

UC Davis

UC Davis Previously Published Works

Title

On the Roles of Stereotype Activation and Application in Diminishing Implicit Bias

Permalink

<https://escholarship.org/uc/item/47s62424>

Authors

Rivers, Andrew M

Sherman, Jeff

Rees, Heather

et al.

Publication Date

2018-01-19

DOI

10.31234/osf.io/7w8yz

Peer reviewed

On the Roles of Stereotype Activation and Application in Diminishing Implicit Bias

Andrew M. Rivers¹, Jeffrey W. Sherman¹, Heather R. Rees¹,
Regina Reichardt², and Karl C. Klauer³

Personality and Social
Psychology Bulletin
2020, Vol. 46(3) 349–364
© 2019 by the Society for Personality
and Social Psychology, Inc
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/0146167219853842
journals.sagepub.com/home/pspb



Abstract

Stereotypes can influence social perception in undesirable ways. However, activated stereotypes are not always applied in judgments. The present research investigated how stereotype activation and application processes impact social judgments as a function of available resources for control over stereotypes. Specifically, we varied the time available to intervene in the stereotyping process and used multinomial modeling to independently estimate stereotype activation and application. As expected, social judgments were less stereotypic when participants had more time to intervene. In terms of mechanisms, stereotype application, and not stereotype activation, corresponded with reductions in stereotypic biases. With increasing time, stereotype application was reduced, reflecting the fact that controlling application is time-dependent. In contrast, stereotype activation increased with increasing time, apparently due to increased engagement with stereotypic material. Stereotype activation was highest when judgments were least stereotypical, and thus, reduced stereotyping may coincide with increased stereotype activation if stereotype application is simultaneously decreased.

Keywords

prejudice/stereotyping, self-regulation, implicit cognition, multinomial modeling

Received January 11, 2018; revision accepted April 28, 2019

A vast literature in psychological science demonstrates that stereotypes influence social judgment (e.g., Hamilton & Sherman, 1994). In part, this influence reflects the fact that stereotypes are readily activated, or made accessible, upon perceiving members of distinct social groups (e.g., Kunda, Davis, Adams, & Spencer, 2002; Kunda & Spencer, 2003). Indeed, stereotype activation frequently exhibits characteristics of automaticity. For example, perceiving a stereotypic image can lead to stereotype activation even when the image is incidental to current goals or is perceived outside of awareness entirely (Kawakami, Dovidio, & Dijksterhuis, 2003; Lepore & Brown, 1997). Considerable research has additionally shown that stereotypes are more likely to influence judgments when processing resources are scarce, supporting the view that stereotyping is an efficient process (Sherman, Macrae, & Bodenhausen, 2000).

Although stereotypic biases are pervasive, in some circumstances, people are able to correct for their biases (e.g., Monteith, Ashburn-Nardo, Voils, & Czopp, 2002; Monteith, Sherman, & Devine, 1998). Thus, there is a conceptual distinction between the activation of stereotypes and their application to judgment (e.g., Gilbert & Hixon, 1991). Activated stereotypes are not always applied and, sometimes, judgments are contrasted or corrected away from activated stereotypes. The distinction between stereotype activation and

application is critical to understanding when stereotypes will or will not bias judgments and to understanding how to effectively intervene to influence stereotyping.

Operating Conditions of Stereotype Activation and Application

Prominent models of stereotyping propose distinct roles for stereotype activation and application. For example, Devine's (1989) dual-process model proposed that stereotype activation is inescapable among individuals with knowledge of those stereotypes. However, the influence of activated stereotypes on judgment can be modulated depending on the perceiver's concurrent processing goals and their ability to

¹University of California, Davis, USA

²University of Regensburg, Germany

³University of Freiburg, Freiburg im Breisgau, Germany

Corresponding Authors:

Andrew M. Rivers, Department of Psychology, University of California, Davis, 1 Shields Avenue, Davis, CA 95616, USA.
Email: amrivers@ucdavis.edu

Jeffrey W. Sherman, Department of Psychology, University of California, Davis, 1 Shields Avenue, Davis, CA 95616, USA.
Email: jsherman@ucdavis.edu

correct for the influence of active stereotypes. Fazio, Jackson, Dunton, and Williams (1995) also emphasize the role of corrective processes following stereotype activation, positing that stereotypes are automatically activated upon encountering members of another group, but that corrective processes can intervene to modulate their influence. Motivations and situational conditions conducive to the operation of corrective processes are critical (Fazio, 1990; Fazio et al., 1995). When people are motivated to correct against biasing influences and conditions are sufficient for the operation of corrective processes, then judgments will deviate from activated stereotypic knowledge.

The models proposed by Devine and Fazio are reflective of a broad consensus that stereotype activation is relatively more automatic (i.e., unintentional, resistant to interference, outside of awareness, and efficient) than corrective processes that prevent the application of stereotypes to judgment (e.g., Hamilton & Sherman, 1994). An important implication is that the conditions that permit stereotype activation are assumed to be less restrictive than are the conditions that permit correction. Thus, factors that selectively interfere with stereotype correction should increase stereotyping by increasing the likelihood that stereotypes will be applied in judgments. Conversely, factors that increase the ability to intervene should reduce stereotyping via reduced application.

Resource-Dependent Stereotype Activation

Although the broad characterization of stereotype activation as automatic and stereotype application as controlled has been the consensual view of stereotyping for some time, there are hints that it may be oversimplified. For example, work by Gilbert and Hixon (1991) showed that the availability of cognitive resources influenced both stereotype activation and application. In particular, the presence of a cognitive load while encountering stereotypic content reduced the extent to which stereotypes were activated. In addition, stereotypic judgments were particularly likely if (a) participants were not cognitively busy during a stereotype activation phase (and, therefore, had stereotypes activated) and (b) were busy during a stereotype application phase (and, therefore, could not correct for stereotypic influence). Completing a simultaneous task appeared to influence both stereotype activation and stereotype application, challenging the prevailing notion that stereotype activation is categorically automatic.

“Automatic” Corrective Processes

Other research has challenged the idea that stereotype-correcting processes are necessarily resource-intensive. Glaser and Kihlstrom’s (2006) Compensatory Automaticity model asserts that habitual correction for stereotypic biases routinizes corrective processes such that preventing the application

of stereotypes need not be deliberate and resource-intensive, and that people can rapidly shift the likelihood that they apply activated stereotypes to judgments. These corrective shifts have been observed even on implicit tasks that are typically thought to preclude intentional, resource-dependent processes (Glaser & Kihlstrom, 2006; Glaser & Knowles, 2008; Sherman, 2006; Sherman et al., 2008).

Providing additional support for efficient stereotype correction, Moskowitz and colleagues (Moskowitz, Gollwitzer, Wasel, & Schall, 1999; Moskowitz & Li, 2011) have shown that stereotype activation can be preconsciously inhibited when people chronically pursue egalitarian goals or when egalitarian goals are made salient. These inhibitory mechanisms can operate under conditions thought to preclude the operation of strategic processes. Specifically, those who pursue egalitarian goals appear to inhibit stereotype activation, even on sequential priming measures using brief stimulus onset asynchronies (SOAs), which significantly limit the time available to correct for stereotyping. Altogether, these results indicate that the characterizations of stereotype activation as automatic and stereotype application (and correction) as controlled may be overly simplified.

Measuring Stereotype Activation and Application

Progress in understanding the impact and nature of stereotype activation and application has been hampered by methodological limitations. In particular, the presumption that stereotype activation is automatic whereas stereotype application or its converse, stereotype correction, is nonautomatic has formed the theoretical rationale for measuring activation and application with different experimental tasks. Because it is assumed to be an automatic process (unintentional, resource independent, lacking awareness, spontaneous), stereotype activation has generally been assessed with implicit measures (e.g., Implicit Association Test [IAT], sequential priming), which are commonly understood to measure the activation of mental associations and preclude the influence of controlled cognitive processes. In contrast, because stereotype correction is assumed to be a nonautomatic process (intentional, resource dependent, involving awareness, strategic), it has generally been assessed with explicit judgment tasks, which are presumed to reflect controlled cognitive processes while minimizing the influence of automatic processing.

An implicit assumption of this “task dissociation” approach (for a review, see Sherman, Krieglmeier, & Calanchini, 2014) is that the tasks used to measure stereotype activation and stereotype application are *process-pure*—that is, that they uniquely measure the process in question without contamination from other mental processes. In other words, the task-dissociation approach assumes that measures of activation (e.g., sequential priming task) reflect only differences in activation and not

differences in application, whereas measures of application (e.g., explicit judgment task) reflect only application. Returning to Gilbert and Hixon's (1991) work investigating the automaticity of stereotype activation and application, we find an example of this approach. Participants in Gilbert and Hixon's (1991) experiments completed a word-fragment completion task in which each fragment could be completed in either a stereotype-relevant or stereotype-irrelevant way. The number of stereotypic word completions was interpreted as an index of stereotype activation, uncontaminated by stereotype application or the prevention thereof. On a subsequent task, participants judged the degree to which their Asian or White interaction partner possessed stereotypic traits. This task was assumed to measure the degree to which participants would apply or prevent the application of stereotypes to their judgments. However, there is now a great deal of evidence that measures of stereotype activation, even implicit measures (e.g., sequential priming, IAT), also reflect the influence of stereotype application and corrective processes (Krieglmeyer & Sherman, 2012; Sherman et al., 2008). Similarly, explicit judgment tasks used to measure stereotype application are necessarily influenced by the extent of stereotype activation. Thus, it is difficult to determine the extent to which performance on either kind of task reflects stereotype activation, stereotype application, or mixtures of both processes. The task-dissociation approach precludes strong conclusions about the conditions under which these processes do or do not occur.

The Present Research

The present research is the first investigation into the operating conditions of stereotype activation and application that does not rely on the task-dissociation approach. In this line of work, we use the Stereotype Misperception Task (SMT; Krieglmeyer & Sherman, 2012), a sequential priming procedure designed specifically to disentangle the joint influences of stereotype activation and application processes. A mathematical model of SMT task performance independently assesses the extents of stereotype activation and application within the same task and under identical conditions, avoiding the inherent problematic assumptions of task dissociation. In this line of work, we investigate the roles of stereotype activation and application under conditions that vary the time available to correct for the influence of stereotypes, providing a direct test of the extent to which the processes of stereotype activation and application operate efficiently and the extent to which they are dependent on the opportunity for intentional processing. We also examine how stereotype activation and application relate to biases in judgment, and how those relationships may be affected by the opportunity for intentional processing.

In the current variant of the SMT, participants judged a series of pixelated target faces according to their level of threat. Prior to viewing these target faces, participants were

briefly exposed to photographs of Black or White male prime faces. In past research, exposure to these Black or White primes influenced how threatening the target faces were judged to be. To examine the influence of processing resources on stereotyping, stereotype activation, and stereotype application, we varied the time separating onset of prime and target images, or SOA. Previous research has shown that the influence of primes on target judgments is reduced with longer SOAs (Fazio et al., 1995). As such, we expected that, as SOA increased, the influence of racial primes on judgments of target images would be reduced. Most importantly, we apply a statistical model that allows us to derive independent estimates of stereotype activation and application to shed light on how these processes relate to stereotyping, and how they are influenced by SOA. The influence of stereotypes may be reduced as SOA increases because stereotype activation is diminished, either due to inhibition of the stereotype (Monteith et al., 1998; Moskowitz & Li, 2011) or to passive decay of activated stereotypic knowledge (Kunda et al., 2002). But increases in SOA might also allow people to better prevent the application of activated stereotypes to their judgments (e.g., Glaser & Knowles, 2008; Sherman et al., 2008). In this case, it is not the extent of stereotype activation that is critical, but the degree to which people apply active stereotypes. Application of the SMT model allows us to independently test how SOA affects each of these mechanisms and the likelihood that they contribute to reductions in stereotypic judgments.

Experiment 1

Participants

In Experiment 1, we sought to collect data from at least 80 participants. In total, 92 undergraduate students¹ at the University of California, Davis participated in the experiment for partial course credit (71.1% female, $M_{\text{age}} = 18.9$ years; 58% Asian, 21% Caucasian, 20% Latino/a). Participants completed the experiment as the second of three unrelated tasks in an hour-long experimental session. Based on previously reported effect sizes from Krieglmeyer and Sherman (2012; $d = 1.04$), Experiment 1 was powered well above .95 to detect stereotypic bias in the SMT procedure. However, we were specifically interested in changes in bias resulting from our within-subjects manipulation of SOA. After choosing our desired sample size, sensitivity analysis in G*Power 3.1 indicated that 80 participants would allow us to detect an effect size of $d_z = .317$, corresponding to a small to medium effect, with power set at $1 - \beta = .80$ (Faul, Erdfelder, Buchner, & Lang, 2009). Power analysis in subsequent experiments relies on observed effects of the SOA manipulation. Two participants were excluded from analyses for using a single key response for all trials.² Including these data in analyses does not change the direction or statistical interpretation of the reported results.

Stimuli

Prime stimuli consisted of photographs of 24 Black and 24 White males, each approximately 20 to 30 years old. Each face was cropped at the base of the neck and superimposed on a plain gray background (see Phills, Kawakami, Tabi, Nadolny, & Inzlicht, 2011). In addition to these, we included a set of “neutral” prime images that contained no racial cues. These consisted of the outline of a face superimposed on a gray background (see Krieglmeier & Sherman, 2012).

Target stimuli consisted of 48 computer-generated face morphs created by Oosterhof and Todorov (2008) to vary systematically in perceived threat. Each of the computer-generated images depicts a male face with morphed facial features that were either two standard deviations above or below the neutral point of threat. These images were distorted with a pixilation filter using photo-editing software to increase their ambiguity. Target images differ objectively along the dimension of evaluation to apply the multinomial processing tree model to the data (see “Multinomial Modeling Analyses” section below). The choice of ± 2 standard deviations was made to be consistent with the original validation of the SMT procedure reported in Krieglmeier and Sherman (2012).³

Procedure

In each session, one to four participants completed experimental procedures individually in separate computer cubicles. After providing informed consent, participants learned that the experiment was concerned with rapid impression formation. Participants then viewed SMT task instructions and completed 12 practice trials of the SMT procedure. Two test blocks of the SMT procedure followed the practice trials. After completing SMT test blocks, participants completed several exploratory measures, a brief demographic questionnaire (age, sex, ethnicity), and an experimental debriefing.⁴

SMT. The SMT developed by Krieglmeier and Sherman (2012) served as the dependent measure. The structure of the SMT is similar to commonly used sequential priming measures (e.g., Affect Misattribution Procedure; Payne, Cheng, Govorun, & Stewart, 2005). In the present experiment, participants were instructed to “respond as quickly as possible,” judging each target face morph as either “more threatening” or “less threatening” than the average target. Each target was preceded by one of the three prime types (Black face, White face, or neutral face outline). Participants were explicitly informed to not respond to the prime faces but to attend to them for later questions. Because of these explicit instructions, any influence of the prime faces on threat judgments is assumed to be unintentional.

The two SMT test blocks consisted of 72 trials each and included a self-paced break between the two blocks. Each trial began with a fixation cross for 500 ms. A prime image

was presented for 150 ms following the fixation cross. On half of the trials, a blank screen was then displayed for 175 ms, followed by onset of a target image. On the other half of the trials, target images were presented immediately after offset of prime images. This created trials with one of two within-subjects levels of SOA: 150 ms or 325 ms. Target images were always shown for 100 ms and were then replaced by a static visual mask that remained on the screen until the participant rendered a judgment. Participants indicated whether they judged the target to be more or less threatening than average by pressing either the “D” or “K” keys on a standard computer keyboard. The following trial began 500 ms after the previous judgment.

Results

Analytic plan. The full design for Experiment 1 was 3 (prime type: Black vs. White vs. neutral) \times 2 (target type: high threat vs. low threat) \times 2 (SOA: 150 vs. 325 ms), with all factors manipulated within subjects. Our analysis is comprised of two stages. In addition to quantifying the effect of SOA on mean levels of racial bias, we use the SMT multinomial processing tree model to disentangle the contribution of multiple component processes (see “Multinomial Processing Analyses” section below for a detailed description).

In the first analytic stage, we investigate whether participants exhibited stereotypic biases in their judgments about target stimuli. The extent of bias on the SMT, called the *SMT Effect*, is determined by comparing responses on trials that include Black versus White prime faces (see Krieglmeier & Sherman, 2012, for validation of this index). We describe the traditional full factorial ANOVA analyses in the supplemental appendix because they are not the focus of the present findings. For the SMT, the design consists of two prime levels (Black vs. White) and two SOA levels (150 vs. 325 ms). This analytic stage quantifies (a) the magnitude of racial bias in participants’ judgments and (b) the degree to which SOA moderates racial bias in participants’ judgments.

In the second analytic stage, we investigate how SOA influences the component processes of stereotype activation and application. Thus, in the second analytic stage, we apply the SMT multinomial processing tree model, which has been developed and validated for exactly this purpose—to investigate patterns of responding across the SMT procedure (Krieglmeier & Sherman, 2012; also see Payne, Hall, Cameron, & Bishara, 2010, for a similar approach with the Affect Misattribution Procedure). To accomplish this, the SMT model analysis includes responses to all three prime and both target types.

SMT effects. Our first research question pertained to the magnitude of stereotypic biases in threat judgments as a function of SOA. The SMT effect, an index of stereotypic bias, is calculated by subtracting the proportion of “more threatening” judgments on White prime trials from the proportion of “more

Table 1. Proportion “More Threatening” Responses as a Function of Prime (White vs. Neutral vs. Black), Target (Low vs. High Threat), and Experimental Level for Experiments 1 to 4.

	White prime		Neutral prime		Black prime	
	Low threat	High threat	Low threat	High threat	Low threat	High threat
Experiment 1						
150-ms SOA	.23 (.21)	.29 (.24)	.27 (.27)	.42 (.30)	.49 (.30)	.54 (.27)
325-ms SOA	.36 (.24)	.40 (.24)	.28 (.28)	.41 (.31)	.54 (.29)	.55 (.26)
Experiment 2						
150-ms SOA	.28 (.20)	.31 (.20)	.25 (.28)	.40 (.29)	.49 (.27)	.51 (.25)
200-ms SOA	.33 (.21)	.36 (.22)	.27 (.29)	.41 (.30)	.51 (.27)	.52 (.30)
325-ms SOA	.38 (.24)	.42 (.23)	.29 (.29)	.39 (.29)	.49 (.25)	.54 (.26)
Experiment 3						
150-ms SOA	.36 (.27)	.36 (.22)	.24 (.30)	.31 (.32)	.33 (.30)	.38 (.28)
325-ms SOA	.38 (.26)	.42 (.29)	.23 (.27)	.34 (.33)	.31 (.31)	.42 (.27)
Experiment 4						
175-ms SOA with mask	.33 (.26)	.38 (.26)	.26 (.28)	.31 (.32)	.49 (.30)	.48 (.31)
175-ms SOA no mask	.31 (.27)	.29 (.27)	.25 (.34)	.32 (.34)	.52 (.30)	.52 (.27)
350-ms SOA with mask	.44 (.27)	.46 (.25)	.29 (.35)	.38 (.36)	.52 (.29)	.50 (.30)
350-ms SOA no mask	.47 (.22)	.45 (.27)	.28 (.34)	.36 (.35)	.44 (.32)	.42 (.30)

Note. Standard deviations appear in parentheses. SOA = stimulus onset asynchrony.

Table 2. Proportion “More Threatening” Response Difference Score (Black–White Prime) and SMT Effect Size Estimates in Experiment 1 by SOA.

Experiment 1	Difference score	SMT effect d_z
150-ms SOA	.25 [.18, .33]	.737
325-ms SOA	.16 [.09, .23]	.494

Note. Values given in brackets denote 95% confidence intervals. SMT = Stereotype Misperception Task; SOA = stimulus onset asynchrony.

threatening” judgments on Black prime trials (Krieglmeyer & Sherman, 2012).⁵ We calculated the SMT effect for both SOA conditions. The data confirmed our primary prediction; the SMT effect was stronger when SOA between prime and target images was 150 ms versus 325 ms, $t(89) = 4.074, p < .001$; Hedges $g_{av} = .268$; 95% confidence interval ($CI_{\text{difference}}$) = [.047, .137] (see Tables 1 and 2). To understand the mechanism(s) underlying this difference in bias, we employed multinomial modeling in the second analytic stage.

Multinomial modeling analyses. We generated a set of equations representing the SMT process model that was developed and previously validated (see Krieglmeyer & Sherman, 2012). The SMT process model estimates four latent parameters (see Figure 1). At the initial branch of the model tree, an *activation* (SAC) parameter captures activation of racial stereotypes. That is, to what extent do primes activate stereotypic schema (e.g., the stronger association of threat with Black than White men)? When stereotypes are active, a parameter representing *application* (SAP) captures whether those stereotypes are applied to judgment or whether judgments are corrected away from them. When

threat stereotypes are not activated ($1 - SAC$), a parameter representing *detection* (D) captures the ability to accurately detect target threat level and respond accordingly. Finally, when threat stereotypes are not activated ($1 - SAC$) and target characteristics are not detected ($1 - D$), a *guessing* parameter (G) captures general tendencies to respond with high- or low-threat judgments.

To further explicate the model, consider the case in which a participant is responding to a trial in which the prime face is Black and the target image is low in threat. In this case, the Black prime activates threat-related stereotypic content with the probability of SAC. If the stereotype is activated and it is applied to the judgment, the participant will render a “more threat” judgment with the probability of $SAC \times SAP$. However, the participant may correct their judgment away from the activated stereotype, rendering a “low threat” judgment with the probability of $SAC \times (1 - SAP)$. If the Black prime does not activate the stereotype, participants may correctly detect the target image, rendering the “low threat” judgment with probability $(1 - SAC) \times D$. If the stereotype is not activated and the extent of threat in the target image is not accurately detected, then the participant may guess “high threat” with probability $(1 - SAC) \times (1 - D) \times G$ or, alternatively, may guess “low threat” with probability $(1 - SAC) \times (1 - D) \times (1 - G)$. Thus, the overall probability of a “high threat” response on this trial is $SAC \times SAP + (1 - SAC) \times (1 - D) \times G$. The probability of a “low threat” response is $SAC \times (1 - SAP) + (1 - SAC) \times D + (1 - SAC) \times (1 - D) \times (1 - G)$. According to the SMT, model estimates for SAP, D, and G are conditional probabilities, in that each parameter’s influence is dependent on the activity (or inactivity) of another parameter (e.g., G is conditional on the inactivity of




Figure 1. Multinomial processing tree of the SMT model.

Note. The top tree for trials with black and white primes and the bottom tree for neutral primes. The table on the right depicts responses as a function of prime and target. The response “more threat” is represented by a “+” sign, and the response “less threat” is represented by a “-” sign. SMT = Stereotype Misperception Task; SAC = stereotype activation; 1 - SAC = absence of stereotype activation; SAP = stereotype application; 1 - SAP = stereotype correction; D = detection of target trait; 1 - D = failure to detect target trait; G = tendency to guess more threat; 1 - G = tendency to guess less threat.

stereotype activation and detection). In contrast, the estimate of SAC represents the unconditional probability of stereotype activation.

Multinomial modeling analyses were conducted using the freely available *multiTree* computer package (Moshagen, 2010). This package implements a maximum likelihood framework to test the goodness of model fit and to estimate parameter values. Each of the four parameters is manipulated in an iterative process until the model’s expected frequencies most closely approximate the observed response frequencies. Parameter estimates vary between 0 and 1, and represent the probability of process involvement. The magnitude of discrepancy between model expectations and observed frequencies is expressed in the G^2 statistic and corresponding p value. A nonsignificant result indicates that any discrepancy between the expected data and the observed data were not detectable.

Modeling results. Frequency counts of more and less threatening responses were aggregated for each of the SMT trial types. We fit the SMT process model equations to aggregated counts from trials in which the SOA was 150 ms and 325 ms. When fit to the data, the model G^2 statistic suggested that the fit of the SMT model was acceptable, $G^2(4) = 6.461$, $p = .167$. To further quantify the magnitude of misfit, we calculated the w coefficient, which can be thought of as the

effect size of model misspecification after controlling for power (see Cressie, Pardo, & Pardo, 2003). The resulting estimate of $w = .022$ indicated that the magnitude of model misfit was small after controlling for power.⁶

Our primary questions of interest center on whether model parameters reliably respond to changes in prime-target SOA. To compare model parameters across the two SOA conditions, we first fit a baseline model in which all parameters from the two SOA conditions were permitted to freely vary. By constraining corresponding parameters across SOA levels, we created nested models to test against the baseline model. The addition of any constraint necessarily reduces fit, increasing G^2 . Large reductions in model fit result in higher ΔG^2 from baseline to nested models. Statistically significant p values extracted from ΔG^2 indicate that the constrained model should be rejected in favor of the baseline model. In other words, significant p values suggest that parameters differ between the levels being compared and should, therefore, be independently estimated (as in the baseline model). It is important to point out that the SMT multinomial model estimates the cognitive processes that are thought to underlie stereotyping, but this modeling approach cannot separately estimate processes that are specific to particular stimuli. Thus, our investigation tests how SOA influences these domain-general stereotyping mechanisms.

Table 3. SMT Multinomial Model Parameter Estimates by Experimental Level in Experiments 1 to 4.

Parameter estimate	D	G	SAC	SAP
Experiment 1				
150-ms SOA	.11 [.08, .14]	.32 [.29, .34]	.35 [.24, .45]	.87 [.75, .99]
325-ms SOA	.13 [.09, .17]	.32 [.30, .34]	.77 [.68, .86]	.60 [.58, .63]
Experiment 2				
150-ms SOA	.10 [.06, .15]	.30 [.26, .33]	.43 [.30, .55]	.73 [.65, .81]
200-ms SOA	.11 [.07, .16]	.32 [.28, .35]	.55 [.42, .68]	.65 [.60, .69]
325-ms SOA	.11 [.06, .16]	.32 [.29, .36]	.67 [.54, .80]	.58 [.55, .61]
Experiment 3				
150-ms SOA	.06 [.01, .11]	.26 [.22, .30]	.37 [.23, .52]	.49 [.43, .56]
325-ms SOA	.10 [.04, .16]	.26 [.22, .30]	.59 [.46, .73]	.51 [.47, .56]
Experiment 4				
175-ms SOA with mask	.05 [-.02, .12]	.28 [.23, .32]	.61 [.48, .74]	.60 [.56, .64]
175-ms SOA without mask	.01 [-.05, .07]	.30 [.25, .34]	.55 [.40, .70]	.71 [.64, .77]
350-ms SOA with mask	.12 [.04, .21]	.29 [.24, .34]	.94 [.82, 1.0]	.52 [.50, .55]
350-ms SOA without mask	.04 [-.03, .12]	.31 [.27, .36]	.66 [.50, .81]	.49 [.45, .52]

Note. Values given in brackets denote 95% confidence intervals. SMT = Stereotype Misperception Task; SAC = stereotype activation; SAP = stereotype application; SOA = stimulus onset asynchrony.

The D parameter could be collapsed across the two SOA levels without a statistical loss of model fit, $\Delta G^2(1) = 0.457$, $p = .499$, $w = .004$. This indicates that D did not detectably differ between the two SOA levels. Consistent with previously reported SMT data, the point estimate for D (see Table 3) was relatively low at both SOA levels, 150 ms ($D = 0.114$; 95% CI = [.083, .145]) and at 325 ms ($D = 0.131$; 95% CI = [.093, .169]). However, that CIs did not overlap with 0 indicates that participants reliably discriminated between high- and low-threat targets.

Similarly, the G parameter did not differ between SOA conditions, $\Delta G^2(1) = 0.040$, $p = .841$, $w < .001$. When no target or prime information was available to inform judgments, 95% CIs of the G parameter did not overlap with 0.5, indicating that participants tended to guess low threat at both 150-ms ($G = 0.317$; 95% CI = [.294, .341]) and 325-ms SOA ($G = 0.321$; 95% CI = [.297, .344]).⁷

Constraining the SAP parameter across the two SOA levels led to a significant loss of fit, $\Delta G^2(1) = 56.950$, $p < .001$, $w = .066$. This indicates that SAP differed significantly between SOA conditions. The tendency to apply active stereotypes was lower at 325-ms SOA ($SAP = .606$; 95% CI = [.584, .629]) than at 150-ms SOA ($SAP = .870$; 95% CI = [.749, .990]). CIs for the SAP parameter did not overlap with 0.5 at either SOA, indicating that active stereotypes tended to be applied (vs. corrected against) at both levels.

The SAC parameter, likewise, differed between the two SOA levels, $\Delta G^2(1) = 34.619$, $p < .001$, $w = .052$. Unexpectedly, stereotype activation was higher at 325-ms SOA ($SAC = .767$; 95% CI = [.676, .858]) compared with 150-ms SOA ($SAC = .345$; 95% CI = [.235, .455]). CIs for SAC did not overlap with 0, indicating that stereotypes were active at both SOAs.

Discussion

Data from Experiment 1 supported our primary prediction that the biasing impact of primes on judgments would decrease as the time separating prime and target images increased. In addition, modeling analyses showed that reduced stereotype application rather than reduced stereotype activation corresponded with reductions in bias. As time between prime and target increased, participants were less likely to apply prime-activated threat stereotypes. In addition, an unexpected finding emerged from Experiment 1; stereotype activation was higher when SOA was longer versus shorter. This means that stereotype activation was highest when judgments showed the least amount of racial bias.

Experiment 2

Rationale

The primary goal of Experiment 2 was to closely replicate the aforementioned exploratory findings from Experiment 1. To be explicit, we now expected that estimates of stereotype activation would again be higher at longer versus shorter SOA. A secondary goal was to more deeply investigate the efficiency of stereotyping processes. In Experiment 1, we sought to understand the relationships among stereotype activation, stereotype application, and stereotypic judgment at the presumed boundaries of strategic responding. Early semantic priming research suggested that SOAs below 400-ms eliminated such processes (Neely, 1977). However, more recent work suggests that a strict threshold of 400-ms SOA is likely untenable (Hutchison, 2007). In fact, Hutchison (2007) found that the use of strategic response strategies could be seen at SOAs as brief as 250-ms. In addition, controlled

faking of responses to evaluative priming tasks is possible, even with a relatively brief SOA (i.e., 280-ms) and brief response windows (i.e., 600-ms; Teige-Mocigemba & Klauer, 2013; also see Teige-Mocigemba & Klauer, 2008). In Experiment 1, we showed that stereotyping processes were diminished with a 350-ms SOA compared with a 150-ms SOA. To further examine the efficiency profile of stereotype control processes, in Experiment 2, we tested whether greater control would be observed with an SOA of 200-ms versus an SOA of 150-ms. Although we did not have strong a priori predictions for the 200-ms condition, we felt such a condition would be informative. If stereotypic bias is reduced by an additional 50-ms (from 150-ms to 200-ms SOA), this would provide evidence that control is possible even at 200-ms, well below Neely's (1977) initial boundary.

To accomplish these goals, we modified Experiment 1's procedure by adding a third 200-ms SOA level. This additional 200-ms SOA level exactly corresponds to conditions used in the SMT's initial validation (see Krieglmeyer & Sherman, 2012). Thus, each of the 144 trials was assigned to 150-ms, 200-ms, or 325-ms SOA.

Participants

Seventy-five undergraduate students at the University of California, Davis participated in Experiment 2 for partial course credit (82.2% Female, $M_{age} = 19.4$ years; 52% Asian, 26% Caucasian, 19% Latino/a, 3% Black). Five participants were excluded from analyses according to our a priori standards. Including all data in analyses does not change the direction or statistical conclusions of the reported results. Based on effect size estimates from Experiment 1 ($d_z = .430$), 45 participants were necessary to detect a similar effect at $1 - \beta = .80$ power. However, we expected that adding an additional within-subjects level of the SOA manipulation would reduce the number of data points informing estimate at each level of SOA. Given this additional uncertainty, we sought to obtain a similar number of participants to Experiment 1. Using G*Power and multiTree power analyses, we estimated that the final sample of 70 provided power greater than $1 - \beta = .95$ to detect effects of the 150-ms versus 325-ms SOA manipulation on the SMT effect, SAP, and SAC at levels similar to Experiment 1 (Faul et al., 2009; Moshagen, 2010). To consider experimental power comparing 150-ms to 200-ms, we conducted sensitivity analyses that showed a sample of 70 would provide .80 power to detect an effect size as small as $d_z = .300$.

Results

SMT effects. Repeated-measures ANOVA indicated that the SMT effect again differed across SOA, $F(1.81, 138)^8 = 4.868$, $p = .011$, $\omega_p^2 = .052$.⁹ Replicating Experiment 1, simple comparisons revealed that the SMT effect was

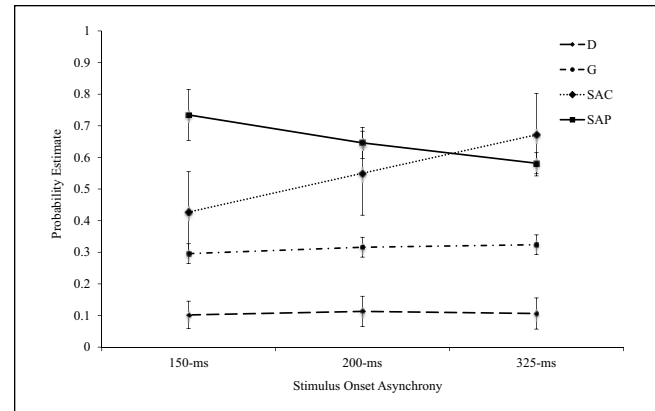


Figure 2. Probability estimates of the parameters SAC (stereotype activation), SAP (stereotype application), D (target detection), and G (guessing) by SOA level (150 vs. 200 vs. 325 ms) in Experiment 2.

Note. Error bars represent 95% confidence intervals. SAC = stereotype activation; SAP = stereotype application; SOA = stimulus onset asynchrony.

stronger at 150-ms SOA than 325-ms SOA, $F(1, 69) = 11.16$, $p = .001$, $g_{av} = .291$; 95% $CI_{difference} = [.038, .150]$. However, the SMT effect did not statistically differ between 150-ms and 200-ms SOA, $F(1, 69) = 2.52$, $p = .117$, $g_{av} = .125$; 95% $CI_{difference} = [-.011, .095]$, or between 200-ms and 325-ms SOA, $F(1, 69) = 2.18$, $p = .144$, $g_{av} = .150$; 95% $CI_{difference} = [-.018, .122]$.

Multinomial modeling analyses. When fit to the data, the model G^2 statistic suggested that the fit of the SMT model was acceptable, and that the magnitude of misfit controlling for power was small, $G^2(6) = 10.519$, $p = .104$, $w = .032$.

The effect of SOA on the SAP parameter replicated, $\Delta G^2(2) = 17.431$, $p < .001$. Active stereotypes were marginally more likely to be applied at 150-ms SOA than 200-ms SOA, $\Delta G^2(1) = 3.777$, $p = .052$, $w = .019$, and reliably more likely to be applied 325-ms SOA, $\Delta G^2(1) = 17.348$, $p < .001$, $w = .040$ (see Figure 2). Furthermore, the SAP estimate at 200-ms SOA was higher than at 325-ms, $\Delta G^2(1) = 4.949$, $p = .026$, $w = .022$. Thus, participants were less likely to apply activated stereotypes at longer versus shorter SOA.

The impact of SOA on the SAC parameter also replicated the results from Experiment 1. SAC parameters could not be constrained across the three levels, $\Delta G^2(2) = 6.823$, $p = .033$, $w = .025$. The SAC estimate at 150-ms SOA did not detectably differ from the estimate at 200-ms, $\Delta G^2(1) = 1.680$, $p = .195$, $w = .012$, but was reliably lower than the estimate at 325-ms, $\Delta G^2(1) = 6.700$, $p = .010$, $w = .025$. SAC estimates at 200-ms SOA and 325-ms SOA did not differ, $\Delta G^2(1) = 1.641$, $p = .200$, $w = .013$. Replicating the exploratory SAC result from Experiment 1 increases our confidence that stereotype activation is higher at longer versus shorter SOA.

Discussion

Experiment 2 provided additional support for the hypothesis that increasing the time between racial primes and target images mitigates the magnitude of racial bias in judgment. Furthermore, stereotype application (SAP) was lower with the 325-ms SOA than the 150-ms SOA. As in Experiment 1, differences in stereotype application rather than activation corresponded with the lower levels of racial bias observed at longer SOA. In fact, when stereotype activation was highest, judgments showed the least amount of racial bias.

Although the comparison between 150-ms and 200-ms was not significant for the SMT effect or SAC, a general trend of decreased bias and SAC was observed. Moreover, SAP was marginally lower at 200-ms SOA versus 150-ms SOA. Thus, it appears that stereotype correction is possible, even at 200-ms.¹⁰

The increase in stereotype activation at longer SOA appears to be at odds with aspects of existing theory and research. There are several reasons to expect that stereotype activation should *decrease* as SOA increases. First, priming research suggests that the activation of concepts fades over time (e.g., Kunda et al., 2002). Second, previous research suggests that stereotypes are relied upon to a greater extent when controlled processing is constrained because they are needed as efficient social judgment tools/heuristics (Macrae, Milne, & Bodenhausen, 1994). As such, when processing resources, such as time, are restricted, stereotype activation should increase to promote efficient decision making. Finally, other work suggests that stereotype activation itself can be inhibited. If this process is time-dependent, then stereotype activation would be reduced at longer SOA. However, the data did not corroborate any of these accounts. Instead, stereotype activation increased as SOA increased. Experiment 3 directly tested one plausible reason why stereotype activation might increase—increased processing of priming imagery.

Experiment 3

Rationale

Experiments 1 and 2 found and replicated the finding that stereotype activation is higher as SOA increases. Experiments 3 and 4 sought to further characterize the nature of this increase. One possibility is that participants might continually process the prime images during the longer interstimulus interval. Greater time and effort spent processing the primes would increase stereotype activation. In Experiment 3, participants completed a simple recognition memory test to determine whether increases in stereotype activation corresponded with better memory for prime images. If primes receive additional processing at longer SOA, then they should be recognized at higher rates in the memory test (Bower & Karlin, 1974).

Participants

Forty-nine undergraduate students at the University of California, Davis participated in Experiment 3 for partial course credit (84.4% Female, $M_{age} = 20.4$ years; 49% Asian, 18% Caucasian, 25% Latino/a, 6% Black). We sought to collect a sample of at least 48 participants according to our preregistered plan on the Open Science Framework¹¹ to achieve .90 power to detect an effect of $d_z = .430$. Seven participants were excluded according to our preregistered criteria, resulting in a final sample of 42. Including all data in analyses does not change the direction or statistical conclusions of the reported results.

Design

The design of Experiment 3 was similar to Experiment 1, with one modification: The SMT consisted of a single block of 96 trials in which participants encountered 16 prime images of each type (Black, White, neutral). This modification permitted us to retain 12 unseen prime images that were used as lures on the recognition test (see below). Each of the primes was randomly chosen to be presented either at the 150-ms or the 325-ms SOA, and was presented once with a high threat target and once with a low threat target.

Prime-recognition task. Directly after completing the SMT procedure, participants completed the recognition memory measure. Twenty-four previously shown “old” prime images and eight “new” lure images were presented in a random order, and participants were asked to judge whether or not each image had appeared in the previous task.

Results

SMT effects. Surprisingly, a paired-samples t test indicated that the SMT effect did not differ across SOA, $t(41) = -.671$, $p = .506$, $g_{av} = -.059$, 95% CI_{difference} = $[-.084, .042]$.¹² Moreover, we did not observe significant racial bias at either SOA level, $ps > .7$. We reflect on this further in the discussion below.

Multinomial modeling analyses. The fit of the SMT model appeared acceptable and the magnitude of misfit controlling for power was small, $G^2(4) = 2.133$, $p = .711$, $w = .023$.

Importantly, longer SOA again led to an increase in SAC, $\Delta G^2(2) = 6.828$, $p = .030$, $w = .041$. In contrast, SOA had no detectable effect on SAP, $\Delta G^2(2) = .246$, $p = .620$, $w = .008$.

Recognition memory. We computed an index of recognition accuracy for each participant by first calculating the proportion of prime images correctly identified as old (*Hits*) and subtracted from that the proportion of lures incorrectly identified as old (*False Alarms*). Accuracy was above chance performance at both 150-ms and 325-ms SOA ($ps < .001$). Critically,

and consistent with our preregistered predictions, recognition memory was higher at 325-ms than at 150-ms SOA, $t(41) = 7.717, p < .001; g_{av} = .886, 95\% CI_{\text{difference}} = [.146, .250]$.

Discussion

Experiment 3 again found evidence that increases in SOA led to increases in stereotype activation. Consistent with the hypothesis that primes received additional processing when SOA was long, recognition memory was better for primes when SOA was 325-ms versus 150-ms. Despite high statistical power for the fully within-subject design, the impact of SOA on the SMT effect and on SAP did not replicate the relationships found in Experiments 1 and 2. Neither did this study find evidence of racial bias in people's judgments. These results were unexpected given the robust stereotypic biases and effects of SOA observed in Experiments 1, 2, and 4. It is unclear whether the absence of an effect is a statistical anomaly or due to an unidentified moderating variable.¹³ To be clear, there were no changes in the paradigm that we expected to change the otherwise robust effect of SOA.

Experiment 4

Rationale

Experiment 3 demonstrated that longer relative to shorter SOA produced greater stereotype activation and led to better memory for prime images. Both of these results are consistent with the idea that participants were continuing to process the primes during the interstimulus interval. Experiment 4 sought to test whether this additional processing was perceptual or conceptual in nature. It is possible that iconic memory increases the length of time that prime images are available in visual working memory on trials with longer versus shorter SOA. Each of the previous experiments presented prime images for 150-ms each, but Sperling's (1960) partial report paradigm demonstrates that iconic representations can persist after offset of visual stimuli for up to 1,000-ms in visual working memory. An iconic memory interpretation would suggest that stereotype activation was highest at longer SOA levels because prime images were accessible in visual working memory for a longer period of time. Orthogonally manipulating the presence of a backward-visual mask while holding SOA constant is a straightforward test of this possible explanation for increasing stereotype activation. An iconic representation explanation predicts that stereotype activation will increase at higher SOA *only when primes are not masked*. In contrast, if increased processing of the conceptual meaning and associations of the primes is responsible, then the increase in activation should be observed regardless of masking.

Participants

Fifty-eight undergraduate students at the University of California, Davis participated for partial course credit (70.7%

Female, $M = 20.6$ years, 59% Asian, 19% Caucasian, 21% Latino/a, 2%^{age} Black). We sought a sample of at least 52 to set power at .80 to detect an effect of $d_z = .4$. According to our a priori criteria, two participants were excluded from analyses. Including all data in analyses does not change the direction or statistical conclusions of the reported results.

Procedure

Experiment 4's procedure was similar to that of Experiment 1, with the following modifications. First, on half of the SMT trials we backward masked prime stimuli for 25-ms. The mask was sized to the same dimensions as prime stimuli and consisted of a visual black and white static pattern. An equivalent 25-ms of blank screen appeared on trials on which prime stimuli were not masked. Thus, the shorter SOA was 175-ms (150-ms prime presentation plus either 25-ms mask or 25-ms blank screen) and the longer SOA was 350-ms (150-ms prime presentation plus either 25-ms mask or 25-ms blank screen plus 175-ms blank screen). The mask factor was manipulated orthogonally, resulting in a 3 (prime type: Black vs. White vs. neutral) \times 2 (target type: high vs. low) \times 2 (SOA: 175- vs. 350-ms) \times 2 (prime mask: masked vs. unmasked) fully within-subjects design. As the SOA manipulation selectively influenced only the SAC and SAP model parameters, we sought to increase the precision of estimates for these parameters by increasing the proportion of trials with race primes relative to trials with the neutral prime. Experiment 4 included 64 Black, 64 White, and 32 neutral prime trials for a total of 160 trials per participant.

Results

SMT effects. SMT effect estimates were compiled and entered into a 2 (SOA: 175- vs. 350-ms) by 2 (mask: masked primes vs. unmasked primes) repeated-measures ANOVA model. Replicating Experiments 1 and 2, there was a main effect of SOA on the SMT effect, $F(1, 55) = 40.113, p < .001, \omega_p^2 = .407$ (see Figure 3).¹⁴ SMT effects were stronger at the shorter 175-ms SOA than the longer 350-ms SOA, $t(55) = 6.333, p < .001, g_{av} = .468$. Although there was no main effect of the mask manipulation, $F(1, 55) = 0.550, p = .461, \omega_p^2 < .001$, an interaction between SOA and mask emerged, $F(1, 55) = 7.565, p = .008, \omega_p^2 = .103$. When primes were not masked, there was a strong effect of SOA on the SMT effect, $t(55) = 5.407, p < .001, g_{av} = .622$. When primes were masked, the effect of SOA on the SMT effect was still significant, but markedly smaller in magnitude, $t(55) = 2.462, p = .017, g_{av} = .221$. Thus, masking primes dampened the influence of SOA on the SMT effect. We then examined simple effects looking within each SOA level. In contrast to the predictions from the iconic memory account, the effect of masking primes changed SMT effects at short SOA, $t(55) = 2.659, p = .010, g_{av} = .265$, and had no detectable impact at long SOA, $t(55) = -1.545, p = .128, g_{av} = -.169$.

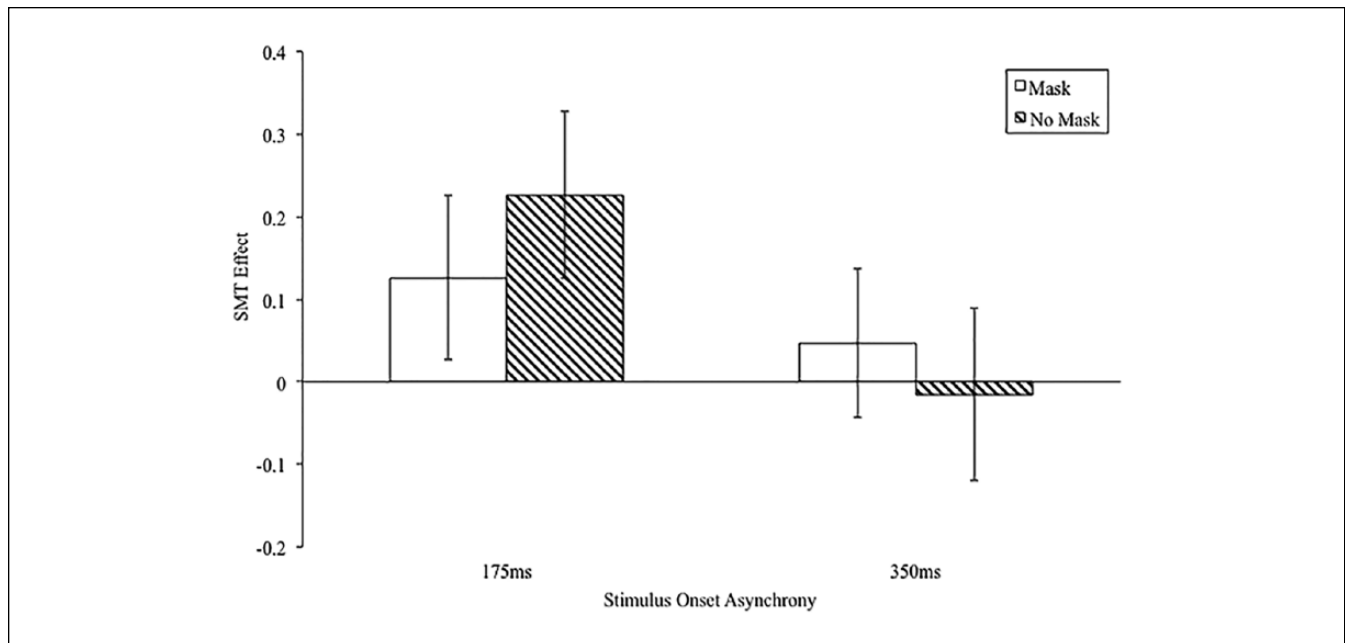


Figure 3. SMT effect by SOA (175- and 350-ms) and backward mask (present and absent) in Experiment 4. Note. Error bars represent 95% confidence intervals. SMT = Stereotype Misperception Task; SOA = stimulus onset asynchrony.

Multinomial modeling analyses. For each SMT model parameter, we generated two models—one that permitted the free interaction of the SOA and mask factors and one that allowed only main effects. The restricted main effects model was then compared against the interaction model, with significant ΔG^2 indicating the presence of an interaction between the two factors. The SMT model provided a good approximation to the data, $G^2(8) = 5.559$, $p = .696$, $w = .026$.

Once again, increasing SOA from 25-ms to 200-ms led to higher SAC, $\Delta G^2(1) = 9.316$, $p = .002$, $w = .032$ (see Table 3). As a main effect, masking prime stimuli produced a small but detectable increase in SAC versus not masking primes, $\Delta G^2(1) = 4.260$, $p = .039$, $w = .019$. This result is the opposite of what would be predicted by the iconic memory account. The critical prediction from the iconic memory interpretation is an interaction between SOA and mask, such that masking primes reduces SAC, but only when SOA is long. There was no evidence of the critical SOA \times Mask interaction, $\Delta G^2(1) = 1.274$, $p = .259$.

Replicating Experiments 1 and 2, SAP was higher when SOA was shorter versus longer, $\Delta G^2(1) = 55.251$, $p < .001$, $w = .078$. Masking prime stimuli decreased SAP versus not masking primes, $\Delta G^2(1) = 35.506$, $p < .001$, $w = .061$. There was an interaction between SOA and mask on the SAP parameter, $\Delta G^2(1) = 11.010$, $p < .001$. Simple comparisons revealed that SOA had a strong effect on SAP when primes were not masked, $\Delta G^2(1) = 21.049$, $p < .001$, $w = .049$. SOA had a significant but attenuated effect on SAP when primes were masked, $\Delta G^2(1) = 9.721$, $p = .002$, $w = .033$.

Recognition memory. Recognition accuracy was indexed using the hits minus false alarms index used in Experiment 3.¹⁵ We entered this into a repeated-measures ANOVA testing the effects of SOA and backward mask. Replicating Experiment 3, recognition for prime images was better at 200-ms SOA compared with 25-ms SOA, $F(1, 55) = 129.155$, $p < .001$, $\omega_p^2 = .692$. Recognition was directionally reduced when primes were backward masked, but this effect did not approach significance, $F(1, 55) = 1.552$, $p = .218$, $\omega_p^2 = .010$. There also was no interaction between SOA and the backward mask, $F(1, 55) = 1.450$, $p = .234$, $\omega_p^2 = .008$.

Discussion

The primary goal of Experiment 4 was to test whether increases in SAC at longer SOA were the result of additional perceptual processing of prime images in iconic memory. If true, masking prime stimuli would interact with the SOA factor. We found no evidence for this critical interaction. Instead, we found that backward masking prime stimuli, if anything, appeared to modestly increase stereotype activation. In concert with the recognition memory results from Experiment 3, this pattern of results suggests that prime images receive greater conceptual, rather than perceptual, processing when SOA is longer.

Within-Paper Meta-Analysis

We sought to quantify the effect of SOA on our three variables of interest across the four reported experiments. To do




Figure 4. Forest plot depicting effect of SOA on SMT effect by experiment.

Note. Error bars represent 95% confidence intervals. SMT = Stereotype Misperception Task; SOA = stimulus onset asynchrony.




Figure 5. Forest plot depicting effect of SOA on SAC by experiment.

Note. Error bars represent 95% confidence intervals. SOA = stimulus onset asynchrony; SAC = stereotype activation.

this, we conducted mixed-effects meta-analytic tests using the “metafor” package in the open source R platform (R Development Core Team, 2010; Viechtbauer, 2010). We aggregated data from short (150- to 175-ms) and long (325- to 350-ms) SOA conditions and converted each effect size statistic to the r correlation coefficient (Lakens, 2013).

The meta-analytic estimate for the effect of SOA on the SMT effect was significant, $Z = 2.639$, $p = .008$, $r = .387$; 95% CI = [.105, .612]. In addition, there was detectable heterogeneity, $Q(3) = 20.294$, $p < .001$, indicating that we

should reject the null hypothesis that the present experiments were examining the same effect of SOA on the SMT effect (Higgins, Thompson, Deeks, & Altman, 2003). In other words, detecting heterogeneity implies that the effect of SOA on the SMT effect was moderated by a third variable (Figure 4).¹⁶

The effect of SOA on the SAC model parameter was significant, $Z = 3.128$, $p = .002$, $r = .190$; 95% CI = [.072, .304]. There was no evidence for heterogeneity, $Q(3) = .225$, $p = .974$ (Figure 5).




Figure 6. Forest plot depicting effect of SOA on SAP by experiment.

Note. Error bars represent 95% confidence intervals. SOA = stimulus onset asynchrony; SAP = stereotype application.

The effect of SOA on the SAP model parameter was significant, $Z = 3.437$, $p < .001$, $r = .208$; 95% CI = [.091, .320]. There was no evidence for heterogeneity, $Q(3) = 4.223$, $p = .238$ (Figure 6).

Taken together, these within-paper meta-analyses indicate that increasing SOA had a reliable impact on reducing both the SMT effect and stereotype application (SAP), as well as on increasing stereotype activation (SAC).

General Discussion

The present research sought to understand the roles of basic mechanisms thought to underlie stereotypic biases in social judgment—stereotype activation and stereotype application. We manipulated the amount of time separating stereotypic primes and social targets, finding that increases in time reduced the magnitude of stereotypic bias. Two mechanistic accounts could have explained this reduction in bias. First, reductions in stereotyping could have corresponded with decays or suppression of stereotype activation. Second, reductions in stereotyping could have corresponded instead with reductions in stereotype application. We found strong evidence for the second possibility. When the two basic stereotyping mechanisms were independently estimated with the SMT processing tree model, shifts in stereotype application corresponded with stereotypic biases in peoples' social judgments. In contrast, stereotype activation appeared to play a subordinate role to stereotype application. In fact, stereotype activation consistently increased with time, even as

stereotypic biases in peoples' judgments decreased. This means that stereotypes were most active when stereotype-congruent biases in judgment were weakest. The present work demonstrates that stereotype activation might not necessarily always result in biased judgment, and that preventing the application of activated stereotypes can be an effective strategy to reduce bias.

Stereotyping: Automaticity, Control, and Operating Conditions

The present results shed light on the operating conditions of the basic mechanisms of stereotype activation and application. We tested our hypotheses using an indirect measure of stereotyping that reflects unintentional racial biases and manipulated SOA at levels assumed to preclude deliberative processing. Under these conditions, we found evidence for the modulation of both stereotype activation and application, indicating that these processes are more dynamic over even short periods of time than previously thought (but see Cunningham, Zelazo, Packer, & Van Bavel, 2007). In terms of stereotype application, this finding is consistent with both the Compensatory Automaticity model of stereotyping (Glaser & Kihlstrom, 2006), which proposes that stereotype-corrective processes themselves can proceed rapidly (i.e., at SOAs assumed to preclude controlled processes; Neely, 1977), and the suggestion from the Quadruple process model of implicit social cognition (Sherman et al., 2008) that associative biases can be overcome relatively

quickly and efficiently (Calanchini & Sherman, 2013). At the same time, we found that even short increases of time between racial primes and target images affected the likelihood of applying stereotypic information in judgments. Participants were most likely to apply active stereotypes when brief periods of time separated prime and target images. Experiment 2 found that stereotype application was diminished at each stepwise increase in time separating racial primes and target images. Thus, stereotype application and control over it, though relatively quick, is not an entirely automatic process.

Interestingly and unexpectedly, we also found that stereotype activation was greater when there was more time between primes and targets. We initially hypothesized that additional time would reduce rather than increase stereotype activation due to enhanced suppression or passive decay. However, there was no evidence in our experiments supporting the hypothesis that stereotype activation decayed or was suppressed as time increased. It is possible that our activation results reflect the fact that stereotype activation is not entirely efficient, and requires time to unfold (Gilbert & Hixon, 1991). We think that it is more likely that these results reflect the fact that the extent of even highly efficient processes can be increased by additional processing. Indeed, we found evidence from a measure of recognition accuracy that primes were attended to and processed more thoroughly at longer versus shorter SOAs. Evidence from the backward masking manipulation suggested that this additional processing is conceptual, rather than perceptual, in nature.

An important caveat is that our conclusions about the operating conditions of stereotype activation and application were based on a specific manipulation of processing resources: time. Other manipulations, such as a cognitive load (Gilbert & Hixon, 1991), may yield different conclusions. An alternative explanation regarding stereotype activation draws on the “goal looms larger” effect (e.g., Goschke & Kuhl, 1993). In contrast to semantic priming where we might expect rapid decay in concept activation, goal priming has been shown to lead to increased accessibility over time until the goal is fulfilled. Based on research on intergroup emotions (e.g., Cottrell & Neuberg, 2005; Mackie & Smith, 2002), it seems likely that our prime pictures may have activated an emotional-motivational state of threat, fear, and corresponding goals for protection. As a consequence, the accessibility of threat representations may have increased over time because they were relevant to goals for protection. Further research is needed to further delineate the conditions under which our observations hold.

Self-Regulating Implicit Biases: Cognitive Mechanisms

Results from this study underscore the importance of researching the basic mechanisms underlying the ability to mitigate stereotypic biases. A comprehensive understanding

of these basic mechanisms will provide the basis for developing effective interventions. For example, consider accounts of stereotyping that suggest that stereotype activation is the primary driver of biased judgments (e.g., ‘cognitive monster,’ Bargh, 1999). In these accounts, when stereotypic knowledge is accessible, judgments are influenced in an assimilative fashion, inevitably producing downstream bias, particularly under conditions thought to interfere with control (e.g., lack of time or cognitive resources; Devine, 1989; Fazio et al., 1995). Indeed, consistent with this proposal, we found that people generally tended to apply active stereotypes across each of the experiments.

However, challenging the activation-dominant perspective, the extent of stereotype activation could not account for the magnitude of peoples’ biases in the current research. The magnitude of racial bias in judgments and stereotype activation were dissociated. This is difficult to reconcile with theorizing that posits a direct and inevitable pathway from stereotype activation to stereotypic bias (e.g., Bargh, 1999). Even when stereotype activation was strong, participants were able to prevent the application of stereotypes under suboptimal conditions that are presumed to interfere with strategic processes. This does not necessarily mean that interventions aimed at reducing stereotype activation via suppression or other means, or by changing stereotypic knowledge outright will be unsuccessful. There is no question that such reductions in activation can reduce bias. At the very least, reducing activation bypasses the need to inhibit the application of stereotypes. Nevertheless, the present results suggest that the relationship between stereotype activation and stereotyping is more nuanced than is sometimes described. Even under conditions that make control difficult, the present results showed that stereotyping was more dependent on the extent of stereotype application than activation. Practically, this suggests that interventions to reduce bias should emphasize the role of practice in correcting for the influence of active stereotypes (e.g., Calanchini, Gonsalkorale, Sherman, & Klauer, 2013; Kawakami, Dovidio, Moll, Hermsen, & Russin, 2000).

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: Preparation of this article was partially supported by an Anneliese Maier Research Award from the Alexander von Humboldt Foundation to J.W.S. and K.C.K.

Notes

1. Data from an additional 12 participants were collected because data collection proceeded more quickly than expected. Analyses were only performed on the full dataset and at no prior time. Sensitivity power analysis indicated that the final

- sample of 90 participants allowed us to detect an effect size of $d_z = .299$ at $1 - \beta = .80$.
- We set two a priori criteria for exclusion of data for each experiment. First, participants who used a single key for every trial were excluded (e.g., Krieglmeyer & Sherman, 2012; Payne, Cheng, Govorun, & Stewart, 2005). Second, we excluded participants whose proportion of “more threatening” responses fell 2.5 standard deviations outside the sample distribution (see Krieglmeyer & Sherman, 2012).
 - Readers can view and download all target images used in this study at: osf.io/pqbhf/
 - See Supplemental Appendix Table A.2 for all demographic and exploratory measures by experiment.
 - ANOVA analyses are detailed in the supplemental appendix. The pattern of ANOVA results (i.e., significant racial bias and decreased bias at longer stimulus onset asynchrony [SOA]) is not changed when target type is included as a factor and neutral is added as a third level within the prime type factor. As stated earlier, the target factor and neutral prime type are included in the experimental design because they are necessary for the Stereotype Misperception Task (SMT) multinomial model (also see Krieglmeyer & Sherman, 2012).
 - In addition, we compared the performance of the SMT model with other frequently used process dissociation models (see Supplemental Appendix Table A.1). The SMT model unambiguously outperformed competitor models on *Akaike information criterion (AIC)*/*Bayesian information criterion (BIC)*/*minimum description length (MDL)* criteria in all experiments, with the exception that the Affect Misattribution Procedure (AMP) model maximized information on *AIC* criterion for Experiment 4 only.
 - The detection and guessing parameters were not meaningfully impacted by experimental manipulations in any of the present experiments. For simplicity, these parameters are not discussed further.
 - Huynh–Feldt correction applied for violation of sphericity.
 - As in Experiment 1, including the target type factor and the neutral level of the prime type factor does not change the pattern of results (i.e., significant racial bias and moderation by SOA level).
 - A sample of 101 participants would provide 80% power to detect a similar impact of SOA (150 ms vs. 200 ms) on SAP (where $w = .019$). A sample of 173 would be required to provide 80% power to detect a similar impact of SOA (150 ms vs. 200 ms) on the SMT effect (where $d_z = .190$).
 - Preregistration available at <https://osf.io/2xws8/>
 - Including the target type factor and neutral level within the prime type factor does not change this pattern of results.
 - One highly speculative possibility, which we do not have adequate data to test, is that the heavily publicized shooting of concertgoers in Las Vegas occurred the week prior to data collection. Having such a salient exemplar of a White male committing a violent act could increase the accessibility of associations between White males and threat, eliminating the otherwise robust stereotypic biases we observed in each of our other experiments. *A priori*, we expected Experiment 3 to produce the same effects of SOA as observed in Experiments 1, 2, and 4.
 - As in each of the previous Experiments, including the target type factor and the neutral level of the prime type factor does not change the pattern of results (i.e., significant racial bias and

moderation by SOA level). There is also a significant three-way interaction between prime, SOA, and mask that corresponds to the interaction between SOA and mask reported here.

- For full transparency, we wish to make clear that the recognition memory measure was not preregistered as in Experiment 3. Nevertheless, our predictions were the same; longer SOA should be expected to correspond with higher SAC (stereotype activation) as well as higher recognition accuracy.
- We conducted follow-up meta-analytic tests to determine if any variables statistically accounted for heterogeneity of SOA on the SMT effect. No clear conclusions emerged from these tests that could account for the heterogeneity observed.

Supplemental Material

Supplemental material is available online with this article.

References

- Bargh, J. A. (1999). The cognitive monster. In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology* (pp. 361-382). New York, NY: Guilford Press.
- Bower, G. H., & Karlin, M. B. (1974). Depth of processing pictures of faces and recognition memory. *Journal of Experimental Psychology*, *103*, 751-757.
- Calanchini, J., Gonsalkorale, K., Sherman, J. W., & Klauer, K. C. (2013). Counter-prejudicial training reduces activation of biased associations and enhances response monitoring. *European Journal of Social Psychology*, *43*, 321-325.
- Calanchini, J., & Sherman, J. W. (2013). Implicit attitudes reflect associative, non-associative, and non-attitudinal processes. *Social and Personality Psychology Compass*, *7*, 654-667.
- Cottrell, C. A., & Neuberg, S. L. (2005). Different emotional reactions to different groups: A sociofunctional threat-based approach to “prejudice.” *Journal of Personality and Social Psychology*, *88*, 770-789.
- Cressie, N., Pardo, L., & Pardo, M. C. (2003). Size and power considerations for testing loglinear models using ϕ -divergence test statistics. *Statistica Sinica*, *13*, 555-570.
- Cunningham, W. A., Zelazo, P. D., Packer, D. J., & Van Bavel, J. J. (2007). The iterative reprocessing model: A multilevel framework for attitudes and evaluation. *Social Cognition*, *25*, 736-760.
- Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology*, *56*, 5-18.
- Faul, F., Erdfelder, E., Buchner, A., & Lang, A. (2009). Statistical power analysis using G*Power 3.1: Tests for correlation and regression analyses. *Behavior Research Methods*, *41*, 1149-1160.
- Fazio, R. H. (1990). Multiple processes by which attitudes guide behavior: The MODE model as an integrative framework. *Advances in Experimental Social Psychology*, *23*, 75-109.
- Fazio, R. H., Jackson, J. R., Dunton, B. C., & Williams, C. J. (1995). Variability in automatic activation as an unobtrusive measure of racial attitudes: A bona fide pipeline? *Journal of Personality and Social Psychology*, *69*, 1013-1027.
- Gilbert, D. T., & Hixon, J. G. (1991). The trouble of thinking: Activation and application of stereotypic beliefs. *Journal of Personality and Social Psychology*, *60*, 509-517.

- Glaser, J., & Kihlstrom, J. F. (2006). Compensatory automaticity: Unconscious volition is not an oxymoron. In R. R. Hassin, J. S. Uleman, & J. A. Bargh (Eds.), *The new unconscious* (pp. 171-195). New York, NY: Oxford University Press.
- Glaser, J., & Knowles, E. D. (2008). Implicit motivation to control prejudice. *Journal of Experimental Social Psychology, 44*, 164-172.
- Goschke, T., & Kuhl, J. (1993). Representation of intentions: Persisting activation in memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19*, 1211-1226.
- Hamilton, D. L., & Sherman, J. W. (1994). Stereotypes. In R. S. Wyer & T. K. Srull (Eds.), *Handbook of social cognition* (2nd ed., Vol. 2, pp. 1-68). Hillsdale, NJ: Lawrence Erlbaum.
- Higgins, J. P. T., Thompson, S. G., Deeks, J. J., & Altman, D. G. (2003). Measuring inconsistency in meta-analyses. *British Medical Journal, 327*, 557-560.
- Hutchison, K. A. (2007). Attentional control and the relatedness proportion effect in semantic priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 33*, 645-662.
- Kawakami, K., Dovidio, J. F., & Dijksterhuis, A. (2003). Effect of social category priming on personal attitudes. *Psychological Science, 14*, 315-319.
- Kawakami, K., Dovidio, J. F., Moll, J., Hermsen, S., & Russin, A. (2000). Just say no (to stereotyping): Effects of training in the negation of stereotypic associations on stereotype activation. *Journal of Personality and Social Psychology, 78*, 871-888.
- Krieglmeyer, R., & Sherman, J. W. (2012). Disentangling stereotype activation and stereotype application in the Stereotype Misperception Task. *Journal of Personality and Social Psychology, 103*(2), 205-224.
- Kunda, Z., Davis, P. G., Adams, B. D., & Spencer, S. J. (2002). The dynamic time course of stereotype activation: Activation, dissipation, and resurrection. *Journal of Personality and Social Psychology, 82*, 283-299.
- Kunda, Z., & Spencer, S. J. (2003). When do stereotypes come to mind and when do they color judgment? *Psychological Bulletin, 129*, 522-544.
- Lakens, D. (2013). Calculating and reporting effect sizes to facilitate cumulative science: A practical primer for t-tests and ANOVAs. *Frontiers in Psychology, 4*, Article 863.
- Lepore, L., & Brown, R. (1997). Category and stereotype activation: Is prejudice inevitable? *Journal of Personality and Social Psychology, 72*, 275-287.
- Mackie, D. M., & Smith, E. R. (2002). *Beyond prejudice: Moving from positive and negative evaluations to differentiated reactions to social groups*. New York, NY: Psychology Press.
- Macrae, C. N., Milne, A. B., & Bodenhausen, G. (1994). Stereotypes as energy-saving devices: A peek inside the cognitive toolbox. *Journal of Personality and Social Psychology, 66*, 37-47.
- Monteith, M. J., Ashburn-Nardo, L., Voils, C. I., & Czopp, A. M. (2002). Putting the brakes on prejudice: On the development and operation of cues for control. *Journal of Personality and Social Psychology, 83*, 1029-1050.
- Monteith, M. J., Sherman, J. W., & Devine, P. G. (1998). Suppression as a stereotype control strategy. *Personality and Social Psychology Review, 2*, 63-82.
- Moshagen, M. (2010). multiTree: A computer program for the analysis of multinomial processing tree models. *Behavior Research Methods, 42*, 42-54.
- Moskowitz, G. B., Gollwitzer, P. M., Wasel, W., & Schall, B. (1999). Preconscious control of stereotype activation through chronic egalitarian goals. *Journal of Personality and Social Psychology, 77*, 167-184.
- Moskowitz, G. B., & Li, P. (2011). Egalitarian goals trigger stereotype inhibition: A proactive form of stereotype control. *Journal of Experimental Social Psychology, 47*, 103-116.
- Neely, J. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology: General, 106*, 226-254.
- Oosterhof, N. N., & Todorov, A. (2008). The functional basis of face evaluation. *Proceedings of the National Academy of Sciences, 105*, 11087-11092.
- Payne, B. K., Cheng, C. M., Govorun, O., & Stewart, B. D. (2005). An inkblot for attitudes: Affect misattribution as implicit measurement. *Journal of Personality and Social Psychology, 89*, 277-293.
- Payne, B. K., Hall, D. L., Cameron, C. D., & Bishara, A. J. (2010). A process model of affect misattribution. *Personality and Social Psychology Bulletin, 36*, 1397-1408.
- Phills, C. E., Kawakami, K., Tabi, E., Nadolny, D., & Inzlicht, M. (2011). Mind the gap: Increasing associations between the self and Blacks with approach behaviors. *Journal of Personality and Social Psychology, 100*, 197-210.
- R Development Core Team. (2010). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Sherman, J. W. (2006). On building a better process model: It's not only how many, but which ones and by which means? *Psychological Inquiry, 17*, 173-184.
- Sherman, J. W., Gawronski, B., Gonsalkorale, K., Hugenberg, K., Allen, T. J., & Groom, C. J. (2008). The self-regulation of automatic associations and behavioral impulses. *Psychological Review, 115*, 314-335.
- Sherman, J. W., Krieglmeyer, R., & Calanchini, J. (2014). Process models require process measures. In J. W. Sherman, B. Gawronski, & Y. Trope (Eds.), *Dual-process theories of the social mind* (pp. 121-138). New York, NY: Guilford Press.
- Sherman, J. W., Macrae, C. N., & Bodenhausen, G. V. (2000). Attention and stereotyping: Cognitive constraints on the construction of meaningful social impressions. *European Review of Social Psychology, 11*, 145-175.
- Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs: General and Applied, 74*(11), 1-29.
- Teige-Mocigemba, S., & Klauer, K. C. (2008). "Automatic" evaluation? Strategic effects on affective priming. *Journal of Experimental Social Psychology, 44*, 1414-1417.
- Teige-Mocigemba, S., & Klauer, K. C. (2013). On the controllability of evaluative-priming effects: Some limits that are none. *Cognition and Emotion, 27*, 632-657.
- Viechtbauer, W. (2010). Conducting meta-analyses in R with the metafor package. *Journal of Statistical Software, 36*(3), 1-48.