Social Cognition, Vol. 34, No. 2, 2016, pp. 119–132

CONTEXT DEPENDENCY AT RECALL: DECOUPLING CONTEXT AND TARGETS AT ENCODING

Debbie S. Ma
California State University, Northridge

Joshua Correll
University of Colorado at Boulder

Bernd Wittenbrink
University of Chicago

Research suggests that context influences judgments and implicit evaluations (e.g., attitudes toward a person depend upon the context in which we encounter that person). Importantly, previous research presents contextual cues at encoding, such that participants learn about a target person in context. Here, we decouple the process of learning about targets from manipulations of context. Participants in our studies learned about the intelligence and athleticism of a novel male target in a neutral context. Subsequently, we implicitly measured participants' attitudes toward the target in an intelligence-related or neutral context (Study 1) and in intelligence-related and athletic-related contexts (Study 2). In both studies, contextual cues (which were absent during encoding) influenced attitudes at retrieval.

Keywords: context effect; impression formation; implicit attitudes

Researchers have established that attitudes can vary considerably as a function of the context in which individuals encounter attitude objects—a phenomenon referred to as *context dependency*. Within the attitudes literature, context has been

Support for this work was provided by National Science Foundation Continuing Grant 0642580 to the second and third authors and the University of Chicago Booth School of Business to the third author. We thank Bernadette Park and the Stereotyping and Prejudice Research Laboratory for insightful comments on this manuscript and Joseph Blass, Miho Goto, and Ronnie Rowe for assistance collecting data.

Correspondence should be addressed to Debbie Ma, California State University, Northridge, Department of Psychology, 18111 Nordhoff Street, Northridge, CA 91330-8255; E-mail: Debbie.Ma@csun.edu.

broadly defined to include information presented prior to attitude assessment (Blair, Ma, & Lenton, 2001; Mitchell, Nosek, & Banaji, 2003; Schuman, Presser, & Ludwig, 1981; Schwarz, Strack, & Mai, 1991), the situation in which an individual renders an evaluative judgment (Cesario, Plaks, Hagiwara, Navarrete, & Higgins, 2010; Lowery, Hardin, & Sinclair, 2001; Rutchick, 2010), and the social setting in which an attitude object is encountered (Rydell & Gawronski, 2009; Wittenbrink, Judd, & Park, 2001). Context dependency has been a concern for as long as psychologists have been measuring attitudes (Thurstone, 1928), in part, as an ongoing discussion about tapping "true" attitudes. If attitudes are flexible enough to change from one situation to another, then how can we determine which attitude is "true"? The advent of implicit measures reignited the debate on whether "true" attitudes could be measured when early theorizing suggested that "automatic processes are effortless and are initiated spontaneously and inescapably" (Devine, 1989, p. 6, see also Bargh, 1997). Over a decade of research now suggests that implicit attitudes are more malleable than initially thought and are highly susceptible to context (Blair, 2002). Clearly context matters, but it is not clear how context exerts influence.

An early demonstration that implicit attitudes depend on context was provided by Wittenbrink and colleagues (2001). In one study, participants were presented with images of Black and White targets in church and ghetto backgrounds and attitudes toward these targets were measured implicitly. Although the target individuals remained constant, Black targets in a church were judged positively, whereas Black targets in a ghetto were judged more negatively. In a related set of studies conducted by Barden, Maddux, Petty, and Brewer (2004), participants showed more positive attitudes toward Black males who were presented in contexts that evoked positive (e.g., church goers, attorneys) versus negative Black male subtypes (e.g., factory workers, prisoners).

Recently, Gawronski and Rydell (Gawronski, Rydell, Vervliet, & De Houwer, 2010; Rydell & Gawronski, 2009) explored the conditions under which implicit evaluations will become sensitive to context. Their studies showed that if, during attitude formation, the perceiver notices that evaluative information systematically covaries with a particular context cue, implicit evaluations will become dependent on that particular cue. Participants in their studies learn evaluative information about an unfamiliar target person, "Bob." The learning trials paired evaluative information about Bob (descriptions of positive, negative, and neutral behaviors) with a context cue (a blue or yellow screen background) and manipulated whether evaluation and context covaried and whether the covariation was salient. If, for example, Bob's behaviors were mostly negative in a blue screen and positive in a yellow screen, and participants' attention was drawn to the covariation, then subsequent implicit attitudes of Bob depended on the context cue. In subsequent research Gawronski, Ye, Rydell, and De Houwer (2014) demonstrated that although initial attitudinal information is represented in a context-free manner, attitude-incongruent information tends to incorporate contextual cues from the environment, because counter-attitudinal information violated expectations thereby increasing attention to context. Their proposal offers a clear mechanism whereby implicit attitudes can become sensitive to context. It also nicely ties the phenomenon to a more general set of effects in associative learning that can even be observed in animals (Bouton, 1994).

The current studies test one of the central predictions derived from the encoding framework: namely that implicit attitudes will vary with a context cue only if that cue was present during attitude formation. That is, in all of the studies by Gawronski and colleagues, participants learn about Bob in the presence of a particular background (Bob-in-blue or Bob-in-yellow). Yet, what about a scenario where a perceiver was exposed to diverse information (e.g., from a friend or from a biography) suggesting that a previously unknown individual named Mike was very smart but athletically incompetent. If Mike were subsequently encountered in a laboratory, would he evoke a positive evaluation, despite the fact that the situational cue was never present during the initial attitude formation? If he were encountered on a basketball court, would the evaluation be more negative?

To investigate whether context cues can produce variation in implicit attitudes, without having been present during encoding, we conducted two studies in which participants learned about the athletic and intellectual abilities of a novel target, "Mike," in a relatively neutral context. Subsequently, we tested for context effects at recall by measuring participants' implicit attitudes toward Mike in either a neutral or classroom context (Study 1) or in an athletic or classroom context (Study 2). In this paradigm, no particularly meaningful context is provided at encoding. If context effects emerge exclusively because participants encode context as they learn about an attitude object, then context should not influence judgments of Mike in our paradigm. Context should affect evaluations of Mike only if contextual cues provided at recall can activate distinct subsets of information about the target. If contexts influence evaluations in this manner, then we expect that thoughts and feelings made accessible by contextual cues may guide evaluations of Mike (e.g., Judd & Lusk, 1984; Schwarz & Bohner, 2001; Tesser, 1978; Wilson & Hodges, 1992).

STUDY 1

METHOD

Participants and Design. One hundred thirty undergraduates participated for course credit or monetary compensation. The average age was 20.70 (SD = 4.69). Due to a computer error, the data from one individual had to be excluded from analyses, leaving 129 participants. The study employed a 3 (learning condition: *athletic-intelligent*, *athletic-unintelligent*, or *unathletic-intelligent*) × 2 (context type: neutral or intelligent) between-subjects design. Positivity toward that attitude object (Mike) was the dependent variable.

Learning Task. The learning task was designed to control the valence and content of attitudes toward Mike. Participants learned about Mike using a learning paradigm developed by Kerpelman and Himmelfarb (1971; cf., Gawronski et al., 2010; McConnell, Rydell, Strain, & Mackie, 2008; Rydell & Gawronski, 2009). During learning, participants viewed a picture of Mike against a plain, white background. The photo representing Mike was randomly selected for each participant from a set of 10 images of young White men, matched in attractiveness. The image was paired with a sentence about Mike's athleticism or intelligence (e.g., "Mike is in peak physical health" or "Mike has a high IQ"[1]). Participants completed 100 trials, half concerning athleticism, half concerning intelligence. Participants were instructed to guess whether statements were *characteristic* or *uncharacteristic* of Mike with a key press. The statement remained on the screen until the participant responded after which it was immediately replaced with the next statement. Half of the athletic and intelligence items were negatively worded (e.g., "Mike does not work out" and "Mike has a low GPA") to maintain participants' engagement with the task. Participants were not given initial guidance, but feedback (i.e., "Correct" or "Incorrect") was provided after each judgment, which helped to shape one of three information sets. One group learned evaluatively positive information about Mike's athletic and intellectual abilities (*athletic-intelligent*). A second group learned Mike was athletic but unintelligent (*athletic-unintelligent*) and a third group learned Mike was unathletic but intelligent (*unathletic-intelligent*). Participants in the latter two conditions were provided with information in which the valence (positive or negative) depended on the domain in question (athleticism or intelligence). The way in which participants learned about Mike was critical to the design for three reasons. First, the context during encoding was relatively neutral (a white background) and did not relate to athleticism or intelligence. Second, context did not vary as a function of valence at learning. Unlike the studies conducted by Rydell and Gawronski (2009), participants in two of our experimental conditions encountered *both* positive and negative information in the same learning environment. Third, in these two mixed-valence conditions, the nature of the information (positive or negative) covaries with domain (e.g., Mike is either athletic-but-not-intelligent or intelligent-but-not-athletic). Conceivably, the positive and negative information in these two mixed-valence conditions could cancel out and lead to relatively neutral (or ambivalent) attitudes. However, the relationship between domain and valence means that the two domains have very different evaluative implications. If an intelligence-related contextual cue (e.g., a classroom setting) selectively activates information about Mike's intelligence, participants in the *athletic-unintelligent* condition should evaluate him more negatively than participants in the *unathletic-intelligent* condition.

Evaluative Priming Paradigm. To test whether context effects emerged at recall, we assessed implicit evaluations by presenting images of Mike in a classroom for half of the participants and in a neutral context for the other half of participants. In the neutral condition, participants were primed with images of Mike and 9 other men who served as foils (because the image of Mike was randomly selected from 10 faces, the 9 unselected individuals for a given participant served as foils) against a white background. This background was the same one used during the learning task. In the classroom condition, primes were situated in classrooms. Six different

1. Two hundred seventy-two activities were pretested by three independent, convenience samples for valence (n = 23), athleticism (n = 22), and intelligence (n = 23). Based on these data, the 25 most and least athletic and intelligent sentences were selected for use in Studies 1 and 2. Selected items were submitted to a regression to ensure that athletic and intelligence sentences were significantly more positive than unintelligent and unathletic items, $F(1, 96) = 359.95$, $p < .001$. Notably, there was no evidence of either a valence difference between athletic or intelligent items on average, $F(1, 96) = .77$, $p = .38$, or of an interaction between valence and domain (athletic vs. intelligent), $F(1, 96) = 1.40$, $p = .24$.

images of classrooms were used and each of the 10 targets (Mike and the 9 foils) was superimposed onto each of the classroom backgrounds, creating a stimulus set of 60 unique person-classroom pairings from which we drew. Instructions for the task were presented on the computer and followed those typically administered in evaluative priming (Fazio, Sanbonmatsu, Powell, & Kardes, 1986; see also Wittenbrink & Schwarz, 2007). Each trial began with a fixation cross (1,000 ms), followed by a prime (350 ms), a blank screen (150 ms), and a valenced target word (i.e., *love*, *sunshine*, *flower*, *music*, *wonderful*, *gift*, *cancer*, *bomb*, *death*, *rotten*, *threat*, and *hell*). Participants categorized words as "good" or "bad" using the keyboard and accuracy feedback was delivered. A 4,000 ms inter-trial interval separated trials. Participants completed 12 practice and 66 test trials, 12 of which featured Mike as the prime (6 times each with positive and negative target words) and 54 of which were foils (each of the 9 foils appeared 3 times each with a positive and negative target word). Key assignment (i.e., which key signified good and bad) was counterbalanced across participants.

Manipulation Check. Participants rated how athletic, intelligent, unathletic, and unintelligent they thought Mike was on a 7-point scale (1 = not at all, 7 = extremely).

RESULTS AND DISCUSSION

Manipulation Check. We calculated explicit athleticism and intelligence scores by subtracting each participant's rating of Mike as unathletic and unintelligent from ratings of his athleticism and intelligence. A one-way ANOVA comparing the athleticism scores of participants by group was significant, $F(2, 127) = 385.11$, $p < .001$, $\eta_p^2 = .89$. Those who learned Mike was athletic reported higher levels of athleticism than those who learned Mike was unathletic, $F(1, 126) = 770.21$, $p < .001$, $\eta_p^2 = .86$. No difference in athleticism was found between the *athletic-intelligent* and *athletic-unintelligent* groups, $F(1, 126) = 0.17$, $p = .68$, $\eta_p^2 = .00$. There were also significant group differences in terms of explicit ratings of Mike's intelligence, $F(2, 127) = 221.40$, $p < .001$, $\eta_p^2 = .78$. Those who learned Mike was intelligent rated him as more intelligent than those who learned he was unintelligent, $F(1, 126) = 442.77$, $p < .001$, $\eta_p^2 = .78$. No difference was found between the *athletic-intelligent* and *unathletic-intelligent* groups, $F(1, 126) = .34$, $p = .56$, $\eta_p^2 = .00$. Results suggest that participants' perceptions were manipulated as intended (Table 1).

Implicit Evaluations. To analyze the evaluative priming data, we excluded trials on which participants made incorrect valence categorizations (2.12%) and trials on which response latencies were shorter than 300 ms (1.59%) or longer than 3,000[2] ms (0.96%; see Wittenbrink & Schwarz, 2007). Latencies were log-transformed and four estimates were calculated. We computed the average latency required to classify a positive target word as "good" and a negative target word as "bad" on trials primed by Mike (Mike-pos and Mike-neg) and on trials primed by foils (foil-pos and foil-neg). From these means, a 2 (valence: positive or negative) × 2 (prime: Mike or foil) within-participant interaction was calculated using the following

2. We also trimmed the data using a narrower response window of 300 ms to 1,000 ms and obtained the same pattern of results, and the critical tests remained significant for both Studies 1 and 2. We therefore used the longer response window to preserve as much of the data as possible.

equation: [(Foil-pos – Mike-pos) – (Foil-neg – Mike-neg)]. This index, which we refer to as *evaluative bias*, reflects the degree to which participants were faster to indicate a word was positive rather than negative following an image of Mike compared to the same difference for foils. Higher values indicate greater positivity toward Mike relative to foils.

Despite learning about Mike in the absence of any salient context, we predicted that context should affect evaluations of Mike. In a neutral context, we expected intelligence and athletic information should contribute equally to evaluations of Mike; however, in a classroom context, we predicted evaluations of Mike in the classroom should predominantly reflect knowledge of his intelligence. Consistent with this, we observed a marginal interaction effect of learning condition × context type, $F(2, 123) = 2.90$, $p = .06$, $\eta_p^2 = .05$. Although this test provided some evidence for our prediction, we conducted an ANOVA comprised of planned contrasts, which provided a more direct and nuanced test of our hypotheses. The first contrast compared participants who learned that Mike was athletic but unintelligent to participants who learned that Mike was intelligent but unathletic. To assess differences between these two groups, we created a contrast reflecting *type of mixed information* (*athletic-unintelligent* = -1, *athletic-intelligent* = 0, and *unathletic-intelligent* = +1). This contrast was critical to our hypothesis because both the *unathletic-intelligent* and *athletic-unintelligent* groups learned a mixture of domain-specific positive and negative information about Mike. In the absence of a salient context, participants who learned mixed-valence information should evaluate Mike similarly. However, in a classroom context, where intelligence is relevant, participants who believe Mike is smart (but unathletic) should evaluate him more positively than participants who think he is athletic (but unintelligent). Because these two groups should differ only when the classroom context is salient, we predict an interaction between context and type of mixed information. The second (residual) orthogonal contrast assessed the difference between participants who learned homogenously positive (*athletic-intelligent*) and participants who learned mixed-valence information (*athletic-unintelligent*, *unathletic-intelligent*). This *homogeneity of information* contrast (*athletic-intelligent* = 2, *unathletic-intelligent* = -1, *athletic-unintelligent* = -1) is not particularly relevant to hypotheses about context effects.

Effects of Type of Mixed Information. Main effects for context and type of information were not significant, Fs(1, 123) = 0.12, 1.12, ps = .74, .29, η_p^2 = .00, 01. As predicted, the context × type of mixed information interaction was significant, $F(1, 123) = 5.62$, $p = .02$, $\eta_p^2 = .04$ (Figure 1). In a neutral context, there was no effect of type of mixed information, $F(1, 123) = 0.90$, $p = .35$, $\eta_p^2 = .01$. The athletic-unintelligent and unathletic-intelligent versions of Mike were evaluated similarly (Ms = 4.02 and -23.16, respectively). But in the classroom context, there was a difference between conditions, $F(1, 123) = 5.66$, $p = .02$, $\eta_p^2 = .04$. Those who learned about a Mike who was smart (but unathletic; $M = 28.79$ ms) had more positive evaluations than those who learned about a Mike who was athletic (but not smart; $M = -47.24$). Thus, in the classroom, evaluations of Mike primarily reflected his intelligence (to the exclusion of his athleticism), presumably because this information was contextually relevant. We also examined the simple effect of context within both of the mixed-valence groups. Both effects were marginal. Participants who learned

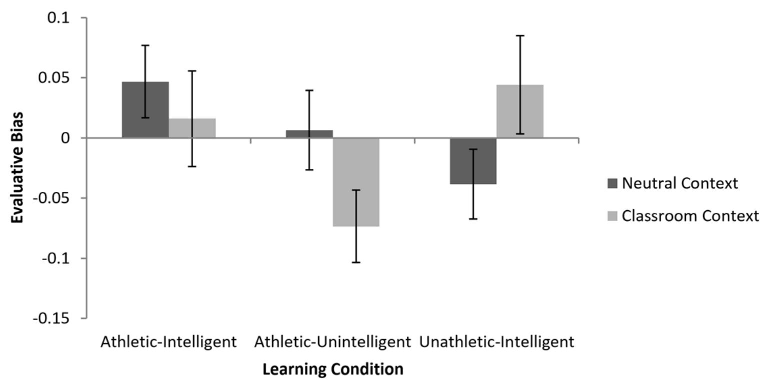

FIGURE 1. Evaluative bias as a function of learning condition in neutral and classroom contexts (Study 1).

Mike was athletic (but not smart) tended to evaluate him more negatively in the classroom than the neutral setting, $F(1, 123)$ 2.99, $p = .09$, $\eta_p^2 = .02$. Conversely, participants who learned Mike was smart (but not athletic) showed the opposite pattern and tended to evaluate Mike more positively in the classroom than neutral context, $F(1, 123) = 2.66$, $p = .10$, $\eta_p^2 = .02$.

Effects of Homogeneity of Information. The main effect of homogeneity of information (the residual contrast) was not significant, $F(1, 123) = 2.49$, $p = .12$, $\eta_p^2 = .02$, although there was trend indicating greater positivity among those who learned uniformly positive information about Mike compared to those who learned mixed-valence information. The context × homogeneity of information interaction did not approach significance, $F(1, 123) = 0.29$, $p = .58$, $\eta_p^2 = .01$.

Finally, we examined one additional contrast, comparing participants in the *athletic-intelligent* and *unathletic-intelligent* groups. Both of these groups learn Mike is intelligent, but differ in their knowledge of him as an athlete. Because athletic ability is irrelevant in the classroom, differences on this dimension should not affect implicit evaluations in this particular context. In line with this prediction, we found no evidence for a difference between these groups in the classroom context, $F(1, 123) = 0.33$, $p = .57$, $\eta_p^2 = .00$. This null result is consistent with the possibility that Mike's athleticism, which was high for one group of participants, but low for another, has little to do with evaluation in a context that emphasizes intellect. The difference between these same two groups was marginal in the neutral context, however: participants who learned uniformly positive information tended to feel more positively toward Mike than those who received mixed valence information, $F(1, 123) = 2.96$, p = .09, $\eta_p^2 = .02$.

TABLE 1. Manipulation Checks of Mike's Athleticism and Intelligence by Condition

	Athleticism		Intelligence	
	Mean	SD	Mean	SD
Study 1				
Athletic-Intelligent	2.59	0.79	2.55	0.76
Athletic-Unintelligent	2.51	1.06	-1.80	1.47
Unathletic-Intelligent	-2.38	0.91	2.69	1.08
Study 2				
Athletic-Unintelligent	2.16	1.26	-1.39	1.26
Unathletic-Intelligent	-1.38	1.66	1.78	1.54

STUDY 2

Study 2 sought to replicate the finding that context can influence attitudes even when encoding occurs in the absence of a salient or apparently meaningful context while providing a more rigorous test of the malleability of attitudes as a function of context during recall. Here, we varied context on a within-participant, trial-by-trial basis and included both athletic- and intelligence-related contexts. We predicted that participants would flexibly utilize knowledge about Mike's athleticism and intelligence, rapidly shifting attitudes depending on context. Additionally, we changed the way in which we manipulated context from Study 1. Rather than superimpose the target person onto intelligent- and athletic-related contexts, we presented images of athletic- or intelligence-related objects near the target person.

METHOD

Participants and Design. Ninety-eight undergraduates participated for course credit or monetary compensation. The average age was 21.02 (SD = 4.22). Three participants were identified as outliers (Cook's Ds were 0.093, 0.092, and 0.049) and one failed to correctly respond to enough trials to calculate the necessary indices, leaving 94 participants in the sample. The study employed a 2 (type of mixed information: *athletic-unintelligent* or *unathletic-intelligent*) × 2 (context type: athletic or intelligent) mixed-model design with context type varying within subject. Evaluative bias was the dependent variable.

Learning Task. We taught participants about Mike using the same paradigm as Study 1, but only included the *athletic-unintelligent* and *unathletic-intelligent* conditions. Two of the White male targets used in Study 1 were removed from the pool of targets because pretests indicated these males did not look convincingly athletic. Of the 8 remaining targets, the 3 who were judged most athletic were eligible to be Mike. Participants were randomly assigned one target who, for them, served as Mike and the remaining 7 targets served as foils during the evaluative priming task. During encoding, Mike was presented against a white background to ensure encoding was independent of the subsequent context manipulation.

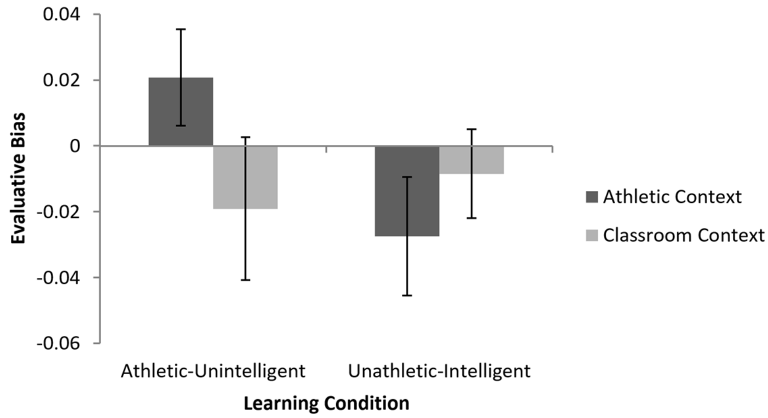

FIGURE 2. Evaluative bias as a function of learning condition in athletic-related and intelligence-related contexts (Study 2).

Evaluative Priming Paradigm. We assessed participants' implicit evaluations of Mike and foils in athletic- and intelligence-related contexts, which we manipulated by placing images of related objects in a randomly selected corner of the target person's photo. Objects included a baseball, football, helmet, soccer ball, tennis ball, racquet, books, calculator, graphing calculator, chess set, and microscope. This context manipulation was chosen to better match the complexity of the contexts (athletic-related contexts tend to be more visually complex and variable than intelligence-related contexts). Each trial began with a fixation cross (1,000 ms), followed by a prime (350 ms), a blank screen (150 ms), and a valenced target word. Valenced words were the same as in Study 1. Participants categorized words as "good" or "bad" using the keyboard (again counterbalanced) and accuracy feedback was delivered. A 4,000 ms inter-trial interval separated trials. Participants completed 33 practice and 216 test trials. The test trials consisted of 24 Mike trials (12 each of positive and negative target trials) and 84 foil trials (42 each of positive and negative target trials) in each of the two contexts. The number of trials was increased because context was varied within participants and to provide a more stable estimate.

Manipulation Check. Participants completed the same manipulation checks from Study 1.

RESULTS AND DISCUSSION

Manipulation Check. Analysis of the manipulation checks indicated that those exposed to information suggesting that Mike was athletic thought he was significantly more athletic than those who learned he was unathletic, $t(89) = 11.33$, $p < .001$, $\eta_p^2 = .59$.[3] Those who learned Mike was intelligent viewed him as significantly more intelligent than those who learned he was unintelligent, $t(89) = 10.77$, $p < .001$, $\eta_p^2 = .59$ (Table 1).

3. Three individuals failed to complete explicit measures, reducing the degrees of freedom.

Implicit Evaluations. We excluded incorrect trials (2.83%) and trials on which the response latencies were shorter than 300 ms (1.20%) or longer than 3,000 ms (0.80%). Remaining latencies were log-transformed and were used to compute evaluative bias scores following the same procedure in Study 1. We predicted participants who learned that Mike was smart (but unathletic) would evaluate him more favorably in the presence of objects like microscopes than athletic equipment, but participants who learned that Mike was athletic (but not smart) would show the opposite pattern. This prediction was supported by a significant type-of-mixed-information × context interaction, $F(1, 92) = 3.83$, $p = 0.05$, $\eta_p^2 = .04$ (Figure 2). Simple-effects tests revealed a marginal effect of context among participants in the *athletic-unintelligent* group, $F(1, 92) = 3.57$, $p = 0.06$, $\eta_p^2 = .04$, indicating a more positive reaction toward Mike in the athletic context (M = 12.65 ms) than in the intelligence context (M = -11.68 ms). The opposite pattern of means was observed in the *unathletic-intelligent* group, with more positive evaluation in the intelligence context (M = -5.41 ms) than the athletic context (M = -21.45 ms), but this effect was not significant, $F(1, 92) = 0.76$, $p = .39$, $\eta_p^2 = .01$. Next, we examined the simple effects of type of mixed information within context. In the athletic context, participants who learned Mike was athletic had more positive evaluations of him than those who learned he was unathletic, $F(1, 92) = 4.28$, $p = 0.04$, $\eta_p^2 = .04$. There simple effect of type of mixed information within the intelligent context was not significant, $F(1, 92) = 0.14$, $p = 0.70$, $\eta_p^2 = .00$. There were no main effects of context, $F(1, 92) = 0.55$, $p = .46$, $\eta_p^2 = .01$ or type of mixed information, $F(1, 92) = 0.99$, $p = .32$, $\eta_p^2 = .01$. Overall, this pattern of results is consistent with the notion that context effects can emerge even when contextual cues are not provided during the encoding process.

GENERAL DISCUSSION

The current research sought to test whether implicit attitudes will vary with a context cue only if that cue is present during attitude formation, a central prediction of the encoding framework of context effects in implicit attitudes (Gawronski et al., 2010). Participants learned about the athleticism and intelligence of a novel target, Mike, in a relatively neutral context. Subsequently, participants' implicit attitudes toward Mike were measured in contexts evocative of athleticism and intelligence. Although participants were not previously exposed to the target in those contexts, they selectively retrieved context-relevant information about Mike. Attitudes reflected information about Mike that was germane to the context presented at retrieval; they were largely unaffected by context-irrelevant information. Ultimately, these data suggest that context effects do not result exclusively from encoding contextual information into representations of attitude objects.

The current studies differed from previous research on context dependency in two key ways. First, unlike past studies that examined context dependency using known social groups (e.g., Blacks; Barden et al., 2004; Wittenbrink et al., 2001), we used a novel attitude object. This was an important step that added to the internal validity of our studies by allowing us to control participants' knowledge of the attitude object. Second, unlike previous research that coupled participants' learning about novel attitude objects with context (Gawronski et al., 2010; Rydell & Gawronski, 2009), participants in our studies learned about Mike in the absence of any immediately meaningful context. This was a crucial element of the current studies for three reasons. First, as we described earlier, individuals generally learn about and refine their attitudes in the same contexts in which they are subsequently measured (Rydell & Gawronski, 2009). In the current design we presented meaningful contextual information only at recall, effectively separating the measurement process from the learning process. Second, and more importantly, although the valence of information presented during encoding depended on an abstract conceptual domain (the domain of intelligence vs. the domain of athleticism), it did not depend on or covary with concrete visual cues. This is an absolutely critical aspect of these studies. If context and valence are always confounded during encoding, participants may develop two discrete representations that merge a specific visual context and an evaluation (e.g., blue-Bob-as-negative and yellow-Bob-as-positive). In our studies, because no meaningful context was presented during encoding, it is improbable that different evaluations of Mike were bound to any particular visual cue. For context effects to emerge at recall, the association between valence, domain, and visual context had to be more abstract. Third, participants in our ambivalent conditions learned both positive and negative information in the same context. This further reduces the likelihood that participants formed separate, context-dependent representations of Mike.

In our view, the novel aspect of the current work is the demonstration that (1) context can constrain the activation of evaluative information on the basis of its semantic relation to that information, and (2) such semantic context effects do not require the contexts to be present during encoding. This finding goes beyond the effects observed in previous research (Gawronski et al., 2010; Gawronski et al., 2014; Rydell & Gawronski, 2009), which are mainly concerned with context effects of meaningless visual cues that simply happened to be present during the encoding of counter-attitudinal information. Although both kinds of contexts may function as retrieval cues for specific pieces of information, the two kinds of contexts effects seem to be different in that context cues do not have to be present during encoding if they have a semantic relation to the content of attitudinal information.

Another important feature of the current research was the fact that attitudes were measured on a trial-by-trial basis in Study 2. As in Study 1, participants did not simply average the positive and negative information they had learned. Rather, they brought to mind whatever information was relevant given the context. Moreover, participants dynamically shifted between positive and negative attitudes on a moment-to-moment basis, responding within a few hundred milliseconds to whatever context was immediately presented.

Taken together, our studies contribute to a larger discussion regarding the stable versus flexible nature of attitudes (for a review, see Schwarz & Bohner, 2001). In particular, some researchers have likened the mental representation of attitudes to a file-drawer—upon encountering an attitude object in context, individuals retrieve the file representing that particular attitude object-context combination (e.g., Fazio, 2007; Wilson & Hodges, 1992). Much of the research we reviewed in the

Introduction could be seen as consistent with this perspective. It is conceivable, for example, that participants have separate "files" for Black church-goers and Black gangbangers (i.e., subtypes; Weber & Crocker, 1983). Likewise, it is possible that, when learning about Bob, participants formed two discrete representations—one based on the initial information learned and a second, oppositely valenced attitude that was stored as a subtype. This idea may also help explain the persistence of stereotypes in the face of counter-stereotypic information, as individuals may create subtypes based on stereotype inconsistent information (Kunda & Oleson, 1995, 1997).

Others have characterized attitudes in terms of connectionist networks (Bassili & Brown, 2005; McClelland, Rumelhart, & the PDP research group, 1986; Smith, 1996), which suggest that stored knowledge is represented in terms of weighted connections between units in a network. According to this view, inputs (e.g., context, mood, goals, etc.) determine which units become active at any given moment (and, just as critically, which units do *not* become active) and thus shape emergent patterns of activation that ultimately yield a summary judgment. Our data seem highly consistent with this connectionist account. When evaluating targets in a neutral context (Study 1), participants responded similarly to smart (but unathletic) and athletic (but unintelligent) versions of Mike. There was no observable difference between the two. A very different pattern emerged when Mike was evaluated in the presence of cues related to these two domains. Attitudes were influenced by contextual cues the participants had never seen before. These cues quickly and flexibly led to the activation of contextually relevant evaluative information. When participants judged Mike in a classroom setting, their attitudes shifted to reflect what they had learned about his intellect. Conversely, judgments in an athletic setting reflected what participants learned about Mike's athleticism.

Despite this, we acknowledge that it is still possible to account for our effects with a discrete, file-drawer model. Although participants learned about Mike in the absence of any real context, they may have spontaneously generated distinct representations for Mike-as-scholar and Mike-as-athlete, which were differentially primed by the classroom and athletic contexts. In our view, such an account offers a far less parsimonious explanation of the data. It is not clear, for example, how readily participants form such discrete representations or how fine-grained these representations could become. Did participants in our studies also form separate representations for Mike-as-English-scholar, Mike-as-math-scholar, and Mike-as-science-scholar? Clearly, questions about the precise nature of how attitudes are structured in memory are beyond the scope of the current studies. Future research on context effects may prove useful for shedding light on these issues.

REFERENCES

Barden, J., Maddux, W. W., Petty, R. E., & Brewer, M. B. (2004). Contextual moderation of racial bias: The impact of social roles on controlled and automatically activated attitudes. *Journal of Personality and Social Psychology, 87*, 5-22.

Bargh, J. A. (1997). The automaticity of everyday life. In R. S. Wyer, Jr. (Ed.), *The automaticity of everyday life: Advances in social cognition* (Vol. 10, pp. 1-61). Mahwah, NJ: Erlbaum.

Bassili, J. N., & Brown, R. D. (2005). Implicit and explicit attitudes: Research, challenges, and theory. In D. Albarracín, B. T. Johnson, & M. P. Zanna (Eds.), *Handbook of attitudes and attitude change* (pp. 543-574). Mahwah, NJ: Erlbaum.

Blair, I. V. (2002). The malleability of automatic stereotypes and prejudice. *Personality and Social Psychology Review, 6*, 242-261.

Blair, I. V., Ma, J. E., & Lenton, A. P. (2001). Imagining stereotypes away: The moderation of implicit stereotypes through mental imagery. *Journal of Personality and Social Psychology, 81*, 828-841.

Bouton, M. E. (1994). Conditioning, remembering, and forgetting. *Journal of Experimental Psychology: Animal Behavior Processes, 20*, 219-231.

Cesario, J., Plaks, J. E., Hagiwara, N., Navarrete, C. D., & Higgins, E. T. (2010). The ecology of automaticity: How situational contingencies shape action semantics and social behavior. *Psychological Science, 21*, 1311-1317.

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

Fazio, R. H. (2007). Attitudes as object-evaluation associations of varying strength. *Social Cognition, 25*, 603-637.

Fazio, R. H., Sanbonmatsu, D. M., Powell, M. C., & Kardes, F. R. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology, 50*, 229-238.

Gawronski, B., Rydell, R. J., Vervliet, B., & De Houwer, J. (2010). Generalization versus contextualization in automatic evaluation. *Journal of Experimental Psychology: General, 139*, 683-701.

Gawronski, B., Ye, Y., Rydell, R. J., & De Houwer, J. (2014). Formation, representation, and activation of contextualized attitudes. *Journal of Experimental Social Psychology, 54*, 188-203.

Judd, C. M., & Lusk, C. M. (1984). Knowledge structures and evaluative judgments: Effects of structural variables on judgmental extremity. *Journal of Personality and Social Psychology, 46*, 1193-1207.

Kerpelman, J. A., & Himmelfarb, S. (1971). Partial reinforcement effects in attitude acquisition and counterconditioning. *Journal of Personality and Social Psychology, 19*, 301-305.

Kunda, Z., & Oleson, K. C. (1995). Maintaining stereotypes in the face of disconfirmation: Constructing grounds for subtyping deviants. *Journal of Personality and Social Psychology, 68*, 565-579.

Kunda, Z., & Oleson, K. C. (1997). When exceptions prove the rule: How extremity of deviance determines the impact of deviant examples on stereotypes. *Journal of Personality and Social Psychology, 72*, 965-979.

Lowery, B. S., Hardin, C. D., & Sinclair, S. (2001). Social influence effects on automatic racial prejudice. *Journal of Personality and Social Psychology, 81*, 842-855.

McClelland, J. L., Rumelhart, D. E., & the PDP research group. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition. Volume II.* Cambridge, MA: MIT Press.

McConnell, A. R., Rydell, R. J., Strain, L. M., & Mackie, D. M. (2008). Forming implicit and explicit attitudes toward individuals: Social group association cues. *Journal of Personality and Social Psychology, 94*, 792-807.

Mitchell, J. P., Nosek, B. A., & Banaji, M. R. (2003). Contextual variations in implicit evaluation. *Journal of Experimental Psychology: General, 132*, 455-469.

Rutchick, A. M. (2010). Deus ex machina: The influence of polling place on voting behavior. *Political Psychology, 31*, 209-225.

Rydell, R. J., & Gawronski, B. (2009). I like you, I like you not: Understanding the formation of context-dependent automatic attitudes. *Cognition and Emotion, 23*, 1118-1152

Schuman, H., Presser, S., & Ludwig, J. (1981). Context effects on survey responses to questions about abortion. *Public Opinion Quarterly, 45*, 216-223.

Schwarz, N., & Bohner, G. (2001). The construction of attitudes. In A. Tesser & N. Schwarz (Eds.), *Blackwell handbook of social psychology, Vol. 1: Intraindividual processes* (pp. 436-457). Oxford, UK: Blackwell.

Schwarz, N., Strack, F., & Mai, H. (1991). Assimilation and contrast effects in part-whole question sequences: A conversational logic analysis. *Public Opinion Quarterly, 55*, 3-23.

Smith, E. R. (1996). What do connectionism and social psychology offer each other? *Journal of Personality and Social Psychology, 70*(5), 893-912.

Tesser, A. (1978). Self-generated attitude change. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 11, pp. 289-338). New York: Academic.

Thurstone, L. L. (1928). Attitudes can be measured. *American Journal of Sociology, 33*, 529-554.

Weber, R., & Crocker, J. (1983). Cognitive processes in the revision of stereotypic beliefs. *Journal of Personality and Social Psychology, 45*, 961-977.

Wilson, T. D., & Hodges, S. D. (1992). Attitudes as temporary constructions. In L. Martin & A. Tesser (Eds.), *The construction of social judgment* (pp. 37-65). Hillsdale, NJ: Erlbaum.

Wittenbrink, B., Judd, C. M., & Park, B. (2001). Spontaneous prejudice in context: Variability in automatically activated attitudes. *Journal of Personality and Social Psychology, 8*, 815-827.

Wittenbrink, B., & Schwarz, N. (2007). *Implicit measures of attitudes.* New York: Guilford.