recall by condition (z-scores)		
	Percent	Raw scores (out of 3)	z-scores	Mixed		All-stereotype
				Practice stereotype information	Practice individuating information	
Focused-on	62	1.87	.67	.56	.71	.71
Irrelevant	42	1.27	-.04	-.03	.04	-.09
Alternative	27	.81	-.60	-.64	-.33	-.71
Effect						
Facilitation	20	.60	.71	.59	.67	.80
Inhibition	15	.46	.56	.61	.37	.62

Note. Facilitation of focused-on set is focused-on minus irrelevant. Inhibition of alternative set is irrelevant minus alternative. N = 38 in each mixed condition and N = 73 in the all-stereotype condition.

associated with decreased inhibition, we conducted a regression using belief in the relevant stereotype to predict the size of the inhibition effect. This regression revealed that the more strongly participants believed in the relevant stereotype, the less RIF they showed, $F(1, 70) = 4.51, p = .04, \beta = -.25$. For example, participants who believed strongly in the stereotype about mothers and who practiced half of Cheryl's maternal S-traits showed relatively little RIF of her unpracticed maternal S-traits.

In the mixed study condition we observed a somewhat similar pattern. In this condition, participants who more strongly believed in the stereotype about the alternative set showed slightly less RIF for those traits, $F(1, 36) = 2.13, p = .15, \beta = -.24$. Although this trend was not statistically significant due to the reduced sample size, the effect size was consistent with the all-stereotype condition and with the results of Experiment 1. As in Experiment 1, there was no relationship between stereotype belief and recall of S-traits from the irrelevant set in either the all-stereotype (average $r = -.05$) or the mixed condition (average $r = -.08$).

Retrieval practice success rate. Overall, participants correctly retrieved 69% of words during the retrieval-practice phase. Participants in the all-stereotype condition retrieved a higher percentage of words (74%) than did participants in the mixed condition⁸ (64%), $F(1, 147) = 5.25, p = .02$. Controlling for retrieval success, however, did not substantively affect the results of any of the analyses discussed above.

Summary. This experiment suggests that rehearsing individuating information about a person may lead one to forget stereotypical information about that person, and vice versa. We also found that practicing stereotypic traits can inhibit memory for unpracticed traits related to the same stereotype; belief moderates this effect such that inhibition is weaker the more strongly participants believe in the stereotype.

General discussion

Our experiments provide evidence that rehearsing information related to one aspect of identity inhibits memory for information related to another aspect of identity. In Experiments 1 and 2, we found that rehearsing stereotypical information related to one aspect of a target's identity inhibited memory for information stereotypically associated with another aspect of the target's identity. This effect emerged regardless of whether participants were given the full compound category (Experiment 1) or only the relevant sub-category (Ex-

periment 2) as a cue during retrieval practice. In Experiment 3, we found that practicing stereotype-relevant traits can inhibit memory for other, unpracticed traits associated with the same stereotype. We also found that practicing stereotypical information about a target person inhibited memory for individuating information about the person, and that practicing individuating information about a person could inhibit stereotype information. In Experiments 1 and 3, the magnitude of RIF participants exhibited was related to their belief in the relevant stereotypes. Participants showed less RIF of words related to stereotypes that they more strongly endorsed. Stereotype belief did not significantly moderate RIF in Experiment 2. It is conceivable that the design of Experiment 2 may have reduced the strength of the relationship between belief in the alternative stereotype and memory for the associated words because the alternative category name was less salient during the retrieval phase, although confirming and explicating this possibility awaits further research.

Do stereotypes behave like typical cognitive categories in these studies? In an important sense, yes: The fact that inhibition was found at all means that when a person-trait pair was practiced, other traits related to the same person were activated and then inhibited, analogous to the process observed for standard categories. However, this basic cognitive process was influenced to some extent by participants' existing social beliefs.

The role of stereotype belief

In our studies, stereotype belief seems to play a role analogous to the role of cross-item integration in standard categories and disanalogous to the role of item strength. In the all-stereotype condition of Experiment 3, we found that the more participants believed in the stereotype, the less inhibition of the unpracticed traits they demonstrated. Yet, if stereotypic items were equivalent to strong items for high-believers and to weak items for low-believers, then we would expect high-believers to show *greater* inhibition of stereotypic words than low-believers (Anderson et al., 1994; also Bäuml, 1998). Thus, item strength fails to explain the relationship between stereotype belief and inhibition of stereotype-consistent words.

This relationship makes sense, however, if belief plays the role of integration. Anderson and McCulloch (1999) found that participants who integrate across the focused-on and alternative sets show less inhibition of the items from the alternative set compared to participants who do not integrate across these two sets. In the all-stereotype condition of Experiment 3, high belief in the relevant stereotype most likely provided an integrative framework connecting the focused-on and alternative sets, thereby reducing inhibition.

⁸ Within the mixed study condition, both participants who studied individuating traits and those who studied stereotypical traits correctly retrieved 64% of words.

We also found that stereotype belief was associated with reduced inhibition when looking at belief in the alternative set only; belief moderated RIF of a stereotype when a different stereotype was practiced (Experiment 1) and when individuating words were practiced (Experiment 3, mixed condition). We think that this result can also be understood in terms of integration, but of a slightly different sort; whereas previous studies have demonstrated that integrating items from the alternative set with items from the focused-on set can reduce RIF, our studies suggest that integrating items from the alternative set with other items from the alternative set may also reduce RIF. We argue that integrating alternative items in this way also promotes mediated retrieval, such that when one alternative item is recalled it is more likely to bring to mind other alternative items.

Why did this form of integration depend on stereotype belief, an individual difference variable, rather than stereotype awareness, which was probably fairly constant across our sample? Although most people may be aware of common cultural stereotypes, Lepore and Brown (1997) have demonstrated that the extent to which people automatically activate a stereotype after being exposed to the relevant category depends on their personal beliefs. Therefore, those participants who believe in the Asian-American stereotype may exhibit automatic activation of trait concepts like “intelligent” and “diligent” whenever they encounter Asian-Americans. The frequent co-activation of these trait concepts may not only strengthen the links between the group and the stereotypic traits, but also between the stereotypic traits themselves. In addition, these trait–trait links may be strengthened by individuals’ causal theories that the traits should co-occur. If believing a stereotype promotes the creation of strong trait–trait links, as this analysis suggests, then high-belief participants may show less inhibition of the alternative set due to these trait–trait links which may, as described above, promote mediated retrieval. This analysis would predict that stereotypic traits should be resistant to inhibition relative to less tightly integrated individuating traits; indeed, we found a slight trend in this direction. In sum, our results suggest that trait words will be relatively difficult to inhibit if they are integrated; that is, if they have been frequently co-activated in the past, as is the case with stereotype-relevant words for those who endorse the stereotype.

Why was there no relationship between stereotype belief and recall of stereotypic traits when these traits belonged to the focused-on or irrelevant categories? In the absence of inhibition, stereotypic traits should be relatively easy to remember, such that the extra activation provided by trait–trait connections would be trivial.

Moving beyond the domain of stereotypes, research on implicit personality theory suggests that traits that

frequently co-occur in individuals but are not inherently similar (e.g., submissive and clumsy; Ashmore, 1981) may show resistance to inhibition just as do words that co-occur because of their common association with a stereotype. Thus, in applying research on non-social categories to person perception, it is crucial to consider how pre-existing social schemas interact with general cognitive processes.

Social implications

Our results have some interesting social implications. The present findings dovetail with a handful of recent studies in suggesting that focusing on one aspect of a person’s identity can inhibit competing aspects of identity (e.g., Macrae et al., 1995). Bodenhausen and Macrae (1998) note that these inhibitory effects emerge as a result of automatic lateral inhibition, in which “the elements that become ‘irrelevant’ or ‘incompatible’ achieve this status by virtue of losing a low-level battle for processing priority” (p. 11). Retrieval-induced forgetting also represents a form of lateral inhibition; participants in our study forgot words from the alternative category not because they tried to suppress these words, but because they tried to remember other words from the focused-on set. Although lateral inhibition is automatically initiated when attention is drawn to one category rather than another, Bodenhausen and Macrae point out that motivation and beliefs may play an important role in determining which category is focused on (or activated) in the first place. Our results imply that stereotype belief also influences the extent to which a competing category is inhibited. Therefore, although lateral inhibition represents a relatively “cold” cognitive process, it too may be influenced by stereotype belief.

Lateral inhibition stands in contrast to more effortful, hierarchical inhibition, in which lower-level processes are reined in by higher-level goals, as in the case of intentional stereotype suppression. Although hierarchical inhibition has received a great deal of attention, its efficacy as a stereotype control strategy is dubious (e.g., Wyer et al., 2000). Rather than struggling to suppress a stereotype about one aspect of an individual’s identity, our research and other recent studies suggest that it may be possible to successfully inhibit an undesirable stereotype by focusing on another aspect of identity.

To illustrate, let us return to the example of the White college student’s perceptions of her African-American professor. Having attended at least a smattering of weekly lectures, the student must be fully cognizant of both the teacher’s racial category and professorial status. Now, suppose a friend asks the student to describe the professor, and the student characterizes the teacher in line with his professorial role. The results of our studies suggest that the process of retrieving information related to the teacher’s professional identity should

inhibit memory for information related to his racial identity. Our research also goes beyond most previous studies in providing evidence that focusing on a person's individuating characteristics can successfully inhibit memory for stereotypic information associated with the person's group membership. Thus, simply retrieving the professor's individuating characteristics ought to reduce the student's tendency to remember traits stereotypic of the professor's race. Our results do not paint a perfectly rosy picture, however. Stereotype-relevant traits seem slightly more difficult to inhibit than individuating traits, particularly for people who believe in the stereotype. Indeed, a target person's individuating traits may be easily overshadowed if perceivers retrieve stereotypic traits.

Our research does, however, contribute to a more nuanced understanding of the dynamics of stereotype activation and inhibition. For example, previous work on stereotype activation would suggest that participants who saw a picture of an Asian-American woman and repeatedly practiced traits stereotypic of her race should have easily recalled her other, unpracticed stereotypic traits, reflecting the strong activation of the stereotype. Yet, in our studies, practicing a woman's racially stereotypic traits made her other, unpracticed stereotypic traits come to mind less easily. This finding highlights the flexibility of stereotype activation and inhibition; rather than occurring inevitably, activation may ebb and wane in the service of perceptual goals. Indeed, the process of retrieval-induced forgetting itself represents a highly flexible process; RIF can be easily elicited (Macrae & MacLeod, 1999) and its effects are not very long-lasting (at least 20 min, but less than 24 h; MacLeod & Macrae, 2001), thereby ensuring that the perceiver's current goals are met. Thus, retrieval-induced forgetting may provide an effective cognitive mechanism that allows social perceivers to avoid remembering stereotypic information about others when the situation calls for them to recall information related to other aspects of identity.

References

- Anderson, M. C., & Bell, T. (2001). Forgetting our facts: The role of inhibitory processes in the loss of propositional knowledge. *Journal of Experimental Psychology: General*, *130*, 544–570.
- Anderson, M. C., Bjork, R. A., & Bjork, E. L. (1994). Remembering can cause forgetting: Retrieval dynamics in long-term memory. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *20*, 1063–1087.
- Anderson, M. C., Bjork, R. A., & Bjork, E. L. (2000). Retrieval-induced forgetting: Evidence for a recall-specific mechanism. *Psychonomic Bulletin & Review*, *7*, 522–530.
- Anderson, M. C., & McCulloch, K. C. (1999). Integration as a general boundary condition on retrieval-induced forgetting. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *25*(3), 608–629.
- Anderson, M. C., & Neely, J. H. (1996). Interference and inhibition in memory retrieval. In E. L. Bjork, & R. A. Bjork (Eds.), *Memory: Handbook of perception and cognition* (2nd ed., pp. 237–313). San Diego, CA: Academic Press.
- Anderson, M. C., & Spellman, B. A. (1995). On the status of inhibitory mechanisms in cognition: Memory retrieval as a model case. *Psychological Review*, *102*, 68–100.
- Ashmore, R. D. (1981). Sex stereotypes and implicit personality theory. In D. Hamilton (Ed.), *Cognitive processes in stereotyping and intergroup behavior* (pp. 37–82). Hillsdale, NJ: Erlbaum.
- Battig, W. F., & Montague, W. E. (1969). Category norms of verbal items in 56 categories. A replication and extension of the Connecticut category norms. *Journal of Experimental Psychology*, *80*, 1–46.
- Bäuml, K. (2002). Semantic generation can cause episodic forgetting. *Psychological Science*, *13*, 356–360.
- Bäuml, K. (1998). Strong items get suppressed, weak items do not: The role of item strength in output interference. *Psychonomic Bulletin & Review*, *5*, 459–463.
- Bäuml, K., & Hartinger, A. (2002). On the role of item similarity in retrieval-induced forgetting. *Memory*, *10*, 215–224.
- Bodenhausen, G. V., & Macrae, C. N. (1998). Stereotype activation and inhibition. In R. S. Wyer (Ed.), *Stereotype activation and inhibition. Advances in social cognition: Vol. 11* (pp. 1–52). Mahwah, NJ: Erlbaum.
- Cohen, J. D., MacWhinney, B., Flatt, M., & Provost, J. (1993). PsyScope: An interactive graphic system for designing and controlling experiments in the psychology laboratory using Macintosh computers. *Behavior Research Methods, Instruments, & Computers*, *25*, 257–271.
- Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality & Social Psychology*, *56*, 5–18.
- Dijksterhuis, A., & van Knippenberg, A. (1995). Timing of schema activation and memory; Inhibited access to inconsistent information. *European Journal of Social Psychology*, *12*, 383–390.
- Dijksterhuis, A., & van Knippenberg, A. (1996). The knife that cuts both ways: Facilitated and inhibited access to traits as a result of stereotype activation. *Journal of Experimental Social Psychology*, *32*, 271–288.
- Dijksterhuis, A., & van Knippenberg, A. (1998). Inhibition, Aberdeen, and other cloudy subjects. In R. S. Wyer (Ed.), *Stereotype activation and inhibition* (pp. 83–96). Mahwah, NJ: Erlbaum.
- Hugenberg, K., & Bodenhausen, G. (2001, February). *Inhibitory processes in self-categorization*. Poster presented at annual meeting of the Society for Personality and Social Psychology, San Antonio, TX.
- Kunda, Z., Miller, D. T., & Claire, T. (1990). Combining social concepts: The role of causal reasoning. *Cognitive Science*, *14*, 551–577.
- Lepore, L., & Brown, R. (1997). Category and stereotype activation: Is prejudice inevitable? *Journal of Personality & Social Psychology*, *72*, 275–287.
- Levy, B. J., & Anderson, M. C. (2002). Inhibitory processes and the control of memory retrieval. *Trends in Cognitive Science*, *6*, 299–305.
- MacLeod, M. D., & Macrae, C. N. (2001). Gone but not forgotten: The transient nature of retrieval-induced forgetting. *Psychological Science*, *12*(1), 148–152.
- 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 & Social Psychology*, *69*, 397–407.
- Macrae, C. N., & MacLeod, M. D. (1999). On recollections lost: When practice makes imperfect. *Journal of Personality & Social Psychology*, *77*, 463–473.
- Macrae, C. N., & Rosevare, T. A. (2002). I was always on my mind: The self and temporary forgetting. *Psychonomic Bulletin & Review*, *9*, 611–614.

- Rudman, L. A., & Borgida, E. (1995). The afterglow of construct accessibility: The behavioral consequences of priming men to view women as sexual objects. *Journal of Experimental Social Psychology*, 31, 493–517.
- Stroessner, S. J. (1996). Social categorization by race or sex: Effects of perceived non-normalcy on response times. *Social Cognition*, 14, 247–276.

- Tajfel, H., & Turner, J. C. (1986). The social identity theory of intergroup behavior. In S. Worchel, & W. G. Astin (Eds.), *Psychology of intergroup relations* (pp. 7–24). Chicago: Nelson-Hall.
- Wyer, N. A., Sherman, J. W., & Stroessner, S. J. (2000). The roles of motivation and ability in controlling the consequences of stereotype suppression. *Personality & Social Psychology Bulletin*, 26, 13–25.