



## Functional modularity in stereotype representation

Kimberly A. Quinn,\* Kurt Hugenberg,<sup>1</sup> and Galen V. Bodenhausen

*Department of Psychology, Northwestern University, 2029 Sheridan Rd., Evanston, IL 60208-2710, USA*

Received 22 January 2003; revised 28 October 2003

Available online 3 December 2003

### Abstract

We employed the retrieval-practice paradigm to test the hypothesis that stereotypes are organized in a meaningful, valence-based way that promotes evaluative coherence. Replicating previous research, we demonstrated that the rehearsal of traits known to describe a target person produced enhanced recall of those practiced traits and reduced recall of other known but non-practiced traits, relative to baseline. However, both the availability of a group label that united the traits within a stereotype and the evaluative consistency of the practiced and non-practiced traits moderated the nature of these effects: although recall of non-practiced stereotypic traits that were evaluatively inconsistent with the practiced traits showed the typical pattern of inhibition, recall of non-practiced stereotypic traits that were evaluatively consistent with the practiced traits was facilitated relative to baseline. We conclude by discussing how the modular representation implied by these findings is functional, potentially fostering the momentary experience of evaluative consistency in person perception.

© 2003 Elsevier Inc. All rights reserved.

*Keywords:* Person memory; Stereotype representation; Facilitation; Inhibition; Retrieval-induced forgetting; Evaluative coherence

### Introduction

Categorization is central to the process of social perception, and abundant evidence confirms that social perceivers activate and use generic information associated with a target's category membership in forming an impression of him or her (e.g., Allport, 1954; Brewer, 1988; Fiske & Neuberg, 1990; Macrae & Bodenhausen, 2000). A presumed advantage of this process is the rapid provision of a coherent, organized, meaningful representation of the target at the expense of relatively little mental effort (e.g., Bodenhausen, Macrae, & Sherman, 1999). This advantage would seem to require that stereotypes possess some kind of internal consistency such that their invocation provides a basis for clear and unequivocal impres-

sions. However, representations of social categories appear to commonly include ambivalent or evaluatively mixed stereotypic content (Fiske, Cuddy, Glick, & Xu, 2002). For example, some groups are considered to be quite competent or even prodigiously talented in culturally valued ways, yet they are also considered to be untrustworthy, dangerous, or socially disliked. Recognizing that a social target is likely to be complex and to have both good and bad qualities seems quite reasonable, but this kind of nuanced thinking is not what one commonly associates with the process of stereotyping. Moreover, the social perceiver's need for simple evaluative consistency has been the core premise of a wide variety of social psychological theories (e.g., Festinger, 1957; Heider, 1958); indeed, perceivers do preferentially form evaluatively coherent impressions of others (e.g., Asch, 1946; Bodenhausen & Macrae, 1994; Hampson, 1998; Srull & Wyer, 1989; Wyer & Srull, 1989). If stereotype representations are often inherently ambivalent, how can they assist in the formation of easily constructed, evaluatively coherent, well-structured social impressions?

Given that perceivers actively strive for evaluatively coherent impressions and yet rely on heterogeneous

\* Corresponding author. Present address: Department of Psychological and Brain Sciences, 6207 Moore Hall, Dartmouth College, Hanover, NH 03755, USA. Fax: 1-603-646-1419.

*E-mail addresses:* [kimberly.a.quinn@dartmouth.edu](mailto:kimberly.a.quinn@dartmouth.edu) (K.A. Quinn), [hugenbk@muohio.edu](mailto:hugenbk@muohio.edu) (K. Hugenberg), [galen@northwestern.edu](mailto:galen@northwestern.edu) (G.V. Bodenhausen).

<sup>1</sup> Present address: Department of Psychology, 136 Benton Hall, Miami University, Oxford, OH 45056, USA.

stereotypes, it would seem functional for stereotypes to be represented in a way that promotes evaluative coherence—namely, in terms of the valence of the stereotypic content. The goal of the current research was to investigate this possibility that stereotypes are organized along the lines of valence, with positive information being integrated with positive information and negative information integrated with negative information.

We chose the retrieval-practice paradigm (Anderson, Bjork, & Bjork, 1994; Anderson & Spellman, 1995) to test our hypothesis regarding the valenced representation of stereotypes. In this paradigm, participants are presented with lists of category–exemplar pairs and are asked to memorize the pairs (e.g., fruit–banana; animal–dog). Following this learning phase, participants practice retrieving a subset of the exemplars in a cued recall task. Following a distracter task, participants then complete a free-recall task in which they retrieve as many exemplars as possible, both practiced and non-practiced, for each category. This paradigm typically yields an effect known as retrieval-induced forgetting (RIF): although repeatedly retrieving or rehearsing factual information leads to enhanced recall of that material, it also leads to inhibited recall of non-practiced but related material (e.g., exemplars from the same category), relative to recall of unrelated material (e.g., non-practiced exemplars from a different category). Thus, repeatedly practicing the combination “fruit–banana” naturally leads to better recall of that category–exemplar pairing, but it also leads to substantially worse recall of other, non-practiced exemplars of the category “fruit” as compared to exemplars of an unrelated category. Because the unpracticed, but related material (e.g., fruit–apple) might otherwise interfere with rehearsal and retrieval of the focal, practiced material (e.g., fruit–banana), the mental system appears adept at arranging for its active inhibition. Not surprisingly, no such effect occurs for categories and exemplars that were previously learned, but are unrelated to the practiced pairs.

The utility of this paradigm for understanding representational structure was illustrated by Anderson and McCulloch (1999). They demonstrated that instructing participants to think about the relations between category exemplars during the learning phase reduced the magnitude of retrieval-induced forgetting; that is, non-practiced information that was conceptually integrated with the practiced information escaped the normally observed inhibition. Anderson and McCulloch suggested that complex knowledge structures composed of highly interconnected components might be especially resistant to inhibition. Stereotypes may be one such structure: although the precise nature of stereotype representation continues to be debated, few would disagree that stereotypes facilitate the integration of social information (Bodenhausen & Macrae, 1998) and that such integration can afford great flexibility in reinterpreting features to

correspond with other information known to be true of targets. If it is the case that stereotypes facilitate such integration, and if this integration reduces the effects of inhibition (Anderson & McCulloch, 1999), then when an activated stereotype guides information integration, such information may be relatively impervious to inhibition. As previously discussed, however, the perceiver’s desire for evaluative consistency may place constraints on this process. Given that perceivers strive for and prefer evaluatively consistent impressions of others (e.g., Bodenhausen & Macrae, 1994), impressions may typically be integrated around an evaluatively coherent subset of the broader stereotype.

More recently, Dunn and Spellman (2003) used the retrieval-practice paradigm to explore inhibitory mechanisms in social information processing (see also Macrae & MacLeod, 1999). Specifically, they investigated how social information is processed when a target is multiply categorizable (e.g., a woman who is both a feminist and a mother). They demonstrated that focusing on one particular categorization (e.g., “feminist”) leads to the inhibition of another categorization (e.g., “mother”). (For similar findings with different paradigms, see Hugenberg & Bodenhausen, *in press*; Macrae, Bodenhausen, & Milne, 1995.) Importantly, the pattern was moderated by the extent to which the perceiver endorsed the target stereotypes: participants who believed strongly in the identity-relevant stereotype demonstrated less inhibition of non-practiced stereotypic information related to that identity. These findings demonstrate that the strength of association between components of a stereotype, as in the case of non-social concepts, has implications for how inhibition plays out. That is, the degree to which features within a category representation are closely linked can constrain the extent to which lateral inhibition is possible.

#### *Overview of the current research*

The evidence provided by Anderson and McCulloch (1999) and Dunn and Spellman (2003) demonstrates that variations in representational structure produce variations in the extent of inhibitory processing. This contingency suggests that one can use the magnitude of inhibition effects—or their presence or absence—to draw inferences about the representational structure of concepts of interest. Our goal was to use the retrieval-practice paradigm to investigate the possibility that stereotype representation is organized in a meaningful way that promotes evaluative coherence. More specifically, we sought to determine whether patterns of facilitation and inhibition in social information processing would be moderated by whether the information to be retrieved from memory was evaluatively consistent versus inconsistent with the perceiver’s ongoing focus of attention. We expected that information that becomes the focus of the perceiver’s attention would necessarily

be facilitated in memory. Whether non-focal information is facilitated or inhibited, however, should be moderated by its relation to the focal information.

To test how the relation between focal and non-focal information influences social information processing, we compared recall following retrieval practice in the presence or absence of an applicable stereotype. We predicted that when no stereotype had been activated about a social target, all non-focal information about that target would be inhibited (i.e., the typical RIF effect). The reasoning was that, in the absence of the organizational framework of a stereotype, non-focal information would not be systematically integrated with the focal information and so could interfere with participants' ability to maintain a focus of attention; inhibition would be required to counter this interference.

We expected a different pattern to emerge, however, when the target information could be organized according to an activated stereotype. We hypothesized that evaluatively inconsistent information, because it would interfere with the perceiver's ability to maintain a focus of attention, would be inhibited in memory (i.e., a typical RIF effect), but that evaluatively consistent information, because it would pose no interference with ongoing attention, would not be inhibited (i.e., an absence of RIF). If, in fact, stereotypic information is integrated and organized along lines of valence, evaluatively consistent unpracticed information may not only be resistant to inhibition, but its recall may actually be facilitated by retrieval practice of related information (i.e., a reversal of the typical RIF effect). When interacting with a feminist, for example, a focus on her determination may subsequently produce both enhanced recall of other positive stereotypic attributes such as "confident" and diminished recall of negative stereotypic attributes such as "angry."

We expected that recall of non-stereotypic information would be uninfluenced by retrieval practice when a stereotype was activated. By virtue of the fact that it is unrelated to the activated mental representation, non-stereotypic information does not compete with focal information and thus presumably should be less vulnerable to lateral inhibition.

## Method

### *Participants*

Participants were 128 undergraduates from Northwestern University who completed the experiment for partial course credit.

### *Procedure*

Participants were recruited for sessions of up to four people, with the entire experiment being administered

on computers housed in individual cubicles, using MediaLab research software.

Participants learned that they would complete a study of "social memory." The computerized instructions indicated that they would be given the names of two target persons (in all cases, David and Susan), along with their self-descriptions and various trait descriptions that had been collected from people who know the target well (e.g., friends, coworkers). Participants then read instructions that included a labeling manipulation, such that participants learned the group membership of one target but not the other. One half of the participants read that David "describes himself as an athlete" and that Susan "says she'd rather let her friends describe her than label herself." The other half of the participants read that Susan described herself as a feminist, but that David did not label himself. This manipulation was designed to activate the relevant stereotype for one target but not for the other.

Participants were told that the first part of the study was a learning phase, in which they would be presented with name-trait pairs. They learned that they should try to remember each trait by relating it to the name of the target with which it was associated. The computer then presented the name-trait pairs in randomized order (e.g., "Susan—determined"), with each pair appearing on the computer screen for 5 s. A total of 24 traits were presented with each target name, including 6 positive stereotypic traits, 6 negative stereotypic traits, 6 positive non-stereotypic traits, and 6 negative non-stereotypic traits. Earlier pilot testing was used to collect lists of traits believed to characterize a variety of social groups, including feminists and athletes. The traits were rated by pilot participants for typicality, and were coded by the experimenters for valence extremity according to the norms provided by Anderson (1968). In the final stimulus set, efforts were made to equate the average stereotypicality of the stereotype-relevant items, and the valence extremity of the positive and negative items. The stimulus items are presented in Table 1.

Participants then completed the practice task, which was presented to them as a memory task. In this task, the computer presented a target name, along with the first two to three letters of an applicable trait (e.g., "Susan—de\_\_\_\_\_"), and participants' task was to recall and type a trait word that they had seen in the learning phase that would complete the word fragment. Participants completed one of four randomly assigned practice tasks: one quarter of participants practiced retrieving positive traits stereotypic of feminists, one quarter practiced retrieving negative traits stereotypic of feminists, one quarter practiced retrieving positive traits stereotypic of athletes, and one quarter of participants practiced retrieving negative traits stereotypic of athletes. In each case, participants were provided with three word fragments from among the six possible trait words.

Table 1  
Stimulus materials

Target	Stereotypic traits		Non-stereotypic traits	
	Positive	Negative	Positive	Negative
Susan	Confident	Aggressive	Casual	Average
	Determined	Angry	Courteous	Forgetful
	Independent	Defensive	Enthusiastic	Gullible
	Liberal	Loud	Hopeful	Nosy
	Outspoken	Masculine	Outgoing	Picky
	Radical	Opinionated	Pleasant	Shallow
David	Active	Arrogant	Agreeable	Disorganized
	Competitive	Irresponsible	Easygoing	Finicky
	Dedicated	Lazy	Gracious	Gossipy
	Popular	Oafish	Optimistic	Naïve
	Social	Stupid	Positive	Normal
	Talented	Wild	Talkative	Superficial

Participants were randomly assigned to practice either the first or final three words of the appropriate list. Following Anderson and Spellman (1995), each word fragment was presented three times, for a total of nine practice trials. It is important to note that one half of the participants practiced retrieving traits relevant to the target whose stereotype label was known; the remainder practiced retrieving traits relevant to the target whose stereotype label was not known.

Following the practice task, participants completed a 5-min distracter task, presented to them as an unrelated experiment.<sup>2</sup> Participants then completed a surprise recall task, in which they were told that they would have a total of 2 min to recall as many traits as possible for both targets. The recall task comprised two blocks; in each, participants were presented with one of the target names (i.e., Susan or David) and reported as many trait words as they could recall. After completing the first block, participants proceeded to complete the second block. Order of presentation of the target names was randomized across participants.

Following completion of the recall task, participants were probed for suspicion, debriefed, and thanked for their participation.

## Results

### *Coding of recall data*

Items reported by participants were coded as correct if they appeared in the original stimulus set and were attributed to the appropriate target; synonyms, intrusions, and items attributed to the wrong target were coded as incorrect.

Each correctly recalled item was coded as falling into one of nine categories, summarized in Table 2. Stereotypic items about the primary target (P) that had previously been practiced were designated P+ items. For the purposes of coding, items consistent with stereotypes about a target's social group were designated as stereotypic regardless of whether or not the target's category label had been activated. Evaluative consistency was defined relative to the valence of the P+ items. Thus, there were four additional types of traits that could be recalled about P: stereotypic, valence-consistent traits (P-SC), stereotypic, valence-inconsistent traits (P-SI), non-stereotypic, valence-consistent traits (P-NC), and non-stereotypic, valence-inconsistent traits (P-NI). Items recalled about the other, control target (O) were designated as stereotypic or non-stereotypic relative to O's social category, but the valence of the items was designated as consistent or inconsistent relative to the valence of the practiced items about P. Thus, there were four possible types of traits for O, summarized in Table 2.

The data were converted to proportions, to account for differences in the number of responses possible within each coding category.

### *Analytic strategy*

We collapsed across the target label (feminist/athlete) and practice target (feminist/athlete) factors to create the theoretically critical conditions where the primary target's stereotype label was available (practice-labeled condition) or not (practice-unlabeled condition). (Note that in the latter condition, the other target's stereotype label was available.)

### *Retrieval-practice effects*

Practicing certain items should facilitate their recall. Thus, we should observe better recall of practiced (P+)

<sup>2</sup> The filler task was a prospective memory task related to neither the feminist nor the athlete stereotype.

Table 2  
Coding categories for recalled traits

<i>Traits describing the primary target (P)</i>	
P+	Practiced stereotypic traits
P-SC	Non-practiced stereotypic items consistent in valence with the P+ traits
P-SI	Non-practiced stereotypic items inconsistent in valence with the P+ traits
P-NC	Non-practiced non-stereotypic items consistent in valence with the P+ traits
P-NI	Non-practiced non-stereotypic items inconsistent in valence with the P+ traits
<i>Traits describing the control target (O)</i>	
(O's traits were never practiced)	
O-SC	Stereotypic items consistent in valence with the P+ traits
O-SI	Stereotypic items inconsistent in valence with the P+ traits
O-NC	Non-stereotypic items consistent in valence with the P+ traits
O-NI	Non-stereotypic items inconsistent in valence with the P+ traits

stereotypic items than unpracticed control stereotypic items of the same valence (O-SC). Indeed, this was the case,  $M_s = .49$  versus  $.30$ , respectively. A  $2$  (practice target: labeled versus unlabeled)  $\times 2$  (practice valence: positive versus negative)  $\times 2$  (trait type: primary versus control) mixed-model analysis of variance (ANOVA), with practice target and practice valence as between-participants factors and trait type as a within-participants factor, confirmed the reliability of this main effect for trait type,  $F(1, 124) = 21.77, p < .001$ .

The analysis also yielded an unexpected main effect for practice valence,  $F(1, 124) = 5.71, p = .018$ , suggesting that recall was better following retrieval practice for negative ( $M = .415$ ) than for positive ( $M = .328$ ) information. This preliminary analysis yielded no other main effects or interactions.

#### *Retrieval-induced forgetting*

Our key hypotheses centered on the recall of unpracticed items. We hypothesized that retrieval practice *without* stereotype activation would lead to the inhibition of *all* non-practiced information about the primary target (relative to the control target), in keeping with the retrieval-induced forgetting literature. However, when a stereotype was activated about the primary target, we expected not to observe inhibition of unpracticed stereotypic traits, provided that those traits are evaluatively consistent with the practiced traits (i.e., P-SC items). On the contrary, recall of such traits should actually be facilitated by practicing the other evaluatively similar stereotypic traits. Consistent with the pattern hypothesized for recall in the absence of stereotype activation, recall of evaluatively inconsistent stereotypic traits in the presence of stereotype activation was expected to show the retrieval-induced inhibition pattern. Recall of non-stereotypic traits, however, was expected to be neither facilitated nor inhibited when a stereotype was activated: by virtue of being unrelated to the stereotype, non-stereotypic traits should not be integrated with the stereotyped and thus should receive no facilitation; by virtue of being irrelevant to the stereotype,

they should provide no competition for the stereotype and thus not be prone to inhibition.

To test our hypotheses, we subjected the recall data to a  $2$  (practice target: labeled versus unlabeled)  $\times 2$  (practice valence: positive versus negative)  $\times 2$  (recall target: primary versus control)  $\times 2$  (trait type: stereotypic versus non-stereotypic)  $\times 2$  (trait valence: practice-consistent versus -inconsistent) ANOVA, with practice target and practice valence as between-participants factors, and all others as within-participants factors.

The analysis yielded a number of main effects and higher-order interactions, which were subsumed within the predicted Practice Target  $\times$  Recall Target  $\times$  Trait Type  $\times$  Trait Valence,  $F(1, 124) = 13.51, p < .001$ .<sup>3</sup> Relevant means are presented in Table 3. To facilitate interpretation of this four-way interaction, we examined separately the Recall Target  $\times$  Trait Type  $\times$  Trait Valence interactions for the practice-unlabeled and practice-labeled conditions.

#### *Recall in the absence of stereotype activation*

We anticipated that when no stereotype was activated during the practice phase, the results would replicate the standard retrieval-induced forgetting phenomenon. That is, we expected recall for the non-practiced traits of the primary target (i.e., all P-traits) to be inhibited, relative to recall for the (non-practiced) traits of the control target (i.e., all O-traits). The within-participants ANOVA for the practice-target unlabeled condition did indeed yield a main effect for target,  $F(1, 62) = 9.85, p = .003$ . Confirming our prediction, recall of the primary target's unpracticed traits ( $M = .19$ ) was reliably

<sup>3</sup> The analysis yielded one higher-order interaction that was not subsumed within the predicted interaction. More specifically, the analysis yielded a Practice Valence  $\times$  Target  $\times$  Trait Valence interaction,  $F(1, 124) = 4.54, p = .035$ . The interaction was neither theoretically meaningful nor readily interpretable: examination of the means indicated that participants in the positive retrieval-practice condition (regardless of label-practice condition) later recalled more negative non-stereotypic traits than did participants in the negative retrieval-practice condition.

Table 3

Mean proportion recall of non-practiced trait items as a function of label practice, target, trait type, and trait valence

	Stereotypic traits (S)		Non-stereotypic traits (N)	
	Primary target (P-)	Other (O)	Primary target (P-)	Other (O)
<i>Practice unlabeled</i>				
Valence-consistent (C)	0.22 (.026)	0.31 (.029)	0.18 (.024)	0.16 (.020)
Valence-inconsistent (I)	0.23 (.028)	0.28 (.024)	0.12 (.014)	0.19 (.022)
<i>Practice labeled</i>				
Valence-consistent (C)	0.42 (.037)	0.29 (.026)	0.17 (.020)	0.20 (.021)
Valence-inconsistent (I)	0.23 (.023)	0.31 (.029)	0.16 (.020)	0.20 (.022)

Note. Standard errors are presented in parentheses.

lower than recall of the control target's unpracticed traits ( $M = .24$ ).

A number of other effects also emerged, however, including a marginal Recall Target  $\times$  Trait Type  $\times$  Trait Valence interaction,  $F(1, 62) = 3.90$ ,  $p = .053$ . We conducted post hoc contrasts to understand the nature of this interaction (depicted in the top half of Table 3). These comparisons were trait- and valence-matched; for example, recall for non-practiced valence-consistent stereotypic trait information about the primary target (P-SC) was compared to recall for non-practiced valence-consistent stereotypic information about the control target (O-SC), and so on.

When the means for the practice-target unlabeled condition were examined individually, the following effects emerged: as predicted, recall for non-practiced stereotypic traits generally favored the control (O) over primary (P-) target; this pattern was reliable for O-SC versus P-SC trait recall,  $t(124) = 2.61$ ,  $p = .01$ , and a trend for O-SI versus P-SI trait recall,  $t(124) = 1.61$ ,  $p = .11$ . Recall for valence-inconsistent non-stereotypic information was also reliably lower for the primary versus control target,  $t(124) = 2.29$ ,  $p = .02$ , as predicted. Recall for valence-consistent non-stereotypic information, however, did not differ reliably as a function of primary/control status,  $p = .60$ .

Thus, when no stereotype was available to organize incoming information, practicing some information tended to produce the inhibition of other information about the practice target, regardless of its consistency with the valence of the practiced traits.

#### Recall following stereotype activation

In the practice-target labeled condition (bottom half of Table 3), the within-participants analysis yielded the predicted Recall Target  $\times$  Trait Type  $\times$  Trait Valence interaction,  $F(1, 62) = 10.08$ ,  $p = .002$ . We conducted planned comparisons to compare, on a within-participants basis, recall for practiced-target versus non-practiced-target information. Again, to provide the most stringent test of our hypotheses, we conducted comparisons that were trait- and valence-matched.

Our key prediction was that, when a stereotype had been activated, the standard retrieval-induced forgetting

effect would not be obtained, and in fact would be reversed, when unpracticed traits were stereotypic and consistent in valence with the practiced stereotypic items. As can be readily seen in Table 3, this was indeed the case; recall of P-SC items was significantly greater than recall of O-SC items,  $t(124) = 4.22$ ,  $p < .001$ , showing a *facilitation* effect of practicing one set of stereotypic traits on recall of other, evaluatively congruent stereotypic traits. Also as predicted, for stereotypic traits that were evaluatively incongruent with the practiced traits, the typical retrieval-induced inhibition effect was evident. That is, P-SI  $<$  O-SI,  $t(124) = -2.58$ ,  $p = .011$ . Recall of the non-stereotypic items did not differ as a function of whether they pertained to the primary or control target,  $ps > .23$ .

Thus, as hypothesized, when a stereotype is available, practicing stereotypic traits facilitates or inhibits the retrieval of other, non-practiced stereotypic traits, depending on the non-practiced traits' evaluative consistency with the practiced set. When a stereotype label was available, practicing certain items improved retrieval not only of those practiced items but also of other non-practiced stereotypic items that were similar in valence; however, retrieval of non-practiced stereotypic items that were *opposite* in valence was inhibited. The a priori organization of the relevant stereotype (of feminists or athletes), when activated, facilitated access to other unpracticed stereotypic content, but only if this content was evaluatively consistent with whichever aspect of the relevant stereotype had been practiced. If the unpracticed stereotypic content was evaluatively inconsistent with the practiced stereotypic information, that unpracticed stereotypic content was inhibited.

#### Additional analyses: Inhibition versus output interference

A possible reinterpretation of our effects is that they are not attributable to retrieval-induced forgetting, but rather to output interference at recall. The first items retrieved during recall can interfere with the retrieval of related material from memory (Roediger & Schmidt, 1980; Tulving & Arbuckle, 1963). In the present research, it is conceivable that following the period of retrieval practice, participants were more able to retrieve

the highly accessible practiced items from memory, producing interference with the retrieval of the related items. If this were the case, then retrieval-induced forgetting should only emerge among participants who first retrieved P+ items during the free-recall task.

To examine this possibility, we conducted two additional analyses. We coded participants' data to reflect the extent to which they began their recall sequences with P+ or P-SC items. We coded the recall positions of these items and calculated the difference between the average recall positions of the two item types. This procedure yielded a difference score with negative values reflecting earlier PSC output and positive values reflecting earlier P+ output; we performed a median split on these scores to categorize participants as early P-SC versus early P+. The inhibition effects for the two groups were calculated in the normal way (i.e., P-SC minus O-SC) and were subjected to a 2 (practice target: labeled versus unlabeled)  $\times$  2 (output order: early P-SC versus early P+) between-participants ANOVA. We also conducted the same analysis with P+ versus P-SI items. Neither analysis yielded significant main effects or interactions for output order ( $ps > .20$ ), indicating that the inhibition effects cannot plausibly be attributed to output interference during the final recall task.

## Discussion

Although focusing on a subset of trait information about a target leads to the inhibition of other trait information about the person, the availability of a social group label that serves to unite the traits within a coherent stereotype moderated the nature of these effects. The typical retrieval-practice effects were replicated when participants did *not* receive a stereotype label. With the addition of a stereotypic label, however, recall of stereotypic traits that were unpracticed and evaluatively inconsistent with the practiced traits was inhibited relative to baseline, whereas recall of unpracticed stereotypic traits that were evaluatively consistent with the practiced traits was facilitated.

Might these effects be attributable to simple affective priming rather than valence-based facilitation and inhibition? The data do not support this interpretation. There was no facilitation of affectively congruent non-stereotypic information about the primary target individual or of affectively congruent stereotypic or non-stereotypic information about the control individual. In general, participants demonstrated similar levels of recall for practice-valence-consistent and -inconsistent non-stereotypic items for both the primary and control target, and practice-valence-consistent and -inconsistent stereotypic items for the control target. Moreover, examination of the means on a target-by-target basis indicated that, regardless of the valence of the focal

retrieval practice traits, recall of non-stereotypic information for both primary and control targets tended to favor negative information, replicating the typical negativity bias in social perception (for a review, see Baumeister, Bratslavsky, Finkenauer, & Vohs, 2001).

### *Implications for stereotype representation and person perception*

To the best of our knowledge, this study is the first to demonstrate that facilitation and inhibition can proceed separately for evaluatively conflicting aspects of a stereotype. (For evidence of differential facilitation of positive and negative components of a stereotype, see Wittenbrink, Judd, & Park, 2001.) The existence of two processes that operate within a category representation is noteworthy, because it suggests that stereotype representation includes not only the features believed to characterize the social group in question, but also important patterns of connections among those features—in this case, stronger associations between evaluatively consistent rather than inconsistent features.

With social category information organized in such a way, inhibitory mechanisms may be able to provide social perceivers with the situational experience of evaluative coherence even when the broader stereotype contains both positive and negative components. For example, in situations where social perceivers recall some positive stereotypic traits about a social target, a stereotype may facilitate the recall of other positive stereotypic traits, while affording the capacity to inhibit otherwise evaluatively incongruent, but still stereotypic, negative traits. These processes have at least two implications for the social perceiver. First, the simultaneous operation of facilitative and inhibitory processes may serve to prevent the experience of ambivalence-related discomfort: by heightening the activation of both focal information and information that is evaluatively consistent with that focal information, and by dampening the activation of information that contradicts the evaluative significance of focal information, the cognitive system prevents inconsistent information from competing with focal information for access to consciousness and upsetting the social perceiver's drive for consistency. Indeed, in recent years, theorists in cognitive and neuropsychological domains have emphasized the importance of inhibitory mechanisms for maintaining stability and focus in the cognitive system (e.g., Anderson & Spellman, 1995; Houghton & Tipper, 1996; Zacks & Hasher, 1994). As previously discussed, lateral inhibition of this sort would be particularly useful as it may serve to prevent social perceivers from experiencing the inconsistency of their beliefs about the social group in question.

Equally important, the processes of the sort documented here enable the social perceiver to respond

flexibility to situational vagaries with little effort. The negative affective implications of ambivalence aside, the social perceiver could potentially encounter difficulty in deciding how to react to social targets if both positive and negative information were competing to influence judgment and behavior. The operation of non-conscious facilitation and inhibition eliminates this potential for response conflict. When situational factors draw the perceiver's attention to a subset of a target's positive stereotypic attributes, for example, stereotypes integrated along the lines of valence will facilitate the recall of other, positively valenced stereotypic information. Simultaneously, lateral inhibition may also act to reduce the likelihood that negative stereotypic information will compete with the positive information, thus reducing interference with the ongoing flow of information processing and behavior.

While social perceivers undoubtedly are able to step back occasionally and reflect upon the ambivalent stereotypes that are attached to various social categories in the abstract, in dealing with specific members of a category, they possess the cognitive wherewithal to sidestep this evaluative incoherence by activating evaluatively consistent subsets of stereotypic knowledge. The mechanisms we have emphasized can augment other theoretical approaches to coping with ambivalence. For example, research on racial ambivalence has provided evidence for response amplification, or a tendency to reduce perceived inconsistency by generating more extreme (positive or negative) responses to racial outgroup members than to racial ingroup members (e.g., Katz, Wackenhut, & Hass, 1986). Response amplification is thought to arise, at least in part, because contextual factors can provide a "boost" to either positive or negative information in the response competition process, allowing one subset of beliefs about a social group to dominate the perceiver's thoughts and behaviors. Our research highlights mechanisms by which this response competition might be resolved. Specifically, stereotype relevant information that is valence-consistent with the focal information is facilitated or "boosted" in the perceiver's memory, while valence-inconsistent stereotypic information may be inhibited. A valence-based, modular representation system of this sort could easily lead to an amplification of experienced positive or negative affect and naturally, an amplification of observed positive or negative behavior. Further research investigating how stereotypes are organized, activated and inhibited promises to illuminate many aspects of person perception and interpersonal behavior.

## References

- Allport, G. W. (1954). *The nature of prejudice*. Cambridge: Addison-Wesley.
- Anderson, N. H. (1968). Likeableness ratings of 555 personality-trait words. *Journal of Personality and Social Psychology*, 9, 272–279.
- 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, and Cognition*, 20, 1063–1087.
- Anderson, M. C., & McCulloch, K. C. (1999). Integration as a general boundary condition on retrieval-induced forgetting. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25, 608–629.
- 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.
- Asch, S. E. (1946). Forming impressions of personality. *Journal of Abnormal and Social Psychology*, 41, 258–290.
- Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad is stronger than good. *Review of General Psychology*, 5, 323–370.
- Bodenhausen, G. V., & Macrae, C. N. (1994). Coherence versus ambivalence in cognitive representations of persons. In R. S. Wyer (Ed.), *Advances in social cognition: Associated systems theory: A systematic approach to cognitive representations of persons* (Vol. 7, pp. 149–156). Hillsdale, NJ: Erlbaum.
- Bodenhausen, G. V., & Macrae, C. N. (1998). Stereotype activation and inhibition. In R. S. Wyer (Ed.), *Advances in social cognition: Stereotype activation and inhibition* (Vol. 11, pp. 1–52). Mahwah, NJ: Erlbaum.
- Bodenhausen, G. V., Macrae, C. N., & Sherman, J. W. (1999). On the dialectics of discrimination: Dual processes in social stereotyping. In S. Chaiken & Y. Trope (Eds.), *Dual process theories in social psychology* (pp. 271–290). New York: Guilford Press.
- Brewer, M. B. (1988). A dual process model of impression formation. In T. Srull & R. Wyer (Eds.), *Advances in social cognition* (Vol. 1, pp. 1–36). Hillsdale, NJ: Erlbaum.
- Dunn, E., & Spellman, B. A. (2003). Forgetting by remembering: Stereotype inhibition through rehearsal of alternative aspects of identity. *Journal of Experimental Social Psychology*, 39, 420–433.
- Festinger, L. (1957). *A theory of cognitive dissonance*. Oxford, UK: Row & Peterson.
- Fiske, S. L., Cuddy, A. J. C., Glick, P., & Xu, J. (2002). A model of (often mixed) stereotype content: Competence and warmth respectively follow from perceived status and competition. *Journal of Personality and Social Psychology*, 82, 878–902.
- Fiske, S. T., & Neuberg, S. L. (1990). A continuum model of impression formation, from category-based to individuating processes: Influence of information and motivation on attention and interpretation. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 23, pp. 1–74). New York: Academic Press.
- Hampson, S. E. (1998). When is an inconsistency not an inconsistency? Trait reconciliation in personality description and impression formation. *Journal of Personality and Social Psychology*, 74, 102–117.
- Heider, F. (1958). *The psychology of interpersonal relationships*. New York: Wiley.
- Houghton, G., & Tipper, S. P. (1996). Inhibitory mechanisms of neural and cognitive control: Applications to selective attention and sequential action. *Brain and Cognition*, 30, 20–43.
- Hugenberg, K., & Bodenhausen, G. V. (in press). Category membership moderates the inhibition of social identities. *Journal of Experimental Social Psychology*.
- Katz, I., Wackenhut, J., & Hass, R. G. (1986). Racial ambivalence, value duality, and behavior. In J. F. Dovidio & S. L. Gaertner (Eds.), *Prejudice, discrimination, and racism* (pp. 35–60). San Diego: Academic Press.
- Macrae, C. N., & Bodenhausen, G. V. (2000). Social cognition: Thinking categorically about others. *Annual Review of Psychology*, 51, 93–120.

- Macrae, C. N., Bodenhausen, G. V., & Milne, A. B. (1995). The dissection of selection in person perception: Inhibitory processes in social stereotyping. *Journal of Personality and Social Psychology*, *69*, 397–407.
- Macrae, C. N., & MacLeod, M. D. (1999). On recollections lost: When practice makes imperfect. *Journal of Personality and Social Psychology*, *77*, 463–473.
- Roediger III, H. L., & Schmidt, S. R. (1980). Output interference in the recall of categorized and paired associates lists. *Journal of Experimental Psychology: Human Learning and Memory*, *3*, 174–188.
- Srull, T. K., & Wyer, R. S., Jr. (1989). Person memory and judgment. *Psychological Review*, *96*, 58–83.
- Tulving, E., & Arbuckle, T. Y. (1963). Input and output interference in short-term associative memory. *Journal of Experimental Psychology*, *72*, 89–104.
- Wittenbrink, B., Judd, C. M., & Park, B. (2001). Spontaneous prejudice in context: Variability in automatically activated attitudes. *Journal of Personality and Social Psychology*, *81*, 815–827.
- Wyer, R. S., Jr., & Srull, T. K. (1989). *Memory and cognition in its social context*. Hillsdale, NJ: Erlbaum.
- Zacks, R. T., & Hasher, L. (1994). Directed ignoring: Inhibitory regulation of working memory. In D. Dagenbach & T. H. Carr (Eds.), *Inhibitory processes in attention, memory, and language* (pp. 241–264). San Diego, CA: Academic Press.