Retrieval cues fail to influence contextualized evaluations

Ryan J. Hutchings, Jimmy Calanchini, Lisa M. Huang, Heather R. Rees, Andrew M. Rivers, Jenny Roth & Jeffrey W. Sherman

Check for updates

Retrieval cues fail to influence contextualized evaluations

Ryan J. Hutchings [a], Jimmy Calanchini[b,c], Lisa M. Huang[a], Heather R. Rees[a], Andrew M. Rivers[a], Jenny Roth[d] and Jeffrey W. Sherman[a]

[a]Department of Psychology, University of California, Davis, Davis, CA, USA; [b]Department of Social Psychology and Methodology, University of Freiburg, Freiburg im Breisgau, Germany; [c]Department of Psychology, University of California, Riverside, Riverside, CA, USA; [d]Department of Psychology, University of Würzburg, Würzburg, Germany

ABSTRACT

Initial evaluations generalise to new contexts, whereas counter-attitudinal evaluations are context-specific. Counter-attitudinal information may not change evaluations in new contexts because perceivers fail to retrieve counter-attitudinal cue-evaluation associations from memory outside the counter-attitudinal learning context. The current work examines whether an additional, counter-attitudinal retrieval cue can enhance the generalizability of counter-attitudinal evaluations. In four experiments, participants learned positive information about a target person, Bob, in one context, and then learned negative information about Bob in a different context. While learning the negative information, participants wore a wristband as a retrieval cue for counter-attitudinal Bob-negative associations. Participants then made speeded as well as deliberate evaluations of Bob while wearing or not wearing the wristband. Internal meta-analysis failed to find a reliable effect of the counter-attitudinal retrieval cue on speeded or deliberate evaluations, whereas the context cues influenced speeded and deliberate evaluations. Counter to predictions, counter-attitudinal retrieval cues did not disrupt the generalisation of first-learned evaluations or the context-specificity of second-learned evaluations (Experiments 2–4), but the counter-attitudinal retrieval cue did influence evaluations in the absence of context cues (Experiment 1). The current work provides initial evidence that additional counter-attitudinal retrieval cues fail to disrupt the renewal and generalizability of first-learned evaluations.

Evaluations can change as a function of context. All major attitude models include a mechanism for explaining how individuals can harbour different (and often inconsistent) evaluations of the same object (Gawronski et al., 2018). In these models, context controls which of the many possible attitudes will be expressed. Though context can moderate evaluations, there are circumstances in which evaluations remain consistent across contexts (Jones & Harris, 1967). For instance, if we learn that a person named Bob is helpful at work, we readily generate expectations that Bob will be a caring father, a supportive coach, and an accommodating client. It is only when Bob violates expectancies by performing unhelpful behaviour that perceivers look beyond

Bob, to the context, to better understand his unexpected behaviour (Roese & Sherman, 2007).

A considerable body of work has investigated these processes (Gawronski et al., 2018; Gawronski, Hu, Rydell, Vervliet, & De Houwer, 2015; Gawronski, Rydell, Vervliet, & De Houwer, 2010; Gawronski, Ye, Rydell, & De Houwer, 2014). In one prominent paradigm, participants learn that a person named Bob performs positive behaviours at the workplace (i.e. Context A), but later learn that he performs negative behaviours at the gym (i.e. Context B). Importantly, Bob performs an equal number of positive and negative behaviours in each context, respectively. The critical question is whether evaluations of Bob learned in Context A will change after learning new,

---

CONTACT Jeffrey W. Sherman ✉ jsherman@ucdavis.edu

counter-attitudinal information in Context B. Typically, evaluations of Bob in Context A remain positive, reflecting the (initial) positive information learned about Bob in that context and despite the negative information learned about him in Context B (cf. ABA renewal: Gawronski et al., 2015). When participants evaluate Bob in a novel context (Context C), evaluations are also positive, reflecting that the initial positive evaluation of Bob generalises to a new context (cf. ABC renewal: Gawronski et al., 2015). Only in Context B are evaluations of Bob negative, reflecting the negative information learned about him in that specific context alone (Gawronski et al., 2015). Taken together, these results suggest that initially learned information is uncontextualized, in that it generalises to evaluations made in a novel context. In contrast, second-learned, counter-attitudinal information is contextualised, in that it only influences evaluations made in the same context in which it was learned.

To account for these findings, retrieval theories have argued that counter-attitudinal (or second-learned) evaluations may be difficult to retrieve from memory outside of the counter-attitudinal learning context, because second-learned evaluations are represented as context-dependent associations (Bouton, 1993, 2010; Gawronski et al., 2018). Specifically, the Representational Theory of Contextualized Attitude Change (Gawronski et al., 2018, 2010) posits that initially-generated evaluations of a target person are encoded as context-independent evaluative representations. Context-independent representations should be activated in all contexts, because the activation of these representation is not linked to a specific context. In contrast, counter-attitudinal behaviours violate expectancies for the target's behaviour and, subsequently, draw attention to other contextual features that can explain the discrepancy (Roese & Sherman, 2007).This enhanced attention to the context binds the context cues and the counter-attitudinal evaluation, such that this evaluation is only activated in the presence of the contextual cues present during counter-attitudinal learning. Importantly, the theory argues that the counter-attitudinal context only indirectly influences evaluations of a target person by unlocking the activation of separate target-evaluation associations. Thus, the effects of context should be specific to target person, because the context cues are never directly associated with the outcome (i.e. positivity or negativity).

In contrast, several learning theories suggest that context cues can influence evaluations by acquiring direct context-evaluation associations (Miller & Matzel, 1988; Rescorla & Wagner, 1972). In this case, the context cues become directly associated with evaluations rather than setting the occasion for the activation of different target cue-evaluation associations. When the context cues and the target cue are combined, evaluations activated by the two cues sum to produce a new evaluation. In contrast to retrieval theories, these accounts suggest that the counter-attitudinal context becomes direct associated with positivity or negativity, and thus, the effects of the counter-attitudinal context should not be limited to the target person. A meta-analysis of experiments on contextualised attitude change demonstrated that contexts can influence evaluations through direct and indirect routes (Gawronski et al., 2015), although most extant research on contextualised attitude change focuses on the indirect route (Gawronski, Ye, et al., 2014).

Reducing contextualization

The finding that second-learned, counter-attitudinal evaluations are contextualised, whereas initially-learned evaluations are resilient and generalise to novel contexts, poses a problem for interventions aimed at affecting long-term attitude change. If, for example, a patient enters a clinic with the goal of overcoming a drug addiction, this work suggests that the effects of clinical treatment will be restricted to the context of the clinic (Context B) and will not persist when the patient returns to her daily life (Context A) or ventures into new situations (Context C). Consequently, substantial, cross-disciplinary research has explored conditions under which learning will or will not be contextualised (Cone, Mann, & Ferguson, 2017; Gawronski et al., 2018; Lai et al., 2014; Rosas, Todd, & Bouton, 2013).

In the current work, we explore a new strategy for decontextualising counter-attitudinal learning by using counter-attitudinal retrieval cues. Retrieval cues are features of the context that facilitate the recall of information learned in that context (Tulving & Thompson, 1973). In the present research, we examined whether a retrieval cue from the context in which counter-attitudinal information is learned can facilitate the retrieval and impact of counter-attitudinal evaluations in other contexts. If a counter-attitudinal retrieval cue can increase the generalizability of counter-attitudinal information, then the cue should influence evaluations of the target in both the initial

and novel contexts in the direction of the counter-attitudinal information. Thus, the goal of the present work was to examine whether counter-attitudinal retrieval cues can serve as a reminder for counter-attitudinal learning and, consequently, attenuate the resilience and generalisation of initial evaluations.

Previous work suggests that counter-attitudinal retrieval cues can enhance the generalizability of counter-attitudinal evaluations in human fear conditioning.[1] For example, Dibbets, Havermans, and Arntz (2008) had participants first view a target cue (e.g. a geometric shape) paired with an aversive outcome (e.g. a loud tone: see also, Vansteenwegen et al., 2006). The cue-outcome pairings were presented in one context (e.g. orange computer screen: Context A). In a second learning phase, participants viewed the target cue without the aversive outcome in a different context (e.g. green computer screen: Context B). During this second phase, participants also viewed an additional cue (e.g. an "&"), which served as a retrieval cue for this new learning. Participants were then tested for their conditioned fear to the target cue in Context A. Presenting the retrieval cue alongside the target cue reduced self-reported expectancies for the aversive outcome. Thus, retrieval cues appear to attenuate the renewal of initial fears within the original learning context. However, recent work has failed to find a retrieval cue effect in a fear conditioning procedure (Quezada et al., 2018).

Current work

Previous research has examined retrieval cues in fear conditioning, but no extant studies have examined the influence of counter-attitudinal retrieval cues in social evaluations, such as evaluations of individuals or groups. Moreover, the current work goes beyond previous retrieval cue research by examining the influence of counter-attitudinal retrieval cues under different measurement conditions. Dual-process theories of attitudes generally agree that measurement conditions determine which cognitive processes can influence evaluations (Fazio, 1990; Gawronski & Bodenhausen, 2006; Strack & Deutsch, 2004). For example, requiring participants to respond quickly should constrain the influence of slower and/or relatively resource-dependent cognitive processes that might otherwise contribute to responses. To date, most of the retrieval cue research in humans has measured self-reported fear under conditions that give participants as much time as they need to

respond. In the present research, participants evaluated target stimuli under two different measurement conditions: with and without time pressure. The speed with which a cognitive process can influence responses is one indication of the efficiency of that process (Moors & De Houwer, 2006). Short response deadlines may also undermine the intentionality of a response, presumably leading to greater influence of more impulsive evaluative responses. If the counter-attitudinal retrieval cue only influences deliberate evaluations, this pattern of results would suggest that retrieval cue effects manifest only when participants have enough time and resources to intentionally consider the influence of the cue.

Experiment 1

The purpose of Experiment 1 was to examine whether evaluations of a target person can be influenced by a counter-attitudinal retrieval cue. In the present research, we operationalised the counter-attitudinal retrieval cue as a wristband worn by participants during counter-attitudinal learning and subsequent attitude measurement.

Methods

Participants
One hundred and sixty-two undergraduates participated in this study in exchange for course credit. Data from three participants were excluded due to computer error and five participants for responding with the same key on every trial of the testing phase.[2] The final sample included 154 participants, 72.08% female, $M_{age} = 19.78$.

Learning procedure
The evaluative learning procedure implemented in the present research was adapted from previous work by Gawronski et al. (2010). In this task, participants passively viewed a series of written descriptions of behaviour attributed to a target person ("Bob"), presented below a picture of Bob, and were instructed to form an impression of him based on the presented information. On each trial, participants viewed a picture of Bob paired with a behavioural statement for 5,000 ms with a 1,000 ms interval displaying a blank-screen between trials. Behavioural information about Bob was presented in two blocks. In the first block of the learning procedure, 40 positive behaviours were attributed to Bob (e.g. "Bob became a

volunteer basketball coach for a children's home") and, in the second block, 40 negative behaviours were attributed to him (e.g. "Bob cheated on a take-home exam from the university"). Thus, all participants viewed an equal number of positive and negative behavioural statements.

Participants wore an orange rubber wristband during the second block of the learning procedure, which served as the counter-attitudinal retrieval cue. Research assistants verified that all participants complied with instructions to put the band on their left wrist after completing the first block of learning and before completing the second block. To examine the influence of the salience of the counter-attitudinal retrieval cue, we instructed participants that the wristband was unrelated to the current task in the low salience condition (i.e. "Please put this band on your left wrist because we want to see how it affects your performance on this task") or related to Bob's behaviour in the high salience condition (i.e. "Please put this band on your left wrist because Bob will behave differently when you have it on compared to when you do not"). Following the evaluative learning procedure, all participants took off their wristbands, which were collected by the experimenter. Next, participants completed an unrelated filler task in which they wrote a description of their kitchen. After three minutes, the computer automatically moved on to the next part of the experiment.

Speeded evaluation task (SET)

The SET was adapted from previous work by Gawronski et al. (2010) and is designed to measure rapid and repeated evaluations. Each trial began with a fixation cross displayed in the centre of the computer screen for 500 ms. The cross disappeared and was immediately replaced with either the same picture of Bob that participants had seen during the learning procedure or pictures of one of four novel male faces (i.e. novel targets). The novel targets were included to test whether the wristband exclusively influenced Bob or whether the cue affected both Bob and the novel targets. Each picture was displayed for 100 ms, then disappeared, and participants were prompted to indicate whether the picture was unpleasant or pleasant using the "A" key of the keyboard or the "5" key of the keypad, respectively. If a response was not registered within 800 ms, participants were instructed to respond more quickly.

The SET was completed in two blocks, each consisting of 30 Bob trials and 30 novel target trials, presented randomly. To examine whether the wristband would influence evaluations, participants wore the wristband for only one of the two blocks of the SET, counterbalanced such that half of the participants wore the wristband during the first block of the SET, and the other half of the participants wore the wristband during the second block of the SET.

Feeling thermometer

Participants then completed a feeling thermometer to assess their deliberate evaluations of Bob. They were instructed to enter a number between 0–100 representing their level of warmth toward Bob, such that zero reflected highly negative evaluations, fifty reflected neutral evaluations, and one hundred reflected highly positive evaluations. To examine the influence of the wristband on deliberate evaluations, half of the participants wore the wristband while completing the feeling thermometer, whereas the remaining participants completed the feeling thermometer without the wristband. Participants who wore the wristband during the second block of the SET also wore the wristband while making deliberate evaluations of Bob.

Data Availability statement

The data and analysis code used in all the reported studies is available in the OSF repository (DOI 10.17605/OSF.IO/C9UXB).

Results

Speeded evaluations

Speeded evaluations of Bob were examined in a 2 (*Wristband at Evaluation:* on, off) × 2 (*Wristband Order*: wristband worn during first block of evaluation, wristband worn during second block of evaluation) × 2 (*Wristband Salience*: low salience, high salience) × 2 (*Target:* Bob, novel targets) mixed ANOVA, with wristband at evaluation and target as within-participants factors and wristband order and wristband salience as a between-participants factor. The wristband salience variable did not have any influence on our results. Therefore, we excluded it from further analyses. Because all participants wore the wristband during the second block of the learning procedure, in which they viewed only negative information about Bob, we predicted that evaluations of Bob would be more negative when participants wore versus did not wear the wristband. Given that

participants wore the wristband when learning about Bob, we predicted that the wristband would influence evaluations of Bob but have no influence on evaluations of the novel targets. An evaluative index was computed separately for Bob and the novel targets for each block based on the proportion of pleasant responses, such that higher values reflect more positive evaluations.[3]

There was a main effect of wristband at evaluation, $F(1, 152) = 3.90$, $p = 0.05$, $\hat{n}_G^2 = .001$. Participants evaluated Bob and the novel targets more negatively when the wristband was on ($M = .52$, $SD = .09$) than when it was off ($M = .55$, $SD = .09$). There also was a significant main effect of target, $F(1, 182) = 18.02$, $p < .001$, $\hat{n}_G^2 = .05$, such that participants evaluated Bob ($M = .46$, $SD = .31$) more negatively than the novel targets ($M = .61$, $SD = .31$).

These main effects were qualified by a significant target × wristband at evaluation interaction, $F(1, 152) = 6.58$, $p = .01$, $\hat{n}_G^2 = .002$ (see Figure 1). To deconstruct this interaction, we separately examined the influence of the wristband on ratings of Bob and the novel targets. Participants evaluated Bob more negatively when the wristband was on ($M = .48$, $SD = .29$) than when it was off ($M = .48$, $SD = .27$), $t(153) = 2.81$, $p = 0.006$, $d = 0.23$, 95% CI [0.07, 0.39]. In contrast, the wristband had no influence on evaluations of the novel targets, $t(153) < 1$.

This two-way interaction was moderated by a significant wristband at evaluation × wristband order × target interaction, $F(1, 152) = 4.20$, $p = .04$, $\hat{n}_G^2 = .001$. To deconstruct this interaction, we separately examined the influence of wristband at evaluation and target for each wristband order condition. When participants wore the wristband during the first block of the SET, there was a significant main effect of target, $F(1, 80) = 14.82$, $p = .0002$, $\hat{n}_G^2 = .08$, but there was no significant main effect of wristband at evaluation nor the predicted wristband at evaluation by target interaction. In contrast, when participants wore the wristband during the second block of the SET, there was a significant main effect of target, $F(1, 72) = 5.04$, $p = .03$, $\hat{n}_G^2 = .03$, and the predicted wristband at evaluation × target interaction was significant, $F(1, 72) = 7.30$, $p = .009$, $\hat{n}_G^2 = .006$. This interaction revealed that participants evaluated Bob more negatively when the wristband was on ($M = .44$, $SD = .28$) than when it was off ($M = .51$, $SD = .31$), $t(72) = 2.82$, $p = .006$, $d = 0.33$, 95% CI [0.09, 0.57]. Evaluations of the novel targets were not significantly influenced by the wristband, $t(72) = 1.19$, $p = 0.24$.

Deliberate evaluations

In contrast to the speeded evaluations, the wristband had no influence on deliberate evaluations of Bob, $t(151) = 0.29$, $p = .77$. Evaluations of Bob when wearing the wristband ($M = 42.34$, $SD = 17.09$) were not different than evaluations of Bob when the wristband was not worn ($M = 41.50$, $SD = 18.18$).

Discussion

In Experiment 1, the counter-attitudinal retrieval cue influenced speeded evaluations of Bob under some conditions. When participants wore the counter-attitudinal retrieval cue during the second block of the SET, participants evaluated Bob more negatively when wearing versus not wearing the retrieval cue. However, when participants wore the counter-attitudinal retrieval cue during the first block of the SET, the retrieval cue had no influence on evaluations of Bob. This pattern of results may have occurred because negative information is stickier than positive information (Ledgerwood & Boydstun, 2014), and the counter-attitudinal retrieval cue was paired with negative information. Nevertheless, Experiment 1 provides initial evidence that retrieval cues can influence evaluations of a target person. Moreover, and in line with predictions, the counter-attitudinal retrieval cue did not influence speeded evaluations of the novel targets. The failure of the counter-attitudinal retrieval cue to generalise to new targets suggests that it did not acquire a general negativity. Instead, it seems that the cue only acts to set the occasion for the expression of different evaluations of Bob (Gawronski et al., 2018). Finally, the counter-attitudinal retrieval cue influenced speeded but not deliberate evaluations of Bob. These results suggest that participants may have corrected against the influence of the counter-attitudinal retrieval cue when given the opportunity to deliberate (Gawronski & Bodenhausen, 2006). Moreover, participants may correct against the influence of the counter-attitudinal retrieval cue because they do not view the cue as a valid basis for evaluations of Bob.

Experiment 2

Experiment 1 provides an initial demonstration that a counter-attitudinal retrieval cue (i.e. the wristband) can shift evaluations in the direction of the information that was paired with the retrieval cue. Building
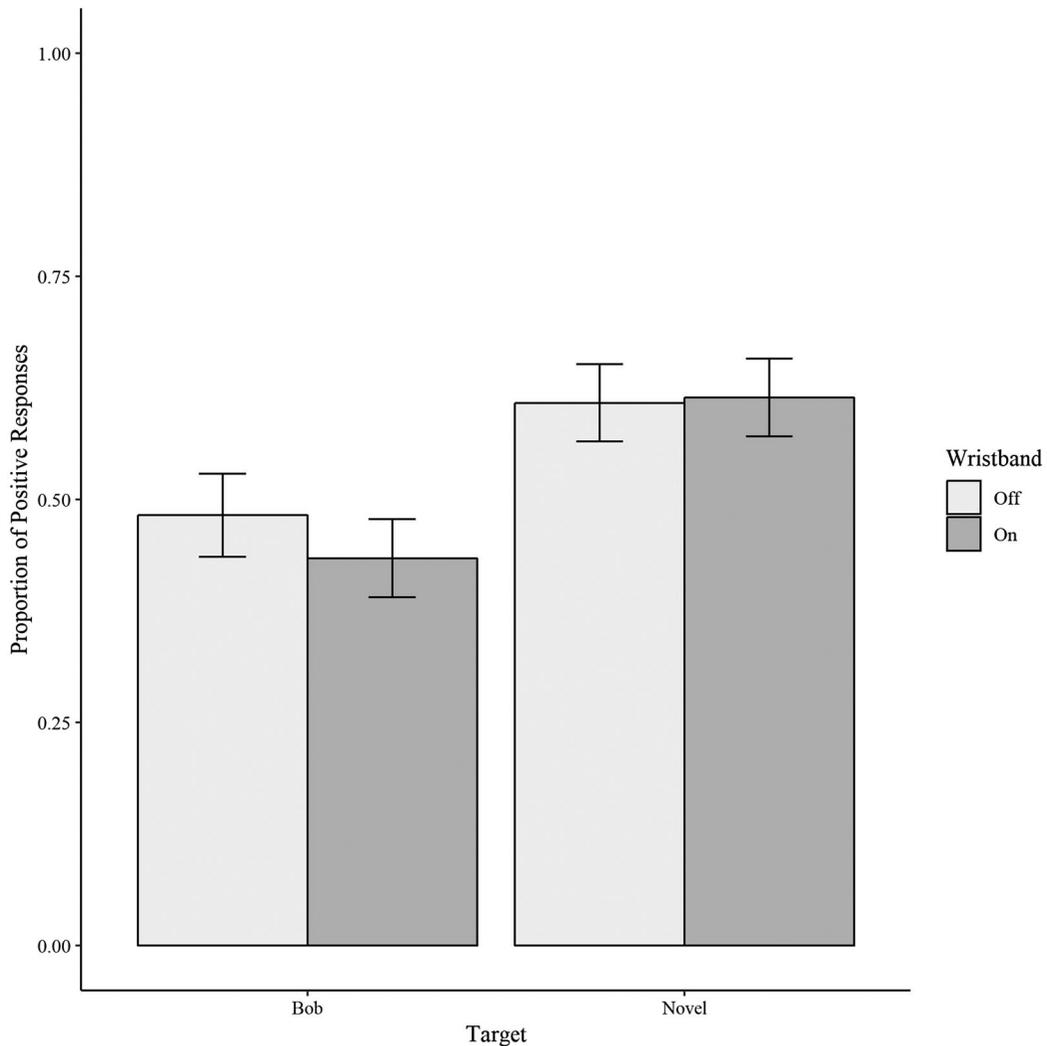
Figure 1. The influence of wristband at evaluation on speeded evaluations of Bob and novel targets in Experiment 1. On = wristband on during speeded evaluation. Off = wristband off during speeded evaluation. Proportion of Positive Responses = the proportion of "more positive" responses on the SET.

upon these findings, the purpose of Experiment 2 was to examine whether counter-attitudinal retrieval cues facilitate the generalisation of second-learned evaluations to the initial-learning context and a novel context. To do so, we manipulated the context in which participants learned positive versus negative information about Bob: Positive behaviours were attributed to Bob in one context, and negative behaviours were attributed to him in a different context. Participants always wore the retrieval cue while learning counter-attitudinal information about Bob. As such, the wristband served as a cue for the negative evaluation formed in this context. We then measured

participants' evaluations of Bob in the initial (positive) context, second-learned (negative) context, and a novel context, both in the absence and presence of the counter-attitudinal retrieval cue. We predicted that the counter-attitudinal retrieval cue would attenuate the renewal of initial positive evaluations of Bob in the first-learned context, and that the cue would enhance the generalizability of negative evaluations to the novel context. Because both the second-learned context and the retrieval cue were present when participants learned negative information about Bob, we predicted that participants would evaluate Bob negatively when measured in the

second-learned context – though whether the retrieval cue would make evaluations in the second-learned context even more negative remained an open question. As in Experiment 1, we predicted that the counter-attitudinal retrieval cue would only influence evaluations of Bob, and not novel targets. Finally, we predicted that the counter-attitudinal retrieval cue would influence speeded, but not deliberate, evaluations.

Method

Participants and design

Based on the findings of Experiment 1, we used G*power 3.1 (Faul, Erdfelder, Lang, & Buchner, 2007) to calculate a target N of 260 with 95% power to detect a main effect of the wristband with a Cohen's d_z of 0.32. Two hundred and sixty undergraduate students participated in this experiment in exchange for course credit. Data from five participants were excluded for not following instructions, nine for pressing the same key on every trial of the testing phase, and one due to computer error. The final sample was 245 participants, 79.76% female, $M_{age} = 19.92$.

The design of Experiment 2 was a 3 (Context: positive context, negative context, and novel context) × 2 (Wristband at Evaluation: on, off) × 2 (Wristband Order: wristband worn during first block of evaluation, wristband-second) × 2 (Target: Bob, novel targets) mixed design, with wristband order as the only between-participants condition.

Learning procedure

The instructions and general setup for the learning procedure mirrored Experiment 1, with the addition of context cues, which we operationalised as the background colour against which the target was presented (see also Gawronski et al., 2010). Participants first viewed a series of forty positive behaviours attributed to Bob, presented against either a blue or yellow background screen colour (i.e. positive context). After completing the first block of the learning procedure, participants were instructed to put on the wristband (i.e. counter-attitudinal retrieval cue). Then, in the second block of the learning procedure, participants viewed a series of forty negative behaviours attributed to Bob in a new context (i.e. either a blue or yellow background, whichever was not presented during the first block: negative context).[4] After finishing both blocks of the learning procedure, participants took off their wristbands and then completed the

same filler task as in Experiment 1 before completing the dependent measures.

Speeded evaluations

The SET was identical to Experiment 1, except for the inclusion of background colours. In each of the two blocks of the SET, participants responded to thirty Bob and thirty novel target trials. On one third of the trials, Bob and the novel targets were depicted against the background colour from the first (i.e. positive) learning block. On an additional third of the trials, Bob and the novel targets were depicted against the second (i.e. negative) learning background colour. On the final third of the trials, Bob and the novel targets were depicted against a novel (i.e. white) background colour. All trials were presented in a random order. The novel context trials examined whether evaluations from the positive or negative learning context would generalise to a new context. To determine whether wearing the wristband would moderate evaluations, participants only wore the wristband for one block of the SET. Half of the participants completed the first block of the SET wearing the wristband, whereas the remaining participants completed the second block of the SET wearing the wristband.

Feeling thermometer

Following the SET, participants completed deliberate evaluations of Bob presented in each of the different contexts (i.e. positive, negative, novel). Participants reported their evaluations of Bob depicted against the positive, negative, and novel context, presented in a randomised order, on a scale ranging from 1–9, with higher values reflecting more positive evaluations. To examine the influence of the wristband on deliberate evaluations, half of the participants wore the wristband while reporting deliberate evaluations, whereas the remaining participants reported deliberate evaluations without the wristband. Participants who wore the wristband during the second block of the SET also wore the wristband when generating deliberate evaluations of Bob.

Results
Speeded evaluations

Rapid evaluations of the target were examined in a 3 (Context: positive, negative, and novel) × 2 (Wristband at Evaluation: on, off) × 2 (Wristband Order: wristband worn during first block of evaluation, wristband worn during second block of evaluation) × 2 (Target: Bob, novel targets) mixed ANOVA, with wristband at

evaluation, context, and target as within-participants factors and wristband order as a between-participants factor. Evaluative indices were computed for each of the three contexts in the same way as in Experiment 1.

In line with predictions, there was a significant main effect of context, $F(1.59, 386.07) = 73.16$, $p < .001$, $\hat{n}_G^2 = .04$.[5] Evaluations were more negative in the negative context ($M = .44$, $SD = .19$) than the novel context ($M = .57$, $SD = .12$), $t(979) = -13.9$, $p < .001$, $d = 0.44$, 95% CI [0.38, 0.51], and evaluations in the novel context were more negative than evaluations in the positive context ($M = .61$, $SD = .15$), $t(979) = 5.03$, $p < .001$, $d = 0.16$, 95% CI [0.10, 0.22]. There also was a significant main effect of target, $F(1, 243) = 12.52$, $p < .001$, $\hat{n}_G^2 = .02$. Participants evaluated Bob more positively ($M = .59$, $SD = .29$) than the novel targets ($M = .49$, $SD = .29$).

There was a significant context × target two-way interaction, $F(2, 486) = 4.20$, $p = .02$, $\hat{n}_G^2 = .004$. Unexpectedly, participants evaluated Bob more positively than the novel targets in all contexts, but the size of this difference was attenuated in the negative context.

There were two significant three-way interactions. There was a significant wristband at evaluation × wristband order × target interaction, $F(1, 243) = 5.23$, $p = .02$, $\hat{n}_G^2 = .0005$. To decompose this interaction, we examined the influence of wristband at evaluation and target separately for each order condition. When participants wore the wristband during the first block of the SET, there was a significant main effect of wristband at evaluation, $F(1, 128) = 4.44$, $p = .04$, $\hat{n}_G^2 = .002$, and a significant main effect of target, $F(1, 128) = 12.63$, $p < .001$, $\hat{n}_G^2 = .04$. These main effects were qualified by a marginally significant wristband at evaluation × target interaction, $F(1, 128) = 3.51$, $p = .06$, $\hat{n}_G^2 = .001$. Whereas Wristband at Evaluation had no influence on evaluations of Bob, $t(128) < 1$, evaluations of the novel targets were more negative when wearing the wristband ($M = .51$, $SD = .25$) than when not wearing it ($M = .48$, $SD = .23$), $t(128) = 2.63$, $p = .01$, $d = 0.23$, 95% CI [0.06, 0.41]. When participants wore the wristband during the second block of the SET, there were no significant main effects or interactions.

There was also a significant wristband at evaluation × wristband order × context interaction, $F(2, 486) = 5.79$, $p = 0.003$, $\hat{n}_G^2 = .0008$. To examine this interaction, we examined the influence of wristband at evaluation and context separately for each order condition. When participants wore the wristband during the first block of the SET, there was a significant

main effect of wristband at evaluation, $F(1, 128) = 4.44$, $p = .04$, $\hat{n}_G^2 = .002$. Participants evaluated both Bob and the novel targets more negatively when wearing the wristband ($M = .55$, $SD = .08$) than when not wearing it ($M = .57$, $SD = .08$). There was also a significant main effect of context, $F(2, 256) = 36.25$, $p < .001$, $\hat{n}_G^2 = .07$. The wristband at evaluation × context interaction was not significant $F(2, 256) = 2.09$, $p = .13$, $\hat{n}_G^2 = .0001$.

When participants wore the wristband during the second block of the SET, there was a significant main effect of context, $F(2, 230) = 37.07$, $p < .001$, $\hat{n}_G^2 = .08$, and this main effect was moderated by a significant wristband at evaluation × context interaction, $F(2, 230) = 3.69$, $p = .03$, $\hat{n}_G^2 = .002$ (see Figure 2). Participants evaluated both Bob and the novel targets more negatively when wearing the wristband in the negative context ($M = .44$, $SD = .17$) than when not wearing it ($M = .40$, $SD = .21$), $t(115) = 2.18$, $p = .03$, $d = 0.20$, 95% CI [0.06, 0.46]. The wristband did not influence evaluations of Bob or the novel targets in the novel context, $t(115) < 1$. Contrary to predictions, participants evaluated Bob and the novel targets marginally more positively when wearing the wristband than when not, $t(115) = 1.76$, $p = .08$.

Deliberate evaluations

Deliberate evaluations of Bob were examined in a 2 (Wristband: on, off) × 3 (Context: positive, negative, neutral) mixed ANOVA, with context as a within-participants factor and wristband as a between-participants factor. In line with the speeded evaluations, there was a significant effect of context on deliberate evaluations of Bob, $F(1.53, 371.64) = 110.00$, $p < .001$, $\hat{n}_G^2 = .13$. Evaluations were more negative in the negative context ($M = 3.89$, $SD = 1.69$) than in the neutral context ($M = 5.15$, $SD = 0.97$), $t(244) = 10.62$, $p < .001$, $d = 0.68$ 95% CI [0.54, 0.82], and evaluations in the neutral context were more negative than in the positive context ($M = 5.74$, $SD = 1.55$), $t(244) = 5.86$, $p < .001$, $d = 0.37$, 95% CI [0.24, 0.50]. There were no other significant main effects or interactions, Fs $(1,243) < 1$.

Discussion

The purpose of Experiment 2 was to examine whether counter-attitudinal retrieval cues attenuate the renewal and generalisation of initial evaluations. In contrast to predictions, the counter-attitudinal retrieval cue did not reduce the renewal or generalisation
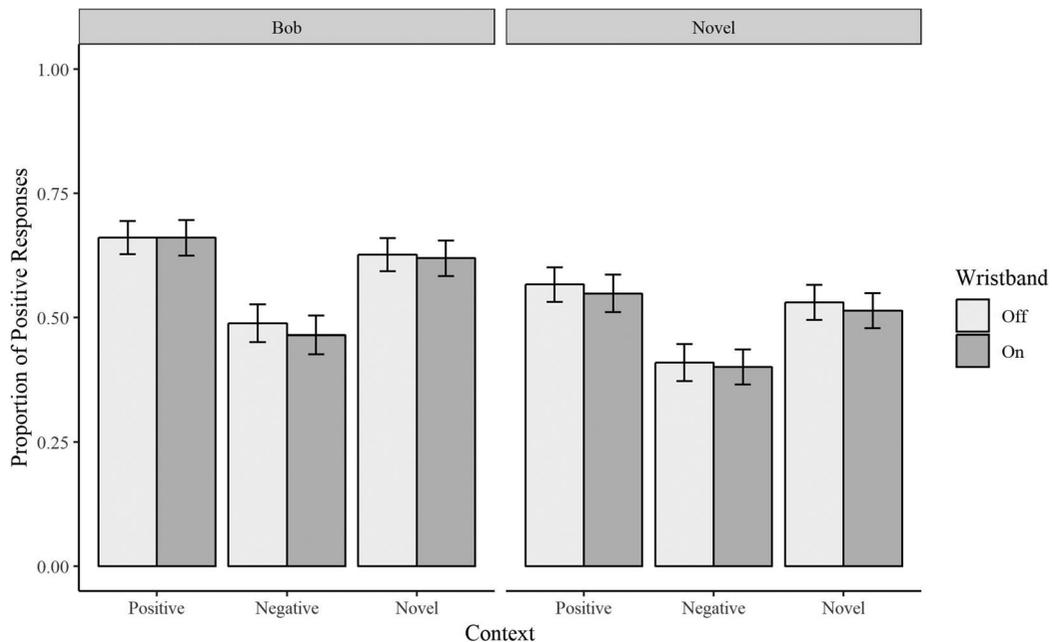
Figure 2. The influence of wristband at evaluation and context on evaluations of Bob and novel targets in Experiment 2. On = wristband on during speeded evaluation. Off = wristband off during speeded evaluation. Proportion of Positive Responses = the proportion of "more positive" responses on the SET. Positive = positive (initial) learning context. Negative = negative (second) learning context. Novel = novel context.

of initial evaluations. That said, the counter-attitudinal retrieval cue did influence evaluations under some conditions. When participants wore the counter-attitudinal retrieval cue during the second block of the SET, they evaluated all targets (i.e. Bob and the novel targets) more negatively when wearing the retrieval cue than when not wearing it. However, this effect was limited to the negative context; the counter-attitudinal retrieval cue did not influence evaluations in the positive or novel contexts. Moreover, when participants wore the counter-attitudinal retrieval cue during the first block of the SET, the retrieval cue did not influence evaluations of Bob or the novel targets. In contrast, the context cues (i.e. background screen colours) strongly influenced evaluations of both Bob and the novel targets.

In contrast to Experiment 1, the counter-attitudinal retrieval cue did not specifically influence evaluations of Bob. Instead, the retrieval cue had a similar effect on evaluations of Bob and the novel targets. Similarly, the context cues influenced evaluations of Bob and the novel targets. This pattern of results conflicts with the predictions of retrieval-based accounts of retrieval cues (Bouton, 1993, 2010; Gawronski et al., 2018). Retrieval-based accounts posit that retrieval cues (i.e. the counter-attitudinal retrieval cue and context cues) set the occasion for different evaluations of a target cue (i.e. the cue presented with the retrieval cue; Bob, in this case). As such, retrieval cues modulate target cue-outcome associations rather than acquire direct associations with the outcome, and thus, retrieval cues should only influence evaluations of the target cue (see also, Holland, 1992). In Experiment 2, the retrieval cues did not exclusively influence the target cue (i.e. Bob). The current pattern of results suggests that the retrieval cues acquired direct associations with positivity and negativity that generalised to novel targets – much like any other stimulus in this task. The evaluations of Bob and the evaluations of the context cues summed to produce distinct evaluations in different contexts.

In Experiment 2, the counter-attitudinal retrieval cue had a much smaller influence on evaluations of Bob when other retrieval cues (i.e. the context cues) were present than they did in Experiment 1, when no other cues were present. One possible interpretation of these findings is that the context cues were more salient than the counter-attitudinal retrieval cue. Indeed, previous research indicates that the salience of context cues at learning moderates their influence on evaluations. Typically, context cues are ignored during initial learning but become salient

when counter-attitudinal information is presented (Gawronski et al., 2010; Gawronski, Ye, et al., 2014). In Experiment 2, we introduced the retrieval cue when counter-attitudinal information was presented, but participants' attention was likely directed to the context cue at this time and, perhaps consequently, away from the retrieval cue. To test this possibility in Experiment 3, we manipulated the salience of the counter-attitudinal retrieval cue. We predicted that explicitly linking the counter-attitudinal retrieval cue to Bob's behaviour would enhance the salience of the cue (Dibbets, Moor, & Voncken, 2013). Moreover, the counter-attitudinal retrieval cue may only influence evaluations of Bob when we link the counter-attitudinal retrieval cue directly to him. Thus, in Experiment 3, we manipulated the salience of the retrieval cue by instructing participants that the retrieval cue was either related or unrelated to the task. The retrieval cue should be more impactful when it is made salient.

Additionally, we predicted that the influence of the counter-attitudinal retrieval cue will generalise to the novel context when the initial-evaluation is contextualised. Previous work has demonstrated that the second-learned evaluation has a greater influence on evaluations in the novel context when the first-learned evaluation is contextualised (Gawronski et al., 2010). Thus, in Experiment 3, we utilised a procedure for contextualising initial-evaluations of Bob.

Experiment 3

In Experiment 2, the counter-attitudinal retrieval cue had limited effects on evaluations of Bob and the novel targets. In Experiment 3, we examined if the counter-attitudinal retrieval cue would be more impactful when it was more salient. Additionally, we examined if the counter-attitudinal retrieval cue would have a greater influence on evaluations in the novel context when the initial-evaluation was contextualised.

Method

Participants and design
Power analyses indicated that our target N for Experiment 3 should be 260 participants to achieve 95% likelihood to detect an effect of the counter-attitudinal retrieval cue with a Cohen's $d_z = 0.32$. We were only able to recruit 241 undergraduate students before the end of the academic term, who participated in

exchange for course credit. We excluded data from three participants for not following instructions, one due to experimenter error, one due to colourblindness, and two for responding with the same-key on every trial of the SET. The final sample was 234 students, 75.32% female, $M_{age} = 19.91$.

Context salience manipulation
All participants completed a manipulation to enhance attention to the context cues and facilitate the contextualisation of initial-evaluations of Bob (adapted from, Gawronski et al., 2010). Specifically, participants viewed a series of twenty positive behaviours attributed to "Jim" presented in one colour context, followed by twenty negative behaviours attributed to him in a different colour context. This paradigm was identical to the "Bob" learning task used in Experiments 1 and 2, with two key differences: the target of learning was "Jim" rather than "Bob", and the background colours for the Jim learning task (i.e. green, brown) were different from the background colours for the two learning blocks of the Bob learning task (i.e. yellow, blue). In doing so, we avoid any potential carry-over effects to the subsequent learning or evaluation tasks based on target or background colour. Because Jim's behaviour covaries perfectly with context cues (i.e. background colours), this manipulation makes context cues salient when participants begin the subsequent Bob learning task and should facilitate the contextualisation of initial-evaluations of Bob (Gawronski et al., 2010).

Learning procedure
The Bob learning task was identical to Experiment 2, with the addition of one manipulation. Participants were given instructions intended to manipulate the salience of the wristband cue (i.e. counter-attitudinal retrieval cue). Participants in the low wristband salience condition were instructed: "Please put this band on your left wrist because we want to see how it affects your performance on this task". In contrast, participants in the high wristband salience condition were told: "Please put this band on your left wrist because Bob will behave differently when you have it on compared to when you do not". The wristband salience manipulations in Experiments 1 and 3 were identical. Although the wristband salience manipulation did not influence responses in Experiment 1, we predicted that wristband salience should influence responses when the context cues and the

counter-attitudinal retrieval cue are simultaneously presented.

Speeded evaluations

The speeded evaluation task was identical to Experiment 2, with one exception: All participants wore the wristband during the second block of the SET. Finally, we did not collect deliberate evaluations of Bob.

Results

SET

Speeded evaluations were examined in a 3 (*Context*: positive, negative, neutral) $\times$ 2 (*Wristband at Evaluation*: off, on) $\times$ 2 (*Wristband salience*: low, high) $\times$ 2 (*Target*: Bob, novel targets) mixed ANOVA with wristband salience as the only between-participants factor. An evaluative index was computed in the same way as in Experiment 2. In contrast to Experiment 2, we predicted that participants would evaluate Bob more negatively when wearing the counter-attitudinal retrieval cue than when not wearing the retrieval cue. To the extent that the context cues overwhelmed the influence of the counter-attitudinal retrieval cue in Experiment 2, we predicted that the retrieval cue would be more impactful when the cue was more versus less salient in Experiment 3.

As in the previous study, there was a main effect of context, $F(1.75, 405.72) = 23.06$, $p < .001$, $\hat{n}_G^2 = .01$.[6] Evaluations in the negative context ($M = .43$, $SD = .16$) were more negative than evaluations in the novel context ($M = .49$, $SD = .11$), $t(935) = 7.79$, $p < .001$, $d = 0.25$, 95% CI [0.19, 0.32], and evaluations in the novel context were more negative than evaluations in the positive context ($M = .51$, $SD = .15$), $t(935) = 2.46$, $p = .01$, $d = 0.08$, 95% CI [0.02, 0.14]. There was a marginally significant main effect of Wristband at Evaluation, $F(1, 232) = 3.72$, $p = .05$, $\hat{n}_G^2 = .0006$. Participants evaluated Bob and the novel targets more negatively when wearing the wristband ($M = .47$, $SD = .09$) than when not wearing the wristband ($M = .49$, $SD = .09$).

There were three significant two-way interactions. There was a significant Wristband at Evaluation $\times$ Wristband Salience interaction, $F(1, 232) = 4.90$, $p = .03$, $\hat{n}_G^2 = .0007$. When the wristband salience was low, the wristband had no influence on evaluations, $t(119) < 1$. In contrast, when wristband salience was high, participants evaluated Bob and the novel targets more negatively when wearing the wristband ($M = .47$, $SD = .07$) than when not wearing the

wristband ($M = .51$, $SD = .07$), $t(113) = 2.94$, $p = 0.004$, $d = 0.28$, 95% CI [0.09, 0.46].

There also was a significant interaction between context and wristband at evaluation, $F(1.94, 450.73) = 3.15$, $p = .05$, $\hat{n}_G^2 = .0005$. Participants evaluated Bob and the novel targets more negatively when wearing the wristband in the negative context ($M = .41$, $SD = .19$) relative to when participants did not wear the wristband ($M = .45$, $SD = .17$), $t(233) = 3.12$, $p = .002$, $d = 0.20$, 95% CI [0.07, 0.33]. The wristband had no influence on evaluations in the novel or positive contexts, $ts(233) < 1$.

Finally, there was a significant interaction between target and wristband at evaluation, $F(1, 232) = 29.77$, $p < .001$, $\hat{n}_G^2 = .004$ (see Figure 3). Participants evaluated Bob more negatively when wearing the wristband than when not wearing it, $t(233) = 5.18$, $p < .001$, $d = 0.34$, 95% CI [0.21, 0.47]. Unexpectedly, participants evaluated the novel targets more positively when wearing the wristband than when not wearing it, $t(233) = 2.33$, $p = .02$, $d = 0.15$, 95% CI [0.02, 0.28].

Discussion

In line with predictions, increasing the salience of the counter-attitudinal retrieval cue enhanced the effectiveness of the retrieval cue. However, we did not predict that increasing the salience of the counter-attitudinal retrieval cue would impact evaluations of Bob and the novel targets. It is possible that explicitly linking the retrieval cue to Bob's behaviour enhanced perceived importance of the cue, and the cue still acquired a direct association with negativity. Although salience influenced both Bob and the novel targets, participants evaluated Bob and the novel targets differently when wearing the counter-attitudinal retrieval cue. When wearing the counter-attitudinal retrieval cue, participants evaluated Bob more negatively than when not wearing it. In contrast, participants evaluated the novel targets more positively when wearing the counter-attitudinal retrieval cue than when not wearing it. As in Experiment 1, the (intended) influence of the counter-attitudinal retrieval cue only extended to a target paired with the retrieval cue during the learning phase. That said, the retrieval cue did not disrupt the effect of context cues on evaluations. Even when wearing the wristband, participants evaluated Bob and the novel targets more positively in both the positive and novel contexts than in the negative context. Moreover,
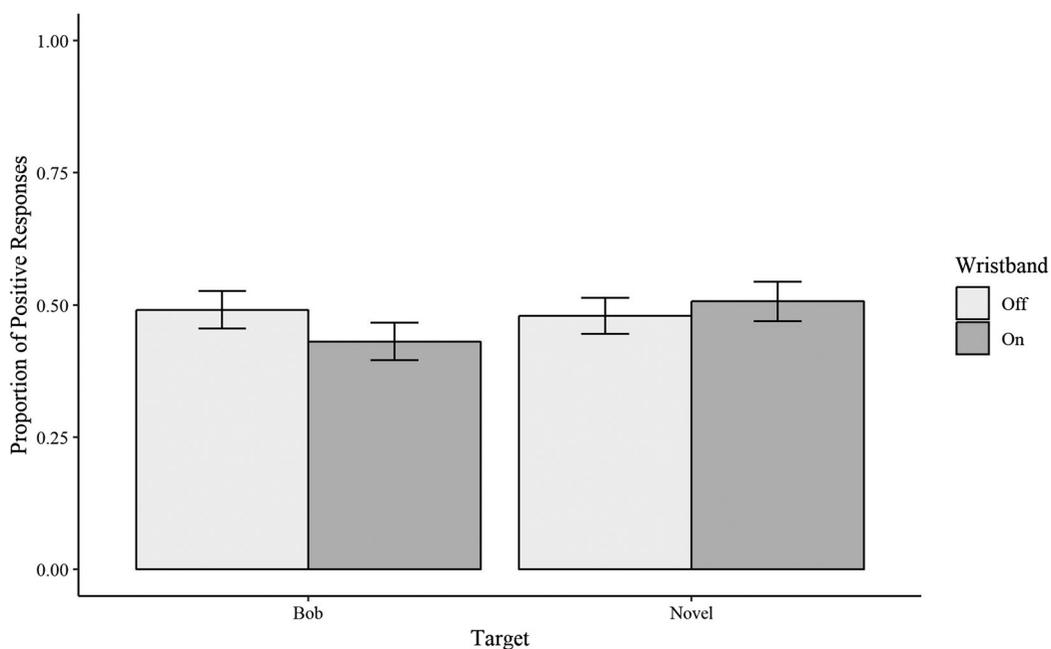
Figure 3. The influence of wristband at evaluation on speeded evaluations of Bob and novel targets in Experiment 3. On = wristband on during speeded evaluation. Off = wristband off during speeded evaluation. Proportion of Positive Responses = the proportion of "more positive" responses on the SET.

the context salience procedure did not reduce the generalisation of initial evaluations to the novel context. As in Experiment 2, the context cues influenced Bob and the novel targets in the same fashion, suggesting that the context cues became directly associated with positivity or negativity.

Experiment 4

In Experiment 3, we found that a counter-attitudinal retrieval cue influenced evaluations of Bob and novel targets on a speeded measure of evaluation, although the retrieval cue only had the predicted effect on evaluations of Bob. In Experiment 3, we did not assess evaluations of Bob and the novel targets on a more deliberate measure of evaluations. Thus, in Experiment 4, we conducted a close replication of Experiment 3 and examined the influence of the counter-attitudinal retrieval cue on both speeded and deliberate measures of evaluation.

Method

Participants and design
Our power analysis indicated that we needed 260 participants to detect the effect of the counter-attitudinal

retrieval cue with a size of d_z of 0.32 with 95% power. We were only able to collect data from 250 participants before the end of the academic term, who participated in the experiment for course credit. Five participants were eliminated for pressing the same key on every trial of the speeded evaluation task, and five participants were eliminated due to computer error. The final sample was 240 participants, 74.17% female, $M_{age} = 20.50$.

Learning procedure
The learning procedure was nearly identical to Experiment 3, with one exception. After completing the Jim learning task, all participants received the high-salience retrieval cue instructions during the second block of the Bob learning task (i.e. "Please put this band on your left wrist because Bob will behave differently when you have it on compared to when you do not").

Speeded evaluations
The SET was identical to Experiments 2 and 3, with one critical exception. Rather than manipulating the wristband within-subjects, participants were randomly assigned to either wear the wristband during the testing phase or to not wear the wristband during

the testing phase. Given the order effects on evaluations evident in Experiments 1 and 2 we sought to use a cleaner, between-participants wristband manipulation in the current experiment.

Deliberate evaluations

Participants completed feeling thermometers assessing their deliberate evaluations of Bob and the novel targets in each context. Participants reported their evaluations on a scale from 1–9. On each trial, participants evaluated Bob or one of the novel targets in the positive, negative, or novel context. The feeling thermometers were presented in a random order. To examine the influence of the wristband, participants either did or did not wear the wristband while generating deliberate evaluations. Participants who wore the wristband during the SET also wore the wristband during deliberate evaluations.

Results

SET

Rapid evaluations of Bob and the novel targets were examined in a 3 (*Context*: negative, positive, novel) × 2 (*Wristband at Evaluation*: on, off) × 2 (*Target*: Bob, novel targets) mixed ANOVA with wristband at evaluation as the only between-subjects variable. An evaluative index was computed in the same fashion as Experiment 3. As in the previous studies, we predicted that participants who wore the wristband during the testing phase would evaluate Bob more negatively than participants who did not wear the wristband during the testing phase.

As in the previous studies, there was a significant main effect of context, $F(1.77, 423.72) = 19.32$, $p < .001$, $\hat{n}_G^2 = .01$.[7] Participants evaluated Bob and the novel targets more negatively in the negative context ($M = .41$, $SD = .16$) than in the novel context

$(M = .47$, $SD = .12)$, $t(479) = 5.93$, $p < .001$, $d = 0.27$, 95% CI [0.18, 0.36], and participants evaluated Bob and the novel targets more negatively in the novel context than in the positive context ($M = .49$, $SD = .15$), $t(479) = 2.29$, $p = .03$, $d = 0.10$, 95% CI [0.01, 019]. Counter to predictions, there were no other significant main effects or interactions (see Table 1).

Deliberate evaluations

Deliberate evaluations of Bob and the novel targets were examined in a 3 (*Context*: negative, positive, novel) × 2 (*Wristband at Evaluation*: on, off) × 2 (*Target*: Bob, novel targets) mixed ANOVA with wristband at evaluation as the only between-subjects variable. In line with Experiment 1, we predicted that the wristband would not influence deliberate evaluations of Bob or the novel targets.

There was a significant main effect of context, $F(1.86, 442.05) = 23.30$, $p < .001$, $\hat{n}_G^2 = .02$.[8] Participants evaluated Bob and the novel targets more negatively in the negative context ($M = 4.30$, $SD = 2.30$) than the novel context ($M = 4.78$, $SD = 2.15$), $t(479) = 3.41$, $p < .001$, $d = 0.16$, 95% CI [0.07, 0.25], and evaluated Bob and the novel targets marginally more negatively in the novel context than the positive context ($M = 5.05$, $SD = 2.32$), $t(479) = 1.91$, $p = .06$, $d = 0.09$, 95% CI [0.002, 0.17]. Moreover, there was a significant main effect of target, $F(1, 238) = 24.13$, $p < .001$, $\hat{n}_G^2 = .04$. Participants evaluated the novel targets ($M = 5.13$, $SD = 1.98$) more positively than Bob ($M = 4.28$, $SD = 2.47$). Finally, there was a significant context × target interaction, $F(1.83, 435.62) = 4.22$, $p = .02$, $\hat{n}_G^2 = .002$. Context had a larger influence on evaluations of Bob, $F(2, 478) = 31.99$, $p < .001$, $\hat{n}_G^2 = .03$, than evaluations of the novel targets, $F(2, 478) = 6.06$, $p = .002$, $\hat{n}_G^2 = .02$.

Discussion

Experiment 4 sought to conceptually replicate the findings of the previous experiment, using a between-participants design and including a deliberate measure of evaluation. However, in Experiment 4, we found no evidence of a counter-attitudinal retrieval cue effect on speeded or deliberate measures. Given the inconsistent effects of the counter-attitudinal retrieval cue across the four experiments reported here, we conducted an internal meta-analysis to examine the effect of the retrieval cue, context cues, and target across the four studies.

Table 1. Analysis of variance table for speeded evaluations in Experiment 4.

Effect	Df_n	Df_d	F	p	$\hat{n}_G^2$
Wristband at Evaluation	1	238	0.47	.49	.0008
Context	2	476	19.11	>.001	.01
Target	1	238	2.31	.13	.004
Wristband at Evaluation × Context	2	476	0.55	.58	.0003
Wristband at Evaluation × Target	1	238	0.19	.66	.0004
Context × Target	2	476	2.03	.13	.0002
Wristband at Evaluation × Context × Target	2	476	0.26	.77	.00003

Internal meta-analysis

Experiment 1 demonstrated that a counter-attitudinal retrieval cue can influence speeded but not deliberate evaluations of Bob, but the retrieval cue had no influence on novel targets. However, Experiments 2, 3, and 4 provide relatively more mixed evidence for the nature, reliability, and generalizability of the counter-attitudinal retrieval cue effects. Consequently, we meta-analyze here the interactive effect of counter-attitudinal retrieval cue, the context cues, and target of evaluation using the procedures outlined in Mcshane and Böckenholt ([2017]). Because the designs of Experiments 2–4 were relatively similar, whereas Experiment 1 did not include context cues, all four experiments cannot be included in a single meta-analysis without collapsing data from the context conditions. Thus, for the sake of completeness, we report below two meta-analyses: one that includes data from all experiments to examine the effects of wristband and target, and another that includes data from Experiments 2–4 to examine the effects of context, wristband, and target.

Speeded evaluations

The first meta-analysis examined the influence of wristband and target on speeded measures across Experiments 1–4, collapsing across the influence of context in Experiments 2–4. We report summary statistics and the estimated within-subjects covariance matrix used to compute all the meta-analyses in the Supplementary Materials. The first meta-analysis was 2 (*Wristband*: off, on) $\times$ 2 (*Target*: Bob, novel) fully within-subjects design. We tested three contrasts: (1) the main effect of wristband, (2) the main effect of target, and (3) the interaction between wristband and target (for descriptive statistics, see Supplemental Table S1; for within-subjects covariance matrix, see Supplemental Table S2). None of the contrasts were statistically significant (*wristband*, $X^2(1) = 0.15$, $p = .696$; *target*, $X^2(1) = 1.12$, $p = .29$; *wristband* $\times$ *target*, $X^2(1) = 0.40$, $p = .53$). Moreover, we computed the heterogeneity across our studies using the I^2 statistic (Higgins & Thompson, [2002]). I^2 statistic captures the percent of variation between the studies that can be attributed to method factors. In behavioural research, an $I^2 < 25\%$ is considered low heterogeneity, an $I^2 = 50\%$ is considered medium heterogeneity, and an $I^2 > 75\%$ is considered high heterogeneity (Higgins & Thompson, [2002]; Mcshane & Böckenholt, [2017]; Pigott, [2012]). In the current studies, there was a high degree of heterogeneity across the four studies, $I^2 = 86.40\%$, 95% *CI* [78.45%, 91.42%]. This estimate suggests that approximately 86% of the variance between the current studies can be attributed to method factors.

The second meta-analysis examined the influence of wristband, context, and target in Experiments 2–4. We excluded Experiment 1 from the meta-analysis, because we did not include context cues in this experiment. The design of the meta-analysis was a 2 (*Wristband*: off, on) $\times$ 3 (*Context*: negative, novel, positive) $\times$ 3 (*Target*: Bob, novel) fully within-subjects design (for descriptive statistics, see Supplemental Table S3; for within-subjects covariance matrix, see Supplemental Table S4). The meta-analysis indicated that there was not a significant main effect of wristband, $X^2(1) = 1.29$, $p = .26$, or target, $X^2(1) < 1$.[9] To examine the influence of the context, we examined the difference between the negative context and the novel and positive contexts collapsing across all other factors. Indeed, there was a meta-analytic difference between negative context and the novel and positive contexts, $X^2(2) = 83.51$, $p < .001$. Moreover, collapsing across all other factors, there was a significant difference between the novel and positive contexts, $X^2(1) = 78.93$, $p < .001$. In summary, the meta-analysis indicated that context cues exerted a substantial influence on evaluations. However, across Experiments 2–4, the meta-analysis did not detect a reliable main effect of the wristband or the target of evaluation.

Next, we examined whether the effects of the wristband and context cues differed for Bob and the novel targets. There was not a significant meta-analytic interaction between wristband and target, $X^2(1) = 1.81$, $p = .178$. Moreover, there was not a significant influence of the target on evaluations in any of the three contexts, $X^2(3) < 1$. Overall, we do not find strong evidence that the effects of the wristband or context cues differs for Bob and novel targets.

We also calculated the heterogeneity across Experiments 2–4 using the I^2 statistic (Higgins & Thompson, [2002]). Across Experiments 2–4, there was a high degree of heterogeneity, $I^2 = 78.97$, 95% *CI* [69.54, 85.48]. This estimate suggests that approximately 79% of the variance between the studies can be attributed to method factors. Thus, the differences between the current studies have a substantial influence on the effectiveness of the manipulations.

Deliberative evaluations

Finally, we examined the meta-analytic effect of the wristband and context on deliberate evaluations of

Bob. Again for the sake of completeness, we performed two meta-analyses: one that examines the effect of the wristband on deliberate evaluations in Experiments 1, 2, and 4, and another meta-analysis that examines the interactive effect of the wristband and context in Experiments 2 and 4. We did not include the target factor because we only collected data on deliberate evaluations of the novel targets in Experiment 4.

The first meta-analysis was a 1×2 (*Wristband*: off, on) fully between-subjects design (for descriptive statistics, see Supplemental Table S5). There was not a significant meta-analytic effect of the wristband on deliberate evaluations of Bob, $X^2(1) < 1$. The second meta-analysis was a 2 (*Wristband*: off, on) $\times$ 3 (*Context*: negative, novel, positive) mixed design with wristband as a between-subjects factor and context as a within-subjects factor (for descriptive statistics, see Supplemental Table S6; for within-subjects covariance matrix, see Supplemental Table S7). There was not a significant main meta-analytic effect of the wristband, $X^2(1) < 1$. To estimate the influence of the context, we examined the difference between the negative context and the novel and positive contexts collapsing across all other factors. There was a significant meta-analytic effect between the novel and negative contexts, $X^2(2) = 27.6$, $p < .001$. Moreover, there was a marginally significant meta-analytic difference between the novel and positive contexts, $X^2(1) = 3.19$, $p = .07$. Wristband did not moderate the effect of context in any of the contexts, $X^2(1) < 1$. Finally, there was a high degree of heterogeneity in the meta-analysis, $I^2 = 82.10\%$, 95% *CI* [64.24%, 91.04%].

General Discussion

Previous research has found that initial evaluations of a person or group generalise to new contexts, whereas counter-attitudinal evaluations are limited to the context in which they are learned (i.e. become contextualised). In multiple models of human and nonhuman learning, retrieval cues (i.e. cues present when learning new, unexpected information) facilitate the generalisation of learning to novel contexts (e.g. Brooks & Bouton, 1994; Vansteenwegen et al., 2006). Theorists have argued that retrieval cue effects occur because these cues facilitate the retrieval of counter-attitudinal cue-outcome associations from memory in novel contexts (Bouton, 1993; Gawronski et al., 2018). In the present research, context cues consistently influenced evaluations, but we failed to find a reliable effect of the

counter-attitudinal retrieval cue on contextualised evaluations. Counter to our hypothesis, internal meta-analysis indicated that the counter-attitudinal retrieval cue did not disrupt the renewal or generalisation of first-learned evaluations. Moreover, we failed to find a meta-analytic effect of the counter-attitudinal retrieval cue on speeded or deliberative measures of evaluation. In the study in which the counter-attitudinal retrieval cue did shift contextualised evaluations (Experiment 3), the retrieval cue exerted an additive influence on evaluations – suggesting that the retrieval cue and context cues summed to produce different evaluations. Despite some mixed results, the overall pattern of findings suggests that counter-attitudinal retrieval cues do not attenuate the renewal and generalisation of first-learned evaluations. Although we did not find a significant influence of the counter-attitudinal retrieval cue, the current work conceptually replicated previous studies demonstrating the renewal and generalizability of first-learned evaluations (for review, see Gawronski et al., 2015). In line with past work, the internal meta-analysis found that participants evaluated the target person and novel targets more positively in the initial (positive) context relative to the second-learned (negative) context (i.e. ABA renewal: Gawronski et al., 2010). Moreover, the internal meta-analysis indicated that participants evaluated the target person and novel targets more positively in the novel context, suggesting that participants generalised their initially (positive) evaluations to a novel context (i.e. ABC renewal: Gawronski et al., 2010). However, the previous literature has found that the context cues exclusively influence the target person, presumably because the context cues set the occasion for the expression of different target person-evaluation associations. Yet, the internal meta-analysis found that the context cues influenced the target person (i.e. Bob) and the novel targets in the same fashion. Given that the context cues are a type of retrieval cue, we now turn to the question of the specificity of retrieval cues, in general, and its implications for theory-building.

Specificity of retrieval cues effects

The Representative Theory of Attitude Change argues that context cues influence evaluations by facilitating the retrieval of counter-attitudinal cue-outcome associations from memory (Gawronski et al., 2018). The theory posits that context sets the occasion for

different evaluations of the target person (for a similar argument, see Holland, 1992). Applying this theory to the current work, the context cues and the counter-attitudinal retrieval cue should only indirectly influence evaluations of Bob by facilitating the activation of Bob-positive or Bob-negative associations. Indeed, several studies have found that the effect of context cues only extends to the target person, suggesting that context serves a special modulatory role for evaluations of the target person (Gawronski et al., 2010). Moreover, recent work demonstrated that context cues will continue to modulate evaluations of the target person even after the context cues are directly associated with a different outcome (Gawronski, Ye, et al., 2014). However, as noted above, we found evidence that the context cues (i.e. the background screen colours) influenced judgments of both the target person (i.e. Bob) and novel targets.

In contrast, the current results do not provide a conclusive assessment of the specificity of the counter-attitudinal retrieval cue. In Experiment 1 and 3, the counter-attitudinal retrieval cue differently impacted Bob and novel targets, but in Experiments 2 and 4, the retrieval cue either had no influence or an indiscriminate influence on evaluations. However, the internal meta-analysis did not find a significant wristband × target interaction. In sum, we found mixed evidence for the specificity of the counter-attitudinal retrieval cue in each experiment, although the internal meta-analysis suggests that the counter-attitudinal retrieval cue may exert a similar influence on Bob and novel targets.

Given that both the counter-attitudinal retrieval cue and the context cue are types of retrieval cues, the current work provides evidence that the effects of retrieval cues, in some procedures, generalise from the target of learning to novel targets. Moreover, this pattern of results is more consistent with accounts of retrieval cues as additional discrete cues that combine with the target cue to produce different evaluations than with accounts that viewretrieval cues as cues that modulate evaluations of a target cue.

Although most work on contextualised attitude change has viewed retrieval cues as modulatory cues, there are several demonstrations in the learning literature of retrieval cues exerting a more direct influence on conditioned responding (for a review, see Urcelay & Miller, 2014). In contrast to retrieval theories, many theories of associative learning view the context as an additional, discrete cue that competes or summates with other cues (Miller & Matzel, 1988;

Rescorla & Wagner, 1972). For instance, these models argue that presentations of the unconditioned stimulus (US) alone in a conditioned chamber can produce context-US associations (i.e. US-preexposure effect: Randich & LoLordo, 1979). These context-US associations compete with discrete target cues presented later in the task. In some instances, previous context-US associations can disrupt new learning about discrete cues and lead to diminished responding to these cues. In other circumstances, the context-US associations can summate with new learning about the discrete cues to produce heightened responses to the these cues (Balaz, Capra, Hartl, & Miller, 1981).

The results of the current work closely align with the context-as-cue model of context. In the current work, we posit that the context cues and, in some circumstances, counter-attitudinal retrieval cue acquired direct associations with positivity and negativity, and the associations activated by these cues competed or summated to produce different evaluations of Bob and the novel targets. The results of Experiment 3 are particularly illuminating to this account. In this study, the counter-attitudinal retrieval cue and context cues had an additive influence on the evaluation of Bob and novel targets. Participants generated the most negative evaluations of Bob and novel targets when both the negative context cue and counter-attitudinal retrieval cue were present. In contrast, a pure-retrieval account of retrieval cues would argue that having a second retrieval cue should not have had an influence on evaluations of Bob because the Bob-negative associations are already activated by the first retrieval cue. Thus, the counter-attitudinal retrieval cue and context cues likely had an effect above and beyond the activation of specific Bob-evaluation associations.

Although the current work provides evidence for a direct mechanism by which retrieval cues influence evaluations, it is likely that retrieval cues influence evaluations through direct and indirect routes (Gawronski et al., 2018; Urcelay & Miller, 2014). Future work should focus on illuminating the conditions under which retrieval cues will directly or indirectly influence evaluations. The learning literature has already identified a handful of important features. For example, contexts are more likely to exert a direct influence when the conditioned trials are presented in rapid succession rather than presented over a longer period of time (i.e. massed versus spaced: Urcelay & Miller, 2010). Further research connecting insights from the learning literature on retrieval cues to the

literature on contextualised attitude change could produce useful insights into the nature of retrieval cues.

Automaticity of retrieval cue effects

Most dual-process models of attitudes posit that measurement conditions constrain the extent to which cognitive processes can influence responses (Gawronski, Sherman, & Trope, 2014). Though past work has examined the influence of context cues on more automatic measures of evaluation (for a review, see Gawronski et al., 2015), the current work extends this literature by also examining the influence of context cues on more deliberate evaluations. We found that context cues consistently influenced evaluations of Bob and the novel targets on speeded and deliberate measures of evaluations, suggesting that the influence of these cues occurs efficiently. Further, given that context influences both speeded and deliberate evaluations, participants likely perceive the context cues as a valid source of evaluative information to use even when they have the time and resources to do otherwise (Gawronski & Bodenhausen, 2006). Given that context influences speeded and deliberate evaluations of Bob and novel targets, it seems that the context cues are associated with general positivity and negativity and, therefore, may be applied to any target in the relevant context. Future research should further specify the mechanisms underlying context effects by exploring the conditions under which context cues will exert a general or target-specific influence on speeded and/ or deliberate evaluations. Finally, given that the counter-attitudinal retrieval cue did not influence speeded or deliberate evaluations in a consistent fashion, it is unclear whether measurement conditions constrain the influence of the cue.

Limitations

There were several limitations in the current work. First, in all four studies, participants learned positive information about the target person followed by negative information about the target person while wearing the counter-attitudinal retrieval cue. Thus, the counter-attitudinal retrieval cue was always paired with negative evaluations. The influence of the counter-attitudinal retrieval cue could change when the retrieval cue is paired with positive rather than negative evaluations. Given the power and

stickiness of negative evaluations (Ledgerwood & Boydstun, 2014; Rozin & Royzman, 2001), the counter-attitudinal retrieval cue could have less impact when it is paired with positive rather than negative evaluations, although recent work has failed to find a negativity bias in impression formation studies (Brannon & Gawronski, 2018). Future work should examine the influence of a counter-attitudinal retrieval cue paired with negative and positive evaluations.

In all four studies, we paired the retrieval cue with counter-attitudinal (or second-learned) evaluations to examine whether the retrieval cue could attenuate the stability and generalizability of first-learned evaluations. However, it is also worthwhile to investigate whether retrieval cues could operate similarly for first-learned evaluations. Indeed, given that the counter-attitudinal cue did not increase the generalizability of second-learned evaluations, future research should examine whether a retrieval cue for first-learned evaluations further bolsters the stability and generalizability of first-learned evaluations. Moreover, such research would illuminate any differences between retrieval cues for first-learned or second-learned evaluations.

Conclusion

The present research provides an initial exploration of the effect of counter-attitudinal retrieval cues on contextualised evaluations. An internal meta-analysis failed to find a reliable influence of counter-attitudinal retrieval cues on evaluations, suggesting that the counter-attitudinal retrieval cues did not disrupt the generalisation of first-learned information. Rather than competing to influence evaluations, counter-attitudinal retrieval cues and context cues may summate to produce different evaluations of a target person. Future research should continue to examine the conditions under which retrieval cues can influence evaluations.

Notes

1. There is also an extensive literature on retrieval cues in nonhuman animal models of associative learning (for a review, see Rosas et al., 2013).
2. We use the same exclusion criteria in all the studies.
3. Responses with a reaction time greater than 800 ms were eliminated in all the studies. Previous work examining speeded evaluations has applied a similar response deadline (Ranganath, Smith, & Nosek, 2008). According to this

criterion, 3.86%, 7.90%, 7.72%, and 4.29% of trials were trimmed from Experiments 1, 2, 3, and 4 respectively.

4. In Study 2 and Study 3, the background screen color contexts were counterbalanced between participants, with some participants receiving a blue positive context followed by a yellow negative context and some participants receiving a yellow positive context followed by a blue negative context. The novel context was always a white background screen color. We found no effects of color, so we report all results collapsed across this variable in Study 2 and Study 3.

5. Maulchy's Test for Sphericity indicated that the sphericity assumption was violated for this test, $W = .74$, $p < .001$. To correct for bias, the Hunyh-Feldt correction was applied.

6. Maulchy's Test for Sphericity indicated that the sphericity assumption was violated for this test, $W = .85$, $p < .001$. To correct for bias, the Hunyh-Feldt correction was applied. This correction was also applied to the Wristband at Evaluation × Context two-way interaction.

7. Maulchy's Test for Sphericity indicated that the sphericity assumption was violated for this test, $W = .89$, $p < .001$. To correct for bias, the Hunyh-Feldt correction was applied.

8. Maulchy's Test for Sphericity indicated that the sphericity assumption was violated for this test, $W = .92$, $p < .001$. To correct for bias, the Hunyh-Feldt correction was applied. This correction was also applied to the Context × Target interaction.

9. We tested the significance of each contrast using a Wald's test derived from the contrast estimate and the contrast variance-covariance matrix.

Disclosure statement

No potential conflict of interest was reported by the authors.

ORCID

Ryan J. Hutchings http://orcid.org/0000-0002-0850-1171

References

Balaz, M. A., Capra, S., Hartl, P., & Miller, R. R. (1981). Contextual potentiation of acquired behavior after devaluing direct context-US associations. *Learning and Motivation*, 12(4), 383–397. doi:10.1016/0023-9690(81)90001-1

Bouton, M. E. (1993). Context, time, and memory retrieval in the interference paradigms of pavlovian learning. *Psychological Bulletin*, 114(1), 80–99. doi:10.1037/0033-2909.114.1.80

Bouton, M. E. (2010). The multiple forms of context in associative learning. In B. Mesquita, L. Feldman-Barrett, & E. R. Smith (Eds.), *The mind in context* (pp. 233–258). New York, NY: Guilford Press.

Brannon, S. M., & Gawronski, B. (2018). In search of a negativity bias in expectancy violation. *Social Cognition*, 36(2), 199–220. doi:10.1521/soco.2018.36.2.199

Brooks, D. C., & Bouton, M. E. (1994). A retrieval cue for extinction attenuates response recovery (renewal) caused by a return to the conditioning context. *Journal of Experimental Psychology:*

Animal Behavior Processes, 20(4), 366–379. doi:10.1037/0097-7403.20.4.366

Cone, J., Mann, T. C., & Ferguson, M. J. (2017). Changing our implicit minds: How, when, and why implicit evaluations can be rapidly revised. *Advances in Experimental Social Psychology*, 56, 131–199. doi:10.1016/bs.aesp.2017.03.001

Dibbets, P., Havermans, R., & Arntz, A. (2008). All we need is a cue to remember: The effect of an extinction cue on renewal. *Behaviour Research and Therapy*, 46(9), 1070–1077. doi:10.1016/j.brat.2008.05.007

Dibbets, P., Moor, C., & Voncken, M. J. (2013). The effect of a retrieval cue on the return of spider fear. *Journal of Behavior Therapy and Experimental Psychiatry*, 44(4), 361–367. doi:10.1016/j.jbtep.2013.03.005

Faul, F., Erdfelder, E., Lang, A.-G. , & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, 39(2), 175–191.

Fazio, R. H. (1990). Multiple processes by which attitudes guide behavior: The MODE model as an integrative framework. In M. Zanna (Ed.), *Advances in Experimental social Psychology* (Vol. 23, pp. 75–109). San Diego, CA: Academic Press.

Gawronski, B., & Bodenhausen, G. V. (2006). Associative and propositional processes in evaluation: An integrative review of implicit and explicit attitude change. *Psychological Bulletin*, 132(5), 692–731. doi:10.1037/0033-2909.132.5.692

Gawronski, B., Hu, X., Rydell, R. J., Vervliet, B., & De Houwer, J. (2015). Generalization versus contextualization in automatic evaluation revisited: A meta-analysis of successful and failed replications. *Journal of Experimental Psychology: General*, 144(4), e50–e64. doi:10.1037/xge0000079

Gawronski, B., Rydell, R. J., De Houwer, J., Brannon, S. M., Ye, Y., Vervliet, B., & Hu, X. (2018). Contextualized attitude change. *Advances in Experimental Social Psychology*, 57, 1–52. doi:10.1016/bs.aesp.2017.06.001

Gawronski, B., Rydell, R. J., Vervliet, B., & De Houwer, J. (2010). Generalization versus contextualization in automatic evaluation. *Journal of Experimental Psychology: General*, 139(4), 683–701. doi:10.1037/a0020315

Gawronski, B., Sherman, J. W., & Trope, Y. (2014). Two of what? In B. Gawronski, J. W. Sherman, & Y. Trope (Eds.), *Dual-process theories of the social mind* (pp. 3–19). New York, NY: Guilford Press.

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. doi:10.1016/j.jesp.2014.05.010

Higgins, J. P. T., & Thompson, S. G. (2002). Quantifying heterogeneity in a meta-analysis. *Statistics in Medicine*, 21(11), 1539–1558. doi:10.1002/sim.1186

Holland, P. C. (1992). Occasion setting in pavlovian conditioning. In D. L. Medin (Ed.), *Psychology of learning and motivation* (Vol. 28, pp. 69–125). San Diego, CA: Academic Press.

Jones, E. E., & Harris, V. A. (1967). The attribution of attitudes. *Journal of Experimental Social Psychology*, 3(1), 1–24. doi:10.1016/0022-1031(67)90034-0

Lai, C. K., Marini, M., Lehr, S. A., Cerruti, C., Shin, J. E. L., Joy-Gaba, J. A., … Nosek, B. A. (2014). Reducing implicit racial preferences: I. A comparative investigation of 17 interventions. *Journal of Experimental Psychology: General*, 143(4), 1765–1785. doi:10.1037/a0036260

Ledgerwood, A., & Boydstun, A. E. (2014). Sticky prospects: Loss frames are cognitively stickier than gain frames. *Journal of Experimental Psychology: General*, *143*(1), 376–385. doi:10.1037/a0032310

Mcshane, B. B., & Böckenholt, U. (2017). Single paper meta-analysis: Benefits for study summary, theory-testing, and replicability. *Journal of Consumer Research*, *43*(6), 1048–1063. doi:10.1093/jcr/ucw085

Miller, R. R., & Matzel, L. D. (1988). The comparator hypothesis: A response rule for the expression of associations. In G. H. Bower (Ed.), *Psychology of learning and motivation* (Vol. 22, pp. 51–92). San Diego, CA: Academic Press.

Moors, A., & De Houwer, J. (2006). Automaticity: A theoretical and conceptual analysis. *Psychological Bulletin*, *132*(2), 297–326. doi:10.1037/0033-2909.132.2.297

Pigott, T. D. (2012). *Advances in meta-analysis*. New York, NY: Springer-Verlag.

Quezada, V. E., Laborda, M. A., Díaz, M. C., Navarro, V. M., Mallea, J., Repetto, P., … Betancourt, R. (2018). Extinction cues do not reduce recovery of extinguished conditioned fear in humans. *International Journal of Psychology and Psychological Therapy*, *18*(1), 39–53.

Randich, A., & LoLordo, V. M. (1979). Associative and nonassociative theories of the UCS preexposure phenomenon: Implications for Pavlovian conditioning. *Psychological Bulletin*, *86*(3), 523–548. doi:10.1037/0033-2909.86.3.523

Ranganath, K. A., Smith, C. T., & Nosek, B. A. (2008). From direct and indirect measurement methods. *Journal of Experimental Social Psychology*, *44*(2), 386–396.

Rescorla, R., & Wagner, A. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In A. H. Black, & W. F. Prokasy (Eds.), *Classical conditioning: Current research and theory* (Vol. 2, pp. 64–99). New York, NY: Appleton-Century-Crofts.

Roese, N. J., & Sherman, J. W. (2007). Expectancy. In A. W. Kruglanski, & E. T. Higgins (Eds.), *Social Psychology: A Handbook of Basic Principles* (2nd ed., pp. 91–115). New York, NY: Guilford Press.

Rosas, J. M., Todd, T. P., & Bouton, M. E. (2013). Context change and associative learning. *Wiley Interdisciplinary Reviews: Cognitive Science*, *4*(3), 237–244. doi:10.1002/wcs.1225

Rozin, P., & Royzman, E. B. (2001). Negativity bias, negativity dominance, and contagion. *Personality and Social Psychology Review*, *5*(4), 296–320. doi:10.1207/s15327957pspr0504_2

Strack, F., & Deutsch, R. (2004). Reflective and impulsive determinants of social behavior. *Personality and Social Psychology Review*, *8*(3), 220–247. doi:10.1207/s15327957pspr0803_1

Tulving, E., & Thompson, D. M. (1973). Encoding specificity and retrieval processes in episodic memory. *Psychological Review*, *80*(5), 352–373. doi:10.1037/h0020071

Urcelay, G. P., & Miller, R. R. (2010). Two roles of the context in Pavlovian fear conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, *36*(2), 268–280. doi:10.1037/a0017298

Urcelay, G. P., & Miller, R. R. (2014). The functions of contexts in associative learning. *Behavioural Processes*, *104*, 2–12. doi:10.1016/j.beproc.2014.02.008

Vansteenwegen, D., Vervliet, B., Hermans, D., Beckers, T., Baeyens, F., & Eelen, P. (2006). Stronger renewal in human fear conditioning when tested with an acquisition retrieval cue than with an extinction retrieval cue. *Behaviour Research and Therapy*, *44*(12), 1717–1725. doi:10.1016/j.brat.2005.10.014