

© 2022 American Psychological Association ISSN: 0022-3514

https://doi.org/10.1037/pspa0000329

How Do Explicit and Implicit Evaluations Shift? A Preregistered Meta-Analysis of the Effects of Co-Occurrence and Relational Information

Benedek Kurdi¹, Kirsten N. Morehouse², and Yarrow Dunham¹ Department of Psychology, Yale University

² Department of Psychology, Harvard University

Based on 660 effect sizes obtained from 23,255 adult participants across 51 reports of experimental studies, this meta-analysis investigates whether and when explicit (self-reported) and implicit (indirectly revealed) evaluations reflect relational information (how stimuli are related to each other) over and above cooccurrence information (the fact that stimuli have been paired with each other). Using a mixed-effects metaregression, relational information was found to dominate over contradictory co-occurrence information in shifting both explicit (mean Hedges' g = 0.97, 95% CI [0.89, 1.05], 95% PI [0.24, 1.70]) and implicit evaluations (g = 0.27, 95% CI [0.19, 0.35], 95% PI [-0.46, 1.00]). However, considerable heterogeneity in relational effects on implicit evaluation made moderator analyses necessary. Implicit evaluations were particularly sensitive to relational information (a) in between-participant (rather than within-participant) designs; when (b) co-occurrence information was held constant (rather than manipulated); (c) targets were novel (rather than known); implicit evaluations were measured (d) first (rather than last) and (e) using an affect misattribution procedure (rather than an Implicit Association Test or evaluative priming task); and (f) relational and co-occurrence information were presented in temporal proximity (rather than far apart in time). Overall, the present findings suggest that both implicit and explicit evaluations emerge from a combination of co-occurrence information and relational information, with relational information usually playing the dominant role. Critically, variability in these effects highlights a need to refocus attention from existence proof demonstrations toward theoretical and empirical work on the determinants and boundary conditions of the influences of co-occurrence and relational information on explicit and implicit evaluations.

Keywords: attitudes, associative theories, implicit evaluations, meta-analysis, propositional theories

Maintaining accurate representations of the environment is crucial for an organism's flourishing and even survival. Such representations enable humans, as well as nonhuman animals, to remember the past, anticipate the future, and perform actions that can bring about desired outcomes and prevent undesired ones. The learning mechanisms giving rise to these representations can be characterized in myriad ways. Nevertheless, one distinction that has been used consistently and fruitfully is that between learning that relies exclusively (or predominantly) on co-occurrence information and learning that, additionally, takes into account relational information (Quilty-Dunn et al., 2021). This distinction has resurfaced in slightly different forms across many seemingly unrelated areas of psychology, including in studies of animal learning (Tolman, 1949), human memory (Collins & Loftus, 1975), and cognitive development (Carey, 2011). Broadly speaking, associative theories (e.g., Dickinson, 2012; Hebb, 1949; Papineau, 2003) posit that much of cognition, and even high-level thought, can be accounted for by relatively simple representations of the fact that two stimuli (such as bread and butter, illness and medicine, or lightning and thunder) co-occur in the environment and the number of such co-occurrences. By contrast, propositional theories (e.g., Fodor, 1975; Gallistel, 1990; Mitchell et al., 2009) direct attention to more complex representations of *how* two stimuli are related to each other. For example, under associative accounts, what matters for learning is only (or primarily) the fact that lightning and thunder often go together in space and time. Propositional accounts, on the other hand, also prominently feature the relationship that two stimuli share with each other: For example, lightning may cause thunder; thunder may cause lightning; or a third variable may cause both.

Benedek Kurdi 🕩 https://orcid.org/0000-0001-5000-0584

The authors thank Patrick Mair for his generous and extremely helpful advice on statistical issues and the *Journal of Experimental Social Psychology* for a prior rejection. The authors also thank the following authors who shared effect sizes or raw data with them and thereby helped ensure that the estimates emerging from this meta-analysis are more accurate: Jeremy Cone, Jan De Houwer, Roland Deutsch, Anne Gast, Bertram Gawronski, Luke Green, Xiaoqing Hu, Sean Hughes, Yong Jae Ko, Nicolas Koranyi, Camilla Luck, Tal Moran, Steven Sweldens, and Christian Unkelbach.

Benedek Kurdi is a member of the Scientific Advisory Board at Project

Implicit, a 501(c)(3) nonprofit organization and international collaborative of researchers who are interested in implicit social cognition.

Benedek Kurdi played a lead role in conceptualization, data curation, formal analysis, investigation, methodology, project administration, visualization, and writing of original draft. Kirsten N. Morehouse played a lead role in conceptualization, data curation, investigation, methodology, project administration, and writing of review and editing. Yarrow Dunham played a lead role in conceptualization, supervision, and writing of review and editing.

Correspondence concerning this article should be addressed to Benedek Kurdi, Department of Psychology, Yale University, 2 Hillhouse Avenue, New Haven, CT 06511, United States. Email: benedek.kurdi@yale.edu Creating representations of stimulus relations is, by definition, more computationally intensive than creating representations of mere stimulus co-occurrences. Whereas the latter type of representation is traditionally thought to contain only symbols for conceptual associations (e.g., BREAD–BUTTER, ILLNESS–MEDICINE, or LIGHTNING– THUNDER) and their strengths, the former also encodes more complex and high-dimensional details of how stimuli are related to each other (Kurdi & Dunham, 2020). As such, although representations reflecting stimulus relations tend to be more computationally costly, they can be vastly beneficial: Depending on the type of relationship that two co-occurring stimuli share with each other, their hedonic consequences for the organism can be diametrically opposed to each other.

For example, a flu vaccine may come to be associated with aversive symptoms because it is causally responsible for those symptoms; alternatively, the same vaccine may come to be associated with aversive symptoms because of its power to prevent those very symptoms (Fan et al., 2021). The evaluative implications of these two cases could not be more different: A stimulus that causes an aversive outcome (such as an illness) is best avoided, whereas a stimulus that prevents that same aversive outcome is best approached even if the two co-occur equally frequently in both cases.

Likewise, in social group cognition, a social category, such as Black Americans, may come to be associated with a negative concept, such as oppression, because they are believed to be oppressive or, conversely, because they are believed to be the targets of oppression (Uhlmann et al., 2006). Again, the evaluative implications are vastly different: A social group that oppresses another social group is worthy of moral condemnation and ostracism, whereas a social group that is being oppressed by another social group is worthy of sympathy and support.

In the study of attitudes, that is, evaluations of stimuli along a positive-negative continuum (Eagly & Chaiken, 1993), learning involving co-occurrence information and relational information are customarily referred to as evaluative conditioning and persuasion, respectively.¹ In evaluative conditioning (Hofmann et al., 2010), a usually initially neutral conditioned stimulus (CS) is paired repeatedly with an intrinsically positive or intrinsically negative unconditioned stimulus (US). As a result of these repeated stimulus pairings, the CS tends to take on the valence of the US: Pairings with a positive stimulus give rise to positive evaluations and pairings with a negative stimulus give rise to negative evaluations. By contrast, in work on persuasion (Petty & Cacioppo, 1981), changes in evaluation occur in response to a verbal communication about the attitude object. For instance, attitudes toward a brand of flu vaccine may shift in a positive direction as a result of pairing the brand repeatedly with a universally admired celebrity endorser (evaluative conditioning) or as a result of a persuasive appeal featuring expert opinions on the vaccine's efficacy (persuasion).

The effectiveness of both types of learning in creating attitude change has been documented hundreds if not thousands of times (Crano & Prislin, 2006; Hofmann et al., 2010). Customarily, evaluative conditioning and persuasion are assumed to exert their effects in fundamentally different ways (Baeyens et al., 1992; Chaiken, 1980; Jones et al., 2009; Levey & Martin, 1975, 1990; Petty & Cacioppo, 1986): evaluative conditioning via automatic stimulus-driven processes leading to the formation of associations between representations of the CS and the US in long-term memory (e.g., VACCINE–NEGATIVE) and persuasion via the effortful formation of propositions about the attitude object (e.g., "I believe that this vaccine is effective"). However, more recently, this clean separation has been questioned. In particular, it has been suggested that evaluative conditioning and persuasion may both rely on similar mechanisms involving the encoding of propositions about stimulus relations (De Houwer, 2018b; De Houwer & Hughes, 2016; Mitchell et al., 2009). This theoretical development has fueled much empirical research on the relationship between these two forms of learning.

The propositional perspective has several consequences for the types of processes assumed to give rise to evaluative conditioning as well as for the conditions under which it is expected to be successful in bringing about attitude change. For example, contrary to the automatic association formation idea, propositional accounts suggest that evaluative conditioning will shift attitudes only to the degree that observers have the motivation and the ability to encode the propositions implied by the stimulus pairings to which they have been exposed. Moreover, crucially for the present purposes, the propositional perspective posits that the effects of co-occurrence information, as conveyed by repeated CS-US pairings, will be modulated by the relational meaning with which observers imbue those pairings. For example, returning to the example above, if observers believe that a vaccine repeatedly co-occurs with negative symptoms because it is causally responsible for those symptoms, they will evaluate the vaccine negatively, but if they believe that it co-occurs with negative symptoms because it can prevent those symptoms, then they will evaluate it positively.

Critically, according to several contemporary theories of evaluative learning and attitude change, whether relational information will influence learning from co-occurrence information, and to what degree, depends on a focal moderator: the type of evaluation that is being investigated (Gawronski & Bodenhausen, 2006; McConnell & Rydell, 2014; Rydell & McConnell, 2006; Smith & DeCoster, 2000; Strack & Deutsch, 2004). Specifically, these theories propose that relational information should be especially likely to modulate or even reverse the meaning of co-occurrence information in driving evaluative learning when evaluations are explicit (or measured under relatively controlled conditions using self-report). However, the likelihood of such reversal, and even modulation, is assumed to be low when evaluations are implicit (or measured under relatively suboptimal conditions, usually using some indirect behavioral index, such as response latencies; De Houwer et al., 2009; Devine, 1989; Fazio et al., 1986; Fazio & Olson, 2003; Greenwald & Banaji, 1995).

This type of view, according to which implicit evaluations should be sensitive exclusively (or at least predominantly) to co-occurrence information, can be traced back to the very beginnings of implicit social cognition research (see Kurdi & Banaji, 2022): Along with the sequential priming paradigms (Meyer & Schvaneveldt, 1971;

¹ In this article, we use the term "attitudes" to refer to evaluative representations (i.e., a latent construct) and the term "evaluations" to refer to performance on measures designed to capture attitudes (i.e., an observable behavior). We use "explicit evaluations" to refer to evaluations that are measured directly (using self-report) and "implicit evaluations" to refer to evaluations that are measured indirectly (using behavioral indices other than self-report). To avoid confusion, we refer to the measures themselves as "direct" versus "indirect" rather than "explicit" versus "implicit" (for a similar approach, see Kurdi & Dunham, 2020).

Neely, 1976) used to measure implicit evaluations in the early studies of the 1980s, the field also inherited the spreading activation models (e.g., Collins & Loftus, 1975) customarily used to interpret the results emerging from such paradigms. Specifically, it was assumed that attitudes were represented via associative strengths between conceptual nodes in long-term memory (e.g., VACCINE–NEGATIVE). Importantly, the conceptual associations underlying implicit evaluation were thought to shift, if at all, then only in response to large amounts of stimulus pairings experienced in the environment, without the flexibility afforded by high-level reasoning processes, such as persuasion (e.g., Bargh, 1999; Devine, 1989; Fiske, 1998; Wilson et al., 2000).

Importantly, the view of implicit evaluations emerging from a slow-learning, purely associative system has been challenged in several lines of early and contemporary empirical work focusing on (a) debiasing interventions (e.g., Blair, 2002); (b) the goal-dependent nature of implicit evaluation (e.g., Ferguson & Bargh, 2008; Moskowitz, 2014); and (c) the rapid revision of implicit evaluations (e.g., Cone et al., 2017; Ferguson et al., 2019). Taken together, these literatures have provided evidence for the remarkable flexibility of implicit evaluations in the face of a variety of inputs going well beyond repeated co-occurrences of stimuli in the environment. By virtue of providing a counterweight to early claims of informational encapsulation and slow associative learning, this body of work serves as an important conceptual backdrop to the present meta-analysis.²

To summarize, much research on evaluative learning and attitude change over the past decades has been driven by a distinction between (a) learning processes thought to register merely the fact that two stimuli co-occur in the environment (co-occurrence information) and (b) learning processes that are also sensitive to the way in which two stimuli are related to each other (relational information). Since the mid-2000s, there has been interest in whether relational information can modulate the effects of co-occurrence information in evaluative learning, and this interest has only intensified in recent years. Within the relevant literature, most attention has been devoted to cases in which the evaluative implications of co-occurrence information and relational information diverge, and the products of evaluative learning are being measured in an indirect rather than a direct way. As such, the present work focuses on evaluative learning under these conditions but also includes explicit evaluations as a secondary point of reference.

The Present Work

Driven by these considerations, the present meta-analysis aims to offer a comprehensive quantitative synthesis of experimental studies investigating the relative contributions of contradictory cooccurrence information and relational information to the change of implicit evaluations. To create an intuitive understanding of the types of studies included in the meta-analytic database, we provide a few illustrative examples below.

Hu et al. (2017a) exposed participants to pairings of pharmaceutical products with positive and negative outcomes (co-occurrence information). Additionally, participants were instructed either that the products caused those outcomes (assimilative relational information) or that they prevented those outcomes (contrastive relational information). The relational information in the cause condition is assimilative because it suggests that the product should be evaluated in line with the valence of the symptom with which it co-occurs. By contrast, relational information in the prevent condition is contrastive because it implies that the product should be evaluated in opposition to the valence of the symptom with which it co-occurs. Thus, if learning is sensitive to relational information, then these two conditions should yield divergent attitudes: The cause condition should give rise to positive evaluations if the outcomes are positive and negative evaluations if the outcomes are negative; the prevent condition should give rise to negative evaluations if the outcomes are positive and positive evaluations if the outcomes are negative.

In a set of experiments by Moran and Bar-Anan (2013), participants were exposed to four families of creatures, two of which co-occurred with positive sounds and two of which cooccurred with negative sounds (co-occurrence information). At the same time, relational information was also manipulated. In particular, two families of creatures caused the sounds to start (assimilative relational information), and two families of creatures caused the sounds to stop (contrastive relational information). Thus, for the crucial comparisons, the pattern of co-occurrence between creatures and sounds was identical, but the relational information was such that some creatures could be inferred to be positive (if they started positive sounds or stopped negative sounds), while others could be inferred to be negative (if they started negative sounds or stopped positive sounds).

Finally, Hughes, Ye, Van Dessel, and De Houwer (2019) had participants undergo an evaluative conditioning procedure using nonword CSs and valenced word USs (co-occurrence information). Prior to this manipulation, participants were shown context pairings consisting of either (a) identical stimuli, suggesting that the CS and the US would share the same meaning (assimilative relational information) or (b) stimuli opposite in meaning, suggesting that the CS and the US would also be opposite in meaning (contrastive relational information). Thus, if learning is sensitive to relational information, a standard evaluative conditioning effect should emerge in the same context pairings condition, whereas the evaluative conditioning effect should be attenuated or even reversed in the opposite context pairings condition. By contrast, if learning reflects only co-occurrence information, then both conditions should produce the same standard evaluative conditioning effect.

The goal of this meta-analysis is to provide a quantitatively rigorous assessment of the power of contrastive relational information in modulating the effects of co-occurrence information on attitude change, and especially the change of implicit evaluations, in cases such as the ones described above. We believe that a quantitative synthesis of this kind will be able to illuminate basic mechanisms of evaluative learning and the nature of implicit evaluations, impose constraints on both existing and future accounts

² However, the studies in this earlier literature tended not to be eligible for inclusion in this meta-analytic database because they usually did not manipulate co-occurrence and relational information about the same target. For example, debiasing studies have manipulated who the experimenter or future interaction partner was (e.g., Lowery et al., 2001; Richeson & Ambady, 2001); studies on goal sensitivity have relied on bodily states such as hunger or thirst (e.g., Ferguson & Bargh, 2004) or high-level motives such as egalitarianism (e.g., Moskowitz et al., 1999); and studies on rapid revision have provided multiple types of language-based evidence rather than a mix of co-occurrence and relational information (e.g., Cone & Ferguson, 2015; Mann & Ferguson, 2015).

of the nature of attitudes and attitude change, and inform theoretically driven interventions that aim to produce momentary or enduring changes in implicit evaluations of existing categories, especially stigmatized social groups.

When it comes to accounts of evaluative conditioning, most early theories, along with several contemporary ones, argue that co-occurrence information, to the exclusion of other sources of evidence, gives rise to this type of effect (Baeyens et al., 1992; Jones et al., 2009; Levey & Martin, 1975, 1990). To the extent that evaluative conditioning effects emerge purely from participants registering stimulus co-occurrences, these effects should be relatively informationally encapsulated and not, or at least not strongly, affected by relational information.³ By contrast, the influence of relational information would be more easily accommodated by theories of evaluative conditioning that emphasize the interplay of co-occurrence information and relational information in giving rise to stimulus evaluations (Gawronski & Bodenhausen, 2018) or those that posit that evaluative conditioning effects are mediated by the formation of propositional representations about stimulus co-occurrences (De Houwer, 2018b).

Similarly, the results of this meta-analysis will inform theories of attitude acquisition and change. As mentioned above, when it comes to implicit evaluation, several classic and contemporary accounts stress the primacy of co-occurrence information, either to the complete or near-complete exclusion of relational information (McConnell & Rydell, 2014; Rydell & McConnell, 2006; Smith & DeCoster, 2000; Strack & Deutsch, 2004), or at least under the assumption that relational information takes a back seat relative to co-occurrence information in shifting implicit evaluations (Gawronski & Bodenhausen, 2006). By contrast, more recent propositional accounts (De Houwer, 2014; De Houwer & Hughes, 2016; Kurdi & Dunham, 2020; Mandelbaum, 2016) tend to highlight the role of propositional reasoning in implicit evaluation. As such, the latter accounts, along with other approaches emphasizing the flexibility of implicit evaluation more generally (e.g., Blair, 2002; Cone et al., 2017; Ferguson et al., 2019; Moskowitz, 2014), would be considerably easier to reconcile with ubiquitous effects of relational information on implicit evaluations.

At the same time, all of these theories converge on the idea that explicit evaluations should be primarily impacted by relational information. As such, we treat explicit evaluations as a secondary dependent variable and as a reference point for gauging the strength of relational effects on implicit evaluation in the present meta-analysis. Nevertheless, this meta-analysis also gives us the opportunity to investigate the relative importance of co-occurrence information and relational information in the context of explicit, rather than implicit, evaluation. Specifically, if results were to robustly deviate from relational information fully dominating the acquisition and change of explicit evaluations, this finding may prompt fundamental revisions to widely accepted accounts of explicit social cognition.

In the context of the theoretical contribution of the present work, four points are worth mentioning. First, we do not mean to suggest that the present data can conclusively arbitrate between associative, hybrid (dual-process), and propositional theories of implicit evaluation, for multiple reasons. None of these theories have been formulated with sufficient specificity to be easily falsifiable; moreover, at least with the use of post hoc assumptions, any broad class of verbal accounts can be made generally compatible with any set of data (De Houwer et al., 2020). At the same time, it is undeniable that the theories mentioned above (as well as early approaches to evaluative conditioning and implicit social cognition) differ from each other in the extent to which they accord relative importance to relational information in evaluative conditioning and in the acquisition and change of implicit evaluations.

Notably, it would not be appropriate for us to treat associative (or hybrid) theories as interchangeable with each other given that the specific accounts belonging to these broad classes differ from each other in the way in which they envision the relationship between implicit and explicit evaluations. At one end of the spectrum, the systems of evaluation model (SEM; McConnell & Rydell, 2014; Rydell & McConnell, 2006) posits that "[...] implicit attitudes and explicit attitudes are the products of different and distinct underlying cognitive processes" (Rydell & McConnell, 2006, p. 996), with the former subserved exclusively by a slowlearning associative system and the latter exclusively by a fastlearning rule-based system. Critically, the SEM allows for little if any cross talk between these two cognitive systems and, as such, would be particularly difficult to reconcile with widespread relational influences on implicit evaluation.

At the other end of the spectrum, the associative–propositional evaluation (APE) model (Gawronski & Bodenhausen, 2006) retains the core idea that implicit evaluations primarily respond to cooccurrences in the environment and explicit evaluations additionally incorporate relational information. However, crucially, the APE model assumes that different cognitive processes (rather than distinct systems) underlie implicit and explicit evaluations and allows for interactions between these two processes. For example, according to the APE model, relational information can sometimes affect implicit evaluations indirectly, as mediated by explicit evaluations. As such, the APE model allows for the possibility of relational influences on implicit evaluations; however, relational influences should be weaker and occur less frequently in implicit than in explicit evaluation.

Finally, Smith and DeCoster (2000) and Strack and Deutsch (2004) occupy intermediate positions by assuming that the cognitive mechanisms underlying implicit and explicit evaluations are largely separate, with the former emerging from associative processes and the latter from symbolic, propositional processes. At the same time, unlike the SEM, both of these models allow for interactions between implicit and explicit evaluations under a limited set of circumstances, specifically when an initially rule-based response (such as calculating the sum of 3 and 2) becomes automatic as a result of protracted practice. Notably, none of the experiments included in the present meta-analysis involved such protracted practice—a point to which we return in the General Discussion section. As such, ubiquitous effects of relational information on implicit evaluation would also be difficult to reconcile with these two accounts.

Second, related to the previous point, in the present work, we focus not only on the overall estimate of the meta-analytic effect size

³ We note that the two predictions are asymmetrical in the sense that specific failures of relational information to influence implicit evaluations cannot be taken to imply that implicit evaluations are generally insensitive to relational information. Relational information may fail to influence implicit evaluations for a host of different reasons related to the nature of that relational information (e.g., it might come from an untrustworthy source or contain a weak argument). In line with this idea, in an exploratory analysis, we probe whether implicit and explicit evaluations tend to shift under similar conditions, and we return to this issue in the General Discussion section. We thank an anonymous reviewer for bringing this point to our attention.

but also use a wide range of moderators (including general studylevel variables, stimulus-related variables, and variables related to the learning process) to account for heterogeneity in effect sizes. We hope that these moderator analyses will provide a better understanding of the boundary conditions of relational influences on implicit (and explicit) evaluation, thereby facilitating the development of a new generation of theories. These theories should be able to make predictions not only about whether, and to what extent, relational information is expected to affect different forms of evaluation, but also about the conditions under which such effects are more or less

Third, although the results of the present work do speak to currently available theories of evaluative conditioning and implicit evaluation, its contribution does not end here. Rather, we believe that the findings emerging from this meta-analysis provide a solid base of evidence for any current or future theory of implicit evaluation. These theories are wide-ranging and include those that are not rooted in the distinction between associative and propositional processes, such as the iterative reprocessing model (Cunningham et al., 2007), the memory systems model (Amodio & Ratner, 2011), the Motivation and Opportunity as Determinants model (Fazio, 1990, 2007), or potential future accounts based on the idea of information compression (Kurdi & Dunham, 2020).

Fourth, and finally, the present data can also inform interventions designed to shift implicit evaluations of existing social categories and other attitude objects toward neutrality. We hope that this metaanalysis will help strengthen bidirectional connections between theoretically driven social cognition work and translational work focused on debiasing efforts. Stronger interconnections between the two fields may help accelerate progress toward their shared goal, namely producing momentary and long-term changes in implicit evaluations of societally relevant targets, including racial outgroups (Forscher et al., 2017; Lai et al., 2016).

Moderator Variables

likely to occur.

In addition to asking whether co-occurrence information or relational information determines implicit evaluations when the evaluative implications of the two are contradictory, we also identified several moderator variables that can be informative with respect to the generality of, and potential boundary conditions on, the overall meta-analytic effect. These moderator variables can be roughly categorized as (a) general study-level variables, (b) variables capturing the types of stimuli used, and (c) variables characterizing different aspects of the learning process. These variables are briefly described below; comprehensive information on interrater reliabilities, definitions, and levels of categorical variables is included in Table 1.

General Study-Level Variables

Publication Status, Type of Sample, Domain, and Measured Construct. The first four general study-level variables publication status of the research report (published vs. unpublished), type of sample used (nonstudent vs. student sample), domain (known vs. novel targets), and measured construct (evaluative vs. semantic)—were included in the meta-analytic database for exploratory purposes, without firm theoretical expectations about their effects. Notably, the publication status variable served as a direct test of publication bias. With regard to measured construct, although the main focus of the present work is on attitudes and not on beliefs or stereotypes due to the former being overrepresented and the latter underrepresented in the relevant primary literature, it has been shown that the attitudes and beliefs overlap with each other to a considerable degree, especially when it comes to implicit social cognition (Kurdi et al., 2019; Phills et al., 2020). In addition to including measured construct (evaluative vs. semantic) as a moderator variable, the effects of including stereotypes or beliefs in the meta-analytic models, along with several other analytic and inclusion decisions, were systematically explored in a multiverse analysis (Steegen et al., 2016; see below).

Order of Measures. The order in which explicit and implicit evaluations were measured (counterbalanced, explicit first, and implicit first) may be theoretically expected to influence the effects of relational information on implicit evaluation. Specifically, the influence of relational information on implicit evaluation could be especially strong when explicit evaluations are measured first due to carryover effects: Explicit evaluations are widely assumed to reflect relational information and, as such, reporting these evaluations first may make it more likely for any relational influence to generalize to implicit evaluations.

Measure of Implicit Evaluation. Among different measures of implicit evaluation, the Implicit Association Test (IAT; Greenwald et al., 1998) and its variants involve effortful categorization of target stimuli, whereas the evaluative priming task (EPT; Fazio et al., 1986) and its variants and the affect misattribution procedure (AMP; Payne et al., 2005) and its variants do not. As such, given that both categorization and propositional reasoning are often assumed to be relatively less automatic processes, it is conceivable that the IAT may be more responsive to relational information than indirect measures not involving stimulus categorization.

Design. Relational information may be more likely to influence implicit evaluations when it is manipulated within rather than between participants (De Houwer et al., 2020) given that the former, but not the latter, design specifically directs participants' attention to the fact that the two targets differ in the relational information characterizing them.

Stimulus-Related Variables

Type of Co-Occurrence Information. Studies have used images (Whitfield & Jordan, 2009), sounds (Moran & Bar-Anan, 2013), statements⁴ (DeCoster et al., 2006), words (Rydell et al., 2006), and other modalities including odors (Koranyi et al., 2013) to

⁴ Initially, we also included narratives (k = 4) in the meta-analytic database but removed them in response to reviewer's feedback pointing out that narratives contain rich relational content. The same reviewer also noted that statements (such as "Bob continually yells at his wife in public") also have relational content and should therefore be removed from the meta-analysis as sources of co-occurrence information. We decided to retain these cases for the following reasons: (a) designs involving behavioral statements are procedurally highly similar to evaluative conditioning designs, the only difference being spatial separation between the CS and US; (b) behavioral statements can shift implicit evaluations by virtue of their co-occurrence structure only (Kurdi & Dunham, 2021); (c) primary authors (e.g., Peters & Gawronski, 2011) often explicitly treat behavioral statements solely as a source of co-occurrence information; and (d) we did not find any difference between behavioral statements and single words as a source of co-occurrence information in moderator analyses (p = .225), suggesting that the minimal relational content of the former is unlikely to be responsible for the effects obtained.

	N.
	p
SI	õ
Ĕ	pr
IS.	-
р]	ĕ
n	at
1	Ξ.
2	Ξ
Ĕ	Se
al	IS.
\$	Ч
÷	õ
F	2
0	5
ŭ	ţ
0	2
OĽ	~
2	· 2
OL	p
÷Ē	an
13	5
8	Se
SS	n
Ž.	1
_	n
'a'	id
·Ĕ	\geq
ő	9
5	.Ħ
Ē.	O
S.	Ę
S.	4
щ	0
gu	ŝ
ü	n
E.	7
Ĕ	n
₹.	8
~	E.
ĥ	ď
-	Ð
5	ţþ
_	5
ĕ	2
þ	~
<u>ad</u> .	5
yr.	ĕ
D.	š
8	Ч
~	P.
• H	лí
nt	te
Je	ш.
H	\$
5	
9	Ë
0	Ë.
iis	ar
	~
r .	Б.
	Ē
	-

 Table 1

 Moderator Variables, Including Study-Level, Stimulus-Related, and Learning-Related Variables

6

KURDI, MOREHOUSE, AND DUNHAM

Omitted levels	Co-occurrence information Diagnosticity information Social role information	I	Counterbalanced		on interrater reliabilities (calculated
Final levels	Causal information Information on logical relations Narrative information Spatiotemporal information Validity information Other	Language Language + observation Observation	Co-occurrence first, relational immediately after Co-occurrence first, relational second Relational first, co-occurrence immediately after Relational first, co-occurrence second Simultaneous Other	Relational creates independent learning Relational modifies co-occurrence	covides definitions and information of
Explanation	What type of relational information was provided to participants?	From what source was relational information acquired?	In what order were the two types of information provided?	Does the relational information modify the co-occurrence information, or does it create independent learning?	the meta-analytic database, the table also pr
Variable name in database	typeRelational	sourceRelational	coOccRelOrder	depend .	nding variable name ir
Interrater reliability	.571 [.494, .649]	.554 [.465, .643]	.502 [.422, .583]	.810 [.744, .875]	variable and the correspc
Variable	Type of relational information	Learning-related variables Source of relational information	Order of co-occurrence and relational information	Dependency between co-occurrence and relational information	Note. In addition to the name of the

Note. In addition to the name of the variable and the corresponding variable name in the meta-analytic database, the table also provides definitions and information on interrater retrabuttes (cauculated via Gwet's AC₁) and the levels of categorical variables. For variables where the number of effect sizes with either explicit or implicit evaluations as the dependent measure was below k = 10, the variable levels were recoded before any analyses were conducted. AMP = affect misattribution procedure; EPT = evaluative priming task; IAT = Implicit Association Test.

SHIFTING EXPLICIT AND IMPLICIT EVALUATIONS

Table 1 (continued)

present co-occurrence information to participants. We did not have strong a priori expectations with regard to this variable. An overview of how co-occurrence information was manipulated in each study is available in online materials (https://osf.io/ybrwt/).

Number of Co-Occurrences. Previous meta-analytic evidence (Hofmann et al., 2010) and direct experimental evidence (Kurdi & Banaji, 2019) suggest that the effect of evaluative conditioning on implicit evaluations does not depend on the number of co-occurrences experienced by participants. However, the number of co-occurrences may still modulate relational effects on implicit evaluation. Specifically, if stimulus pairings become overlearned as a result of a large number of exposures, implicit evaluations could be less likely to additionally incorporate relational information.⁵

Duration of Co-Occurrence Information. The duration of cooccurrence information variable was included for a theoretical reason: specifically, to delineate a group of studies in which USs were presented subliminally. In these studies, the duration of stimulus presentation and the type of information presented were confounded with each other, with co-occurrence information shown for very short durations (subliminally) and relational information presented for considerably longer durations (supraliminally). As such, if these studies produce evidence in favor of the dominance of co-occurrence information, this result may be due to the dominance of subliminal information; if they produce evidence in favor of the dominance of relational information, this result may be due to the dominance of supraliminal information.

Type of Relational Information. The types of relational information relayed to participants were diverse, including information about causality, such as one group of creatures starting and the other group of creatures stopping the occurrence of a melody or a scream (Moran & Bar-Anan, 2013); diagnosticity information, such as verbal instructions informing participants that stimulus pairings either express the deep underlying character of the target groups or have been randomly generated (Kurdi & Banaji, 2019); logical relations, such as context pairings suggesting that the CS and the US are the same or opposite to each other in meaning or valence (Hughes, Ye, Van Dessel, & De Houwer, 2019); narrative information, such as a vignette describing a novel target as a heroic social worker (Mann et al., 2020); spatiotemporal information, such as CSs and USs being presented relatively close to each other or relatively far away from each other (Hughes, Mattavelli, & De Houwer, 2018); and validity information, such as attaching versus not attaching the negation operator "not" to a valenced adjective serving as a US (DeCoster et al., 2006). An overview of how relational information was manipulated in each study is available in online materials (https://osf.io/ybrwt/).

Variables Related to the Learning Process

Source of Relational Information. Although the source of relational information in the studies included in this meta-analysis was generally language, in some cases, participants were prompted to acquire relational information from direct observation alone. For example, in Study 2 of Kurdi, Morris, and Cushman (2022), some CSs were causally responsible for the appearance of USs, whereas other CSs were merely correlated with it. This difference was never verbally described to participants; rather, it had to be inferred from the physical display itself. In other studies, participants learned relational information from a combination of verbal information

and direct observation. For example, in the remaining experiments of Kurdi, Morris, and Cushman (2022), the direct observation period was preceded by extensive verbal instructions.

According to De Houwer et al. (2020), relational information is especially likely to modulate implicit evaluations when it is presented verbally, rather than nonverbally. The reason for this prediction is that in the former case, relational information is more blatant, whereas in the latter case, incorporating relational information into implicit evaluations depends on participants making the inference that the relational information should be used to modify the meaning of the stimulus pairings. Participants may not spontaneously do so or may do so to different degrees.

Order of Co-Occurrence and Relational Information. A theoretical prediction may be formulated with regard to the order of co-occurrence and relational information. Specifically, it has been suggested that relational information may exert stronger effects on implicit evaluation if it is presented closer in time to co-occurrence information (De Houwer et al., 2020; Kurdi & Dunham, 2020). The reason for this prediction is that temporal proximity may make it easier to integrate the two types of information with each other in working memory.

Dependence Between Co-Occurrence and Relational Information. In most included studies, relational information was used to modify the meaning of stimulus pairings. For example, as mentioned above, Moran and Bar-Anan (2013) paired families of novel creatures with pleasant or unpleasant sounds; when relational information was assimilative, the families started sounds of a specific valence, and when relational information was contrastive, the families stopped sounds of that valence. However, in a small number of experiments, the relational information introduced entirely new details about the target unrelated to the co-occurrence information. For example, Rydell et al. (2006) presented participants with valenced word unconditioned stimuli (co-occurrence information) and verbal narratives (relational information) about the same targets. The two were thematically fully distinct, and the relational information did not, in any way, reference the cooccurrence information.

Similar to the measured construct variable, the dependence between co-occurrence and relational information variable was included in the multiverse analysis. We did so to account for the fact that cases in which relational information created new, independent learning instead of directly modifying the meaning of stimulus pairings may be seen as less directly theoretically informative. When the two types of information are unrelated to each other, the direction of the effect may more strongly depend on the relative potency of each rather than the conceptual distinction between co-occurrence information and relational information. For example, Mann et al. (2020) exposed participants to pairings of a novel target with human scream unconditioned stimuli and then told participants that the same target was a heroic social worker. In this case, co-occurrence information and relational information were both extreme in valence. However, the outcome may have been biased in favor of relational information if the stimulus pairings had been less extreme, for example, if they had involved mildly valenced line drawings (Kurdi & Banaji, 2017).

⁵ This variable was coded in response to a reviewer comment and, as such, it is not mentioned in the preregistration.

9

Method

Open Science Practices

Open materials and data for this article, including the metaanalytic database, R code for reproducing statistical analyses, and different types of supplementary information, are available for download from the Open Science Framework (OSF; https://osf .io/42re5/). Virtually all inferential analyses were preregistered; those that were not preregistered are explicitly marked as exploratory below. The preregistration document (https://osf.io/w23rd/) and preregistered analysis script (https://osf.io/dmfy9/) are available from the OSF. The reporting in this article follows the reporting standards for research in psychology, as they apply to meta-analyses (APA Publications and Communications Board Working Group on Journal Article Reporting Standards, 2008).

Literature Search

The research reports potentially eligible for inclusion in the metaanalytic database were obtained from three sources: (a) iterative citation mining using the research reports included in a recent narrative review by Kurdi and Dunham (2020) as its starting point; (b) a reproducible online search conducted using the PsycINFO search engine (https://www.apa.org/pubs/databases/psycinfo); as well as (c) an open call sent out to the Society for Personality and Social Psychology listserv and targeted messages sent to the corresponding authors of research reports already included in the meta-analysis, asking them to contribute unpublished effect sizes. A step-by-step description (https://osf.io/ky4uw/) and flowchart (https://osf.io/89gye/) of the screening process, including the list of the PsycINFO search terms used, the full list of screened research reports along with reasons for exclusion (https://osf.io/zg2p7/), and the final list of included research reports (https://osf.io/a6mez/), are available in the online materials.

May 1, 2020 served as the cutoff date for the literature search. That is, we strove to identify and include all relevant research reports published (or completed) before this date in the metaanalytic database. A total of 3,842 potentially eligible unique research reports were screened for inclusion. Of these, 69 were found to be theoretically eligible for inclusion and 51 contained eligible effect sizes that were used in the final analyses. Theoretically eligible reports were not included in final analyses if (a) the effect sizes could not be computed (k = 2); (b) data were not available (k = 2); (c) the corresponding author did not respond to the request for data (k = 3); or (d) none of the individual effect sizes conformed to the inclusion criteria specified below (k = 7). Moreover, in response to reviewer's feedback, we removed studies from the meta-analytic database if they presented co-occurrence information embedded in narratives (k = 4), given that narratives have rich relational content.

Study-Level Inclusion and Exclusion Criteria

We sought to include in this meta-analysis all research reports that investigate the effects of contradictory co-occurrence and relational information on implicit evaluations in an experimental paradigm involving random assignment of participants to conditions. Explicit evaluations were also included in the meta-analytic database but, given the theoretical focus of the present work, the availability of explicit evaluations was not a criterion for inclusion. Each potentially eligible research report was screened independently by two coders (the first and second authors). If a coder was able to conclude that the research report was ineligible for inclusion on the basis of the title and abstract of the research report alone, then an exclusion decision was made immediately. If the title and abstract did not provide sufficient information, the coder reviewed the entire report to determine eligibility.

Coders relied on a decision tree to determine whether a research report was eligible for inclusion. Specifically, to be included, a research report was required to fulfill the following criteria, which were assessed by each coder in the same fixed order: (a) the research report had to be written in English; (b) the research report had to report at least one original empirical result; (c) the participants of the empirical study had to be adult humans above 18 years of age; (d) the empirical investigation had to use an indirect measure, including the EPT (Fazio et al., 1986), the IAT (Greenwald et al., 1998), the AMP (Payne et al., 2005), and their variations; (e) the empirical investigation had to use an experimental procedure involving random assignment to a condition designed to shift responding on one of the indirect measures mentioned above; and (f) at least one such condition had to present co-occurrence information and relational information with contradictory evaluative implications. If the research report did not meet a criterion at any level of the decision tree, it was deemed ineligible, and compliance with the remaining criteria was not investigated.

The following did not qualify as eligible indirect measures: physiological measures, including skin conductance, startle eyeblink, and electroencephalography; parameters derived from process modeling, such as the quadruple process model (Conrey et al., 2005); and indirect measures not capturing evaluative or semantic content related to a target, such as a measure of metacognitive certainty (Petty et al., 2006). Research reports that would otherwise have been eligible but used only direct but no indirect measures of evaluation (e.g., Förderer & Unkelbach, 2012) were also not included.

Disagreements regarding inclusion/exclusion decisions and the reasons for exclusion were settled by discussion. Prior to these discussions, interrater reliability between the two coders regarding inclusion/exclusion decisions, calculated using Gwet's agreement coefficient (AC₁) rather than Cohen's κ to account for large differences in base rates (Gwet, 2010), was excellent, AC₁ = .994, 95% CI [.991, .996] (percent agreement: 99.38%). In addition, interrater reliability regarding specific reasons for exclusion also reflected high levels of agreement, AC₁ = .872, 95% CI [.861, .882] (percent agreement: 87.83%).

Coding of Effect Sizes and Moderator Variables

Coding Process and Interrater Reliabilities

Effect sizes and moderator variables (see above) were independently extracted from research reports by two coders. Disagreements about both effect sizes and moderator variables were resolved via discussion. Interrater reliabilities were calculated on the basis of each coder's independent coding, finalized before any discussions took place. Each coder's independent coding, prior to any discussions, is available in online materials for effect sizes (https://osf.io/sa 4u7/) and moderator variables (https://osf.io/4v23t/). The interrater reliability for effect sizes, calculated via Gwet's AC₁ on the basis of the final Hedges' *g* effect size, was good, AC₁ = .745, 95% CI [.712, .778] (percent agreement: 74.52%). Similarly, interrater reliability for the variance of the Hedges' *g* effect size, derived from the sample size and standard error and crucial for the calculation of meta-analytic weights, was also good, AC₁ = .824, 95% CI [.795, .853] (percent agreement: 82.43%). Interrater reliabilities for moderator variables (see Table 1), also calculated via Gwet's AC₁, ranged from .502 (order of co-occurrence and relational information) to .978 (measured construct), with a median of .876. That is, all interrater reliabilities for moderate-to-excellent range.

Establishing Independent Samples

In coding effect sizes, a distinction was made between relevant and irrelevant experimental manipulations. The former was used to split study participants into independent samples, whereas the latter was not. Manipulations qualified as relevant if they either (a) differed by the content of the co-occurrence and/or relational information provided to participants (see below) or (b) corresponded to different levels of a moderator variable included in the metaanalytic database.

In some cases, the effect of a manipulation not relevant to the coding of moderator variables was significant for one dependent measure (e.g., implicit evaluations) but not the other (e.g., explicit evaluations). In these cases, the effect sizes were set to be equal across conditions on the measure for which statistics were not reported separately. However, when the manipulation was relevant to moderator variable coding, we sought to obtain effect sizes broken down by manipulation from the study authors. Conditions reflecting theoretically relevant manipulations that did not appear in a sufficient number of studies were added to the meta-analytic database as separate independent samples without a corresponding moderator variable. Such cases included the salience manipulation used in Boucher and Rydell (2012); the high-level processing instructions used in Moran et al. (2015); and the meaningful versus simple negation manipulation used in Johnson et al. (2018).

Eligibility of Effect Sizes

As shown in Table 2, crossing co-occurrence information (positive vs. negative) and relational information (assimilative vs. contrastive) with each other results in four cells: assimilative positive information (A), assimilative negative information (B), contrastive positive information (C), and contrastive negative

Table 2

Combinations of Co-Occurrence Information (Positive vs. Negative) and Relational Information (Assimilative vs. Contrastive)

Combinations	Co-occurrence information: positive	Co-occurrence information: negative
Relational information: assimilative Relational information: contrastive	A C	B D

Note. The A, B, C, D quantities could not be directly obtained for most studies; as such, contrasts involving four different combinations of these quantities were extracted.

information (D). To return to an example used earlier, the condition in Moran and Bar-Anan's (2013) study in which a family of creatures started a pleasant sound would qualify as an instance of A (assimilative positive information); the condition in which a family of creatures started an unpleasant sound would qualify as an instance of B (assimilative negative information); the condition in which a family of creatures stopped a pleasant sound would qualify as an instance of C (contrastive positive information); and the condition in which a family of creatures stopped an unpleasant sound would qualify as an instance of D (contrastive negative information).⁶

From these four combinations of co-occurrence and relational information, four comparisons of interest can be derived (see Table 3). These comparisons were coded in such a way that positive values reflect the effects of relational information over and above co-occurrence information on explicit or implicit evaluations. However, as explained in more detail in the caption of Table 3, the A–B comparison should be positive under both a co-occurrence hypothesis (according to which co-occurrence information alone determines evaluations) and under a relational hypothesis (according to which co-occurrence information are integrated with each other in driving evaluations).⁷ By contrast, predictions diverge for the remaining three comparisons: Under the relational hypothesis, D–C, A–C, and D–B should all be positive; under the co-occurrence hypothesis, D–C should be negative and A–C and D–B should be zero.⁸

In addition, we also included a smaller number of effect sizes that did not conform to the scheme outlined above. Specifically, in these studies, relational information was used not to reverse the effects of co-occurrence information but rather to denote different degrees to which the CS and the US are similar, equivalent, or related to each other. For example, Hughes, Ye, Van Dessel, and De Houwer (2019) exposed all participants to the same CS-US pairings, but participants in different conditions received different types of relational information about the CS-US relationship: Some were told that it was arbitrary, others were told that it was predictive, and a third group were told that it was causal. In studies where the manipulation involved different degrees of a positive relationship, the condition with the weakest form of relationship (e.g., arbitrary) was included using the code X, the condition with the intermediate form of relationship (e.g., predictive) using the code Y, and the condition with the strongest form of relationship (e.g., causal) using the code Z. The comparisons of interest in these studies were Y-X, Z-Y, and Z-X. Under the relational hypothesis, all these

⁶ A, B, C, and D are theoretical quantities of interest that were generally not directly available for most studies given that implicit evaluations tend to be measured in relative, rather than absolute, ways.

⁷ We use "co-occurrence hypothesis" to refer to the idea that implicit evaluations emerge from co-occurrence information alone to the exclusion of relational information and "relational hypothesis" to refer to the idea that implicit evaluations emerge from joint consideration of co-occurrence and relational information. As such, these terms are shorthands used to describe patterns of expected effects succinctly rather than full-fledged theories about the nature of implicit social cognition endorsed by any particular author.

⁸ A reviewer of this work pointed out that the D–C comparison constitutes an especially stringent test given that in this case, co-occurrence and relational information directly contradict each other, whereas for the A–C and D–B comparisons, co-occurrence information is held constant. As such, in online materials (https://osf.io/r5q4w/), we report analyses focusing on the D–C comparison only. With few exceptions, the results of moderator analyses remained unchanged.

Table 3Comparisons Extracted From Primary Studies

Comparison	Co-occurrence hypothesis	Relational hypothesis
A–B	A - B > 0	A - B > 0
D-C	D–C < 0	D-C > 0
A–C	A-C = 0	A-C > 0
D–B	D-B = 0	D-B > 0

Note. A-B is expected to be positive under both hypotheses, but for different reasons. Under the co-occurrence hypothesis, positive co-occurrence information (A) should result in more positive evaluations than negative cooccurrence information (B). Under the relational hypothesis, A-B should be positive because both A and B denote assimilative relational information, with A being positive and B being negative. The predictions of the two hypotheses differ for the remaining comparisons. Under the co-occurrence hypothesis, D-C should be negative because C carries positive information and D carries negative information. By contrast, under the relational hypothesis, D-C should be positive because the contrastive relational information should alter the meaning of the stimulus pairings. Finally, under the cooccurrence hypothesis, A-C and D-B should be zero given that they each carry the same co-occurrence information, and they differ only with respect to relational information. By contrast, under the relational hypothesis, positive information should give rise to more positive evaluations when relational information is assimilative (A) rather than contrastive (C); negative information should give rise to more positive evaluations when relational information is contrastive (D) rather than assimilative (B).

comparisons should be positive; under the co-occurrence hypothesis, they should be zero.

In identifying eligible effect sizes and sources of co-occurrence information and relational information within eligible studies, we relied on the definitions included in the introduction. Specifically, manipulations that conveyed (generally valenced) information about a target via repeated pairings of stimuli were identified as co-occurrence information, and manipulations that conveyed information about the way in which two stimuli are related to each other were identified as relational information. Applying these two definitions yielded unambiguous classifications for 39 out of 51 research reports that provided effect sizes for the metaanalytic database (corresponding to 76%; see https://osf.io/ybrwt/). In the remaining 24% of cases, we used two principles to determine eligibility: (a) inclusivity (i.e., seeking to include more, rather than fewer, effect sizes in the meta-analytic database) and (b) reliance on authors' own reasoning. These two principles led us to include in the meta-analytic database (a) statements (e.g., "Rolanda is helpful," k = 11 research reports) as a source of co-occurrence information although statements have (minimal) relational content and (b) spatiotemporal aspects of stimulus pairings (such as the physical distance between them, k = 1 research report) as a source of relational information. Given the conceptual ambiguity inherent in these decisions, the effects of including these effect sizes in the meta-analytic database were investigated in the multiverse analysis (see below).

Effect sizes were not included in the meta-analytic database if (a) two types of co-occurrence information but no relational information were presented (certain conditions of Hu et al., 2017b); (b) individual difference measures were used to assign participants to groups (Kurdi & Dunham, 2021); (c) a manipulation involved different combinations of co-occurrence and relational information such that the unique contribution of each could not be determined, for example, the same target was paired both with positive behaviors revealed to be characteristic (A) and negative behaviors revealed to be uncharacteristic (D; Bading et al., 2020; Brannon & Gawronski, 2017; Calanchini et al., 2013; Kurdi & Dunham, 2021; Rydell et al., 2007; Rydell & McConnell, 2006); (d) the same task was used for learning and testing (Kawakami et al., 2000); or (e) the study investigated generalization to a trait (Förderer & Unkelbach, 2016) or target (Hughes, Barnes-Holmes, et al., 2018) that was not included in the learning phase.

Eligible effect sizes were extracted from the research reports in line with the following hierarchy: (a) if a Cohen's d measure of effect size was reported, it was extracted directly; (b) if a t test with the corresponding degrees of freedom was reported, these two values were converted into Cohen's d; (c) if an F test with the corresponding degrees of freedom was reported, the square root of the F value (corresponding to the t value) was converted into Cohen's d; (d) if means and standard deviations were reported, the mean was divided by the standard deviation to yield Cohen's d (two-sample designs only); and (e) if means and standard errors were reported, the standard errors were converted to standard deviations and step (d) was applied. In each of these cases, the direction of the effect (i.e., the sign of the effect size) was additionally verified.

If effect sizes corresponding to A–B, D–C, A–C, or D–B were directly available, then they were immediately extracted from the research report. Otherwise, effect sizes corresponding to A, B, C, or D were extracted and the contrasts of interest calculated. In a final step, Cohen's *d* effect sizes were transformed into Hedges' *g*, which provides an unbiased estimate of the standardized mean difference. The formulas used to transform effect sizes into a common metric and to calculate the variance of the effect sizes were based on Borenstein (2009) and are available in online materials (https://osf.io/aqf93/).

To generate a conservative estimate of the effects of relational information, if a manipulation check was administered and data were reported for both the full sample and the sample that passed the manipulation check, we included effect sizes derived from the full sample. For example, Moran et al. (2017) reported results both for the full sample and after excluding participants who did not accurately recall the evaluative information that they had encountered in the learning phase. In this case, we included data from the full sample rather than the subsample. Moreover, if type of information was manipulated within participants, we included only the effect sizes that were influenced by both co-occurrence and relational information. For example, if at Time 1, participants were exposed to an evaluative conditioning manipulation but relational information was only provided at Time 2, then Time 1 data were not included in the meta-analytic database.

We generally used zero (i.e., a neutral evaluation) as a baseline to calculate effect sizes. However, if a no-intervention control condition was available (either in a within-participant design at Time 1 or in a between-participant design), then the control condition was used as a baseline. In cases where the control condition was not a no intervention control (i.e., either co-occurrence or relational information was provided to participants), the control condition was not considered when calculating the effect size. In this way, whenever possible, the effect sizes included in the meta-analysis account for baseline differences in stimulus evaluation, allowing for a more precise estimate of the meta-analytic effect. If the estimate of an otherwise eligible effect size could not be obtained directly from the research report, then the corresponding author was contacted with a request to provide the effect size or raw data that allowed for the calculation of the effect size. Corresponding authors were first contacted on May 17, 2021. If an author did not respond, reminders were sent following a 1-month delay and a 2-month delay. A total of 409 effect sizes could be directly extracted from research reports and 322 effect sizes were identified as missing. Of these missing effect sizes, we were able to obtain and include 266 effect sizes in the meta-analytic database,⁹ corresponding to a success rate of 82.61%. The remaining effect sizes (k = 56) could not be added to the meta-analytic database because they were impossible to compute (k = 26; 8.07%), the data were unavailable (k = 10; 3.11%), or the corresponding author did not respond to our request (k = 20; 6.21%).

Analytic Strategy

Statistical analyses were conducted using the metafor package (Viechtbauer, 2010) in the R statistical computing environment. Open data (https://osf.io/swkdy/) and analysis code (https://osf.io/9xcpf/) are available for download from the online materials. All inferential analyses were preregistered. Analysis code was written after the meta-analytic database was finalized but used randomly simulated values for effect sizes to keep the analytic strategy neutral to the outcome. Readers interested in exploring the data and fitting alternative models to them can do so using a freely available Shiny app (https://meta-analysis-evaluations-shift.shinyapps.io/Kurdi_e t_al_2022/).

Meta-Analytic Effect Size

The overall meta-analytic effect was calculated via stepwise model fitting using a meta-analytic mixed-effects model. Step 1 included only random intercepts for studies; Step 2 additionally included a main effect for type of evaluation (explicit vs. implicit); Step 3 additionally included a main effect for comparison (A–B, A– C, D–B, and D–C); and Step 4 additionally included a Measure × Comparison interaction. Incremental gains in model fit were determined using likelihood-ratio tests, and the best-fitting model retained and interpreted. Marginal means and prediction intervals (PIs) were calculated for each level of the retained variables (or their combinations, if applicable). In addition, the same model fitting steps were separately implemented with X, Y, and Z effect sizes as the dependent variable.

Multiverse Analysis

All meta-analyses involve a range of inherently subjective decisions about inclusion of primary effect sizes and the calculation of the meta-analytic effect size. To ascertain the robustness of the inferences made from the meta-analytic database, we conducted a multiverse analysis (Steegen et al., 2016) using the main effect sizes of interest (A–B, A–C, D–B, and D–C) with implicit and explicit evaluations each as the dependent measure. In the multiverse analysis, we relied on three main types of features to generate new estimates of the meta-analytic effect size: (a) moderator variables, (b) authorship, and (c) model fitting. **Moderator Variables.** Moderator variables were included in the multiverse analysis only to the extent that they could have served as the basis for reasonable alternative inclusion criteria. The effects of the remaining moderator variables were investigated only in meta-analytic moderator models (see below).

On each iteration of the multiverse analysis, we either included or excluded effect sizes where (a) the effect size originated from unpublished research reports, given that such reports may be seen as lower quality because they have not undergone peer review; (b) indirect measures were identified as such by authors but did not fall squarely within the inclusion criteria, such as speeded self-report and savings on relearning; (c) the construct measured was semantic rather than evaluative, given the main focus of the present work on attitude (rather than belief) acquisition and change; (d) statements were identified as sources of co-occurrence information, given that statements also have relational content; (e) co-occurrence information was presented subliminally and relational information was presented supraliminally, given the confounding between the two variables; (f) timing and location of the CS-US pairings were identified as sources of relational information, given that these variables can be seen as features of the co-occurrence information rather than as separate sources of relational information; and (g) the relational information was unrelated to the co-occurrence information, given that these studies may be seen as less directly theoretically informative (see above).

Authorship. On each iteration of the multiverse analysis, we either included or excluded effect sizes (a) originating from research reports by each unique first author and, independently, (b) originating from research reports by Tal Moran who alone contributed 33% of effect sizes to the meta-analysis.

Model Fitting. Each data set emerging from the unique combination of the features above was analyzed twice, once using the metafor package mentioned above and used for all remaining analyses and once using the robumeta package (Fisher & Tipton, 2015), which relies on slightly different analytic assumptions. Analyses were conducted by fitting intercept-only models separately for each type of dependent measure (implicit evaluations vs. explicit evaluations) and each comparison (A–B, A–C, D–B, and D–C).

As such, for each type of effect, we obtained 8,960 estimates and subsequently removed duplicate analyses for each measure– comparison pair, resulting from the fact that the data sets created were not unique due to dependencies among the features used to create them. For each measure–comparison pair, we report number of unique analyses, the mean effect size, as well as measures of variability, including the standard deviation, minimum, and maximum.

Publication Bias

Meta-analytic investigations customarily report analyses of publication bias to account for the possibility that nonsignificant effects are missing from the literature, thus skewing both the overall estimate of the effect size and moderator analyses. We generally believe that such analyses are important to conduct. Nevertheless, in the context of the present study, the outcome of any available test

⁹ Fifteen of these effect sizes were later removed because they were deemed ineligible over the course of the review process, thus resulting in a final effect size total of 660.

of publication bias should be treated with extreme caution, for multiple reasons.

First, the authors of the primary studies are known to have different theoretical commitments; as such, in a way that is highly unusual for other meta-analyses, some authors may have been motivated to suppress significant effects rather than nonsignificant effects. Second, tests of publication bias are known to result in severely biased estimates in the presence of high degrees of heterogeneity (Harrer et al., 2021), which is another feature of the present meta-analytic database. Third, most primary studies included two types of dependent measure: implicit evaluations and explicit evaluations. Because it is widely accepted that relational information should influence explicit evaluation, when it did not, authors may have assumed that the manipulation failed and thus the results may have ended up in the file drawer. Importantly, in this case, the relevant results with implicit evaluations as the dependent measure would be missing from the literature due to censoring on the direct measure and not the indirect measure itself.

Considering these and other complexities of the data-generating process, we believe that any currently available standard test of publication bias would yield ambiguous, if not fully uninterpretable, results in the context of the present data. Nonetheless, we report two tests of potential publication bias below: a direct comparison of effects extracted from published and unpublished research reports and an exploratory test of correlation between sample size and effect size (Levine et al., 2009). We chose the correlation-based method due to its relative simplicity and the fact that it theoretically allows for an unusual finding of significant (rather than nonsignificant) findings missing from the published literature. At the same time, we hope that future analysts will be able to probe the presence of publication bias in the meta-analytic database using more sophisticated methods once they become available.

Moderator Analyses

Before conducting analyses using moderator variables, the distributions of those variables were inspected, and any levels containing less than 10 observations for either explicit evaluations or implicit evaluations as the dependent measure were recoded. Table 1 provides further information about the original and recoded scale levels of moderator variables. Due to the small number of X, Y, and Z effect sizes included in the meta-analytic database, moderator models included only A–C, D–B, and D–C effect sizes. A–B effect sizes were omitted from the main moderator models given that both the co-occurrence hypothesis and the relational hypothesis predict a positive A–B effect; as such, this comparison cannot inform about the relative plausibility of the two hypotheses. However, these analyses were conducted in separate models and are reported in online materials (https://osf.io/r5q4w/).

Similar to the estimation of the overall meta-analytic effect size, the effects of moderator variables were investigated via stepwise model fitting using meta-analytic mixed-effects models. Step 1 included only random intercepts for studies; Step 2 additionally included a main effect for type of evaluation (explicit vs. implicit); Step 3 additionally included a main effect for the given moderator; and Step 4 additionally included a Measure × Moderator interaction. Incremental gains in model fit were assessed using likelihood-ratio tests and the best-fitting model retained and interpreted. To facilitate interpretation, marginal means and PIs were calculated for each level of the retained variables (or their combinations, if applicable). In addition, all pairwise comparisons were probed for statistical significance.

Results

Meta-Analytic Effect Size

Distribution of Effect Sizes

The distribution of effect sizes broken down by type of dependent measure (explicit evaluations vs. implicit evaluations) and comparison (A–B, A–C, D–B, and D–C) is shown in Figure 1.

With explicit evaluations as the dependent measure (top row), as expected, effect sizes were all positive for the A-B and A-C comparisons, virtually all positive for the D-B comparison, and mostly positive for the D-C comparison. The positive effect sizes for the A-B comparison can be seen as a manipulation check, given that for this comparison, the co-occurrence information and relational information did not differ in evaluative implications. For the remaining three comparisons, in line with expectations shared across all currently available theories of explicit social cognition, the distribution of effect sizes provides strong evidence for the dominance of relational information over co-occurrence information in the acquisition and change of explicit evaluations. Unexpectedly, the D-C effect sizes exhibited noticeably larger degrees of variability than any other type of effect size, including a considerable number of negative effects indicative of the dominance of cooccurrence information over relational information.

The pattern of implicit evaluations (bottom row) was similar, although the effect sizes were overall smaller. Specifically, for the A–B comparison, effect sizes were all positive; for the A–C and D–B comparisons, effect sizes were overwhelmingly positive; and, finally, for the D–C comparison, effect sizes were highly variable, but 60% of effects were positive and 40% were negative. As such, implicit evaluations exhibited the expected patterns of sensitivity when the evaluative implications of co-occurrence information and relational information converged with each other (A–B). When they diverged, relational information tended to dominate over co-occurrence information; although (similar to explicit evaluations), this effect was stronger for the A–C and D–B comparisons than for the D–C comparison.

Meta-Analytic Model

The best-fitting meta-analytic model contained random effects for studies, accounting for dependencies between effect sizes, and a Measure (explicit evaluations vs. implicit evaluations) \times Comparison (A–B, A–C, D–B, and D–C) interaction.

With explicit evaluations as the dependent measure, the A–B comparison significantly differed from zero and the effect size was large, Hedges' g = 1.03, 95% CI [0.95, 1.12]. The 95% PI (Borenstein et al., 2017), that is, the interval in which 95% of effects are expected to fall, indicated some measure of heterogeneity, but effect sizes were expected to be consistently positive, 95% PI [0.31, 1.76]. The marginal mean for the A–C comparison also significantly differed from zero, with a very large estimated effect, g = 1.30, 95% CI [1.22, 1.39], and no negative effects predicted to occur, 95% PI [0.57, 2.03]. A similar pattern was found for the D–B comparison, although the effect was somewhat smaller, g = 1.15,

Figure 1

Effect Size Distributions by Dependent Measure (Explicit Evaluations in the Top Row vs. Implicit Evaluations in the Bottom Row) and Comparison (A-B, A-C, D-B, and D-C)



Note. Positive values denote evaluations in line with the relational information. Horizontal lines are individual effect sizes, with the length of the lines proportional to the meta-analytic weight. The dots show the mean meta-analytic effect sizes derived from the corresponding model, and error bars are 95% confidence intervals around the mean. For better visibility, effect sizes above Hedges' g of 4.0 and below -3.0 (<2% effect sizes) were removed. See the online article for the color version of this figure.

95% CI [1.07, 1.23], 95% PI [0.42, 1.88]. Finally, we found a statistically significant but noticeably smaller (medium-sized) effect for the D–C comparison, g = 0.46, 95% CI [0.38, 0.53]. Moreover, this time, the 95% PI included negative effects, 95% PI [-0.27, 1.18], indicating the expectation that co-occurrence information may dominate over relational information at least some of the time.

With implicit evaluations as the dependent measure, the A–B comparison significantly differed from zero and the effect was of a medium size, Hedges' g = 0.55, 95% CI [0.47, 0.63]. Unlike for explicit evaluations, negative effects were expected to emerge at least some of the time, 95% PI [-0.18, 1.28]. The marginal mean for the A–C comparison also significantly differed from zero, with a small corresponding effect, g = 0.39, 95% CI [0.30, 0.48], and the expectation of some negative but mostly positive effects, 95% PI [-0.34, 1.12]. A similar pattern was found for the D–B comparison, although (similar to explicit evaluations) the effect was somewhat smaller, g = 0.29, 95% CI [0.20, 0.37], 95% PI [-0.44, 1.02].

Finally, we found a statistically significant but very small effect for the D–C comparison, g = 0.13, 95% CI [0.06, 0.21], with the expectation of a considerable number of negative effects, 95% PI [-0.59, 0.86].

Planned contrasts involved explicit versus implicit evaluations within each comparison and the three crucial comparisons (A–C, D–B, and D–C) within each type of evaluation. With respect to the first type of contrast, effects on explicit evaluations were consistently larger than on implicit evaluations, including for the A–B comparison, $\chi^2(1) = 667.78$, p < .001; the A–C comparison, $\chi^2(1) = 1013.09$, p < .001; the D–B comparison, $\chi^2(1) = 1169.43$, p < .001; and the D–C comparison, $\chi^2(1) = 507.75$, p < .001.

Within explicit evaluations, all pairwise contrasts between comparisons were significant, with A–C producing a larger effect than D–B, $\chi^2(1) = 41.87$, p < .001; A–C producing a larger effect than D–C, $\chi^2(1) = 1608.93$, p < .001; and D–B producing a larger effect than D–C, $\chi^2(1) = 1193.01$, p < .001. The pattern for implicit

evaluations as the dependent measure was the same, with A–C producing a larger effect than D–B, $\chi^2(1) = 12.14$, p < .001; A–C producing a larger effect than D–C, $\chi^2(1) = 103.09$, p < .001; and D–B producing a larger effect than D–C, $\chi^2(1) = 8.55$, p = .003.

As such, overall, we found evidence for the dominance of relational information over co-occurrence information in determining the direction of explicit and, more importantly, implicit evaluations. However, relational effects on implicit evaluations were smaller than relational effects on explicit evaluations and differed from each other in size depending on the specific comparison investigated. Nevertheless, interestingly, the pattern of differences across comparisons was the same for explicit and implicit evaluations. We return to this pattern of results in detail in the General Discussion section.

Secondary Analyses

In secondary analyses, we sought to establish whether the same pattern of results would generalize to effect sizes that we coded as X, Y, and Z, expressing differing degrees of a relationship, rather than the distinction between assimilative and contrastive relational information (as for the main effect sizes discussed above). Given the small number of Y effect sizes, only X and Z effect sizes were included in the model.

The best-fitting meta-analytic model contained random effects for studies, accounting for dependencies between effect sizes, and a Measure (explicit evaluations vs. implicit evaluations) × Comparison (X vs. Z) interaction. Within explicit evaluations, the Z comparison (strong relational information), g = 1.09, 95% CI [0.96, 1.21], 95% PI [0.68, 1.50], produced a significantly larger effect than the X comparison (same co-occurrence information but weak relational information), g = 0.84,95% CI [0.72, 0.97], 95% PI [0.43, 1.25], $\chi^2(1) = 51.82$, p < .001. The pattern of means was similar for implicit evaluations as the dependent measure, with the Z comparison, g = 0.79, 95% CI [0.67, 0.91], 95% PI [0.38, 1.20], producing a significantly larger effect than the X comparison, g =0.62, 95% CI [0.50, 0.74], 95% PI [0.21, 1.03], $\chi^2(1) = 22.55$, p < 100.001. As such, in line with the results obtained for the main comparisons above, both types of evaluation seemed sensitive to the strength of relationship communicated to participants while holding co-occurrence information constant.

Exploratory Analysis of Explicit–Implicit Convergence

Given that explicit and implicit evaluations showed similar overall patterns of sensitivity to co-occurrence and relational information, we explored whether the similarity also emerged at the level of studies, that is, whether implicit evaluations changed more in a specific sample when explicit evaluations changed. To do so, we fit a meta-analytic moderator model to a subset of studies and conditions for which both explicit and implicit measures were available (k = 266).

The level of alignment was, again, remarkable. When explicit evaluations were at neutrality, implicit evaluations were too, $\beta_0 = -0.10, 95\%$ CI [-0.20, 0.01], z = -1.76, p = .078. Moreover, the effect sizes emerging for explicit evaluations were significantly and strongly predictive of the effect sizes emerging for implicit evaluations, $\beta_1 = 0.81, 95\%$ CI [0.77, 0.86], z = 34.43, p < .001. In other

words, when explicit evaluations shifted (strongly), implicit evaluations also shifted (strongly) and vice versa.

Multiverse Analysis

As shown in Table 4, the number of unique analyses performed differed widely depending on the measure–comparison combination, ranging from 336 for the implicit A–C comparison to 3,052 for the implicit D–C comparison. The reason for this discrepancy is that moderators took on different ranges of unique values and showed different patterns of dependencies as a function of the specific measure–comparison combination.

For the A–B, A–C, and D–B comparisons, the multiverse analysis underscored the robustness of the findings that we obtained in the main meta-analytic model above, such that for both explicit and implicit evaluations as the dependent measure, mean effect sizes were consistently positive and exhibited little variability ($SDs \le$ 0.24 for explicit evaluations and $SDs \le$ 0.08 for implicit evaluations).

By contrast, for the D–C comparison, mean effect sizes showed considerable variation from iteration to iteration, ranging from g =0.06 to g = 0.90 for explicit evaluations and from g = -0.11 to g =0.09 for implicit evaluations. As such, for this comparison, metaanalytic estimates seem less robust to exclusion and analytic decisions, which may have been expected given the larger number of valid iterations alone. At the same time, mean effect sizes were positive for both comparisons. Moreover, crucially, we obtained little evidence for a significantly negative overall effect, which would have been predicted by a pure co-occurrence hypothesis.

Publication Bias

The size of the effect for the crucial A–C, D–B, and D–C comparisons did not depend on the publication status of the research report ($ps \ge .123$), thus not suggesting the presence of publication bias. With respect to A–B comparisons, effect sizes did not differ by publication status for implicit evaluations, and effect sizes were, on average, larger in unpublished than in published studies for explicit evaluations (see supplementary analyses).

Table 4

Results of the Multiverse Analysis by Dependent Measure (Explicit vs. Implicit) and Comparison (A–B, A–C, A–D, and D–C)

Type of evaluation	Comparison	Nunique	М	SD	Minimum	Maximum
Explicit	A–B	700	1.60	0.12	1.14	1.97
Explicit	A–C	488	1.42	0.24	0.85	1.95
Explicit	D–B	1,024	1.22	0.21	0.66	1.66
Explicit	D–C	2,620	0.58	0.12	0.06	0.90
Implicit	A–B	804	0.52	0.05	0.39	0.63
Implicit	A–C	336	0.34	0.07	0.24	0.56
Implicit	D–B	712	0.27	0.08	0.10	0.51
Implicit	D–C	3,052	0.01	0.04	-0.11	0.09

Note. $N_{\text{unique}} = \text{number of unique models}; M = \text{mean effect size (in Hedges' g units)}; SD = \text{standard deviation of effect size estimates across models; minimum = minimum effect size estimate; maximum = maximum effect size estimate.$

As a second measure of publication bias, we correlated sample sizes with effect sizes separately for explicit and implicit evaluations. The correlation was near zero and not statistically significant in both cases, r = -0.04, t(307) = -0.73, p = .467 for explicit evaluations, and r = -0.03, t(285) = -0.56, p = .575 for implicit evaluations. We obtained similar results when the analyses were restricted to the crucial A–C, D–B, and D–C comparisons. With the considerable caveats mentioned above, these two analyses converge on providing no evidence of systematic publication bias (favoring either null results or significant results) in the meta-analytic database.

Moderator Analyses

The results of moderator analyses are presented in detail in Table 5. The moderator analyses reported in the main text focus on comparisons where the evaluative implications of co-occurrence information and relational information diverge (A–C, D–B, and D–C). However, where informative for contrast, we also refer to supplementary analyses involving the A–B comparison.

Study Sample

Study sample did not significantly moderate the effect of relational information on implicit evaluations. By contrast, the effect of relational information on explicit evaluations was stronger among students than among nonstudents.

Domain

Domain significantly moderated the effect of relational information on implicit evaluations such that the effect was significant and positive for novel targets, whereas it was negative and not significantly different from zero for known targets. Supplementary analyses suggest that this pattern was unique to the crucial comparisons; the strength of the effect did not differ between known and novel targets for the A–B comparison. In addition, explicit evaluations exhibited the opposite pattern: The effects of relational information were significantly stronger for known than for novel targets.

Order of Measures

The order of measures significantly moderated the effect of relational information on implicit evaluations: Unexpectedly, the strongest effects of relational information were observed when implicit evaluations were measured first, the weakest effects were observed when explicit evaluations were measured first, and counterbalanced designs produced intermediate effect sizes. Notably, however, this pattern of moderation was not unique to the crucial comparisons nor to implicit evaluations: The largest effects were observed with implicit evaluations measured first for the A–B comparison (see supplementary analyses) and with explicit evaluations as the dependent measure.

Measure of Implicit Evaluation

The measure of implicit evaluation significantly moderated the effect of relational information on implicit evaluations: Unexpectedly, the mean effect was significantly different from zero only for the AMP and related measures; on the IAT, the predicted effects were more likely to be positive than negative (but not significantly so); and on the EPT, positive and negative effects were equally likely to emerge. Notably, this pattern of moderation was unique to the crucial comparisons; for the A–B comparison, the strongest effects were produced by the IAT, followed by the AMP, and then the EPT, in the order of the psychometric strength of the three measures.

Measured Construct

Measured construct did not significantly moderate the effect of relational information on implicit evaluations. By contrast, the effect of relational information on explicit evaluations was somewhat stronger for evaluative than for semantic attributes.

Design

Design significantly moderated the effect of relational information on implicit evaluations such that the effects of relational information were stronger in between-participant than in withinparticipant designs. Explicit evaluations exhibited the same pattern.¹⁰

Type of Co-Occurrence Information

Type of co-occurrence information significantly moderated the effect of relational information on implicit evaluations such that the effect was strongest for verbal types of co-occurrence information (statements and words) and weaker for nonverbal types of co-occurrence information (sounds and especially images). Explicit evaluations exhibited a somewhat different pattern: Here, statements and sounds produced the strongest effects, followed by words and images.

Number of Co-Occurrences

The number of co-occurrences did not significantly moderate the effect of relational information on implicit evaluations. By contrast, unexpectedly, the effect of relational information on explicit evaluations was significantly stronger for larger (rather than smaller) numbers of co-occurrences.

Co-Occurrence Duration

The duration of co-occurrence information did not significantly moderate the effect of relational information on implicit evaluations. By contrast, the effect of relational information on explicit evaluations was stronger when co-occurrence information was presented subliminally rather than supraliminally.

¹⁰ The variable used in this analysis relied on the type of test at retrieval. An anonymous reviewer alerted us to the fact that although the effects of relational information may be probed between participants at test, relational information may have been manipulated within participants at encoding. We were able to identify one research report with k = 12 effect sizes for implicit evaluation and k = 24 effect sizes for explicit evaluation that followed this type of procedure (Moran et al., 2017). The results remain unchanged with this research report excluded from analyses.

	÷
shers.	broadly
publi	inated
s allied	dissem
of it	be
one (ot to
01 (IS DC
ation	and
ssoci	user
cal A	idual
ologia	indiv
yche	the
P	of
ican	use
ne Amei	personal
by th	the J
ted l	for
pyrigh	solely
nent is co	intended
s docun	ticle is
This	This ar

Table 5

Moderator Variable Models for the Crucial (A-C, D-B, and D-C) Comparisons With Implicit and Explicit Evaluations as the Dependent Variable

				Impl	icit evaluations			Exp	licit evaluations	
Variable	Model	Levels	k	Estimate	95% CI	95% PI	k	Estimate	95% CI	95% PI
Publication status	Measure	Published	166				178			
		Unpublished	37		I		59			
Study sample	Interaction	Nonstudents	68 125	0.12	[-0.01, 0.26]	[-0.71, 0.96]	115	$0.68_{\rm a}$	[0.54, 0.81]	[-0.16, 1.51]
Domain	Interaction	Students Known	94 10 10	-0.13	[-0.39, 0.27]	[-0.06, 0.39]	771	0.00b	[1.58, 2.23]	[0.99, 2.82]
		Novel	169	$0.18_{\rm b}$	[0.08, 0.28]	[-0.68, 1.04]	226	$0.77_{\rm a}$	[0.67, 0.86]	[-0.09, 1.63]
Order of measures	Main effects	Counterbalanced	76	$0.13_{ m b}$	[-0.02, 0.27]	[-0.76, 1.01]	127	$0.74_{ m b}$	[0.60, 0.89]	[-0.14, 1.63]
		Explicit first	35	-0.11_{a}	[-0.26, 0.04]	[-0.99, 0.77]	49	$0.51_{\rm a}$	[0.36, 0.65]	[-0.38, 1.39]
		Implicit first	92	$0.27_{\rm b}$	[0.15, 0.39]	[-0.61, 1.15]	61	0.88_{c}	[0.76, 1.01]	[0.00, 1.77]
Measure of implicit	Moderator	AMP family	46	$0.32_{\rm b}$	[0.21, 0.42]	[-0.21, 0.84]				
evaluation		EPT family	70	$-0.01_{\rm a}$	[-0.11, 0.10]	[-0.53, 0.52]				
		IAT family	79	$0.07_{\rm a}$	[-0.01, 0.15]	[-0.45, 0.59]				
-		Other	×,	$0.46_{\rm b}$	[0.15, 0.78]	[-0.14, 1.06]	2			
Measured construct	Interaction	Evaluative	191	0.16	[0.07, 0.25]	[-0.66, 0.98]	177	0.77b	[0.68, 0.86]	[-0.05, 1.59]
-		Semantic	12	0.25	[-0.06, 0.56]	[-0.62, 1.12]	23	0.50 _a	[0.19, 0.81]	[-0.37, 1.37]
Design	Interaction	Between	67	0.4/b	[U.35, U.6U]	$\begin{bmatrix} -0.44, 1.39 \end{bmatrix}$	66	1.83b	[1. /2, 1.95]	[0.92, 2. /4]
E		W IUNIN	1/4	0.03ª	[-0.04, 0.15]	[06.0, 0.80]	7/1	0.51a	[0.42, 0.01]	[-0.39, 1.42]
1 ype of co-occurrence	Interaction	Images	70	70'0-	[-0.17, 0.13]	[C/.0, 6/.0–]	çç Ç	. 9C.0	[0.43, 0.73] [0.11_0.001	$\begin{bmatrix} -0.19, 1.30 \end{bmatrix}$
IIIIOIIIIauoii			215	0.12%	[-0.0, 0.07]	[-0.50, 1.11]	7 00	40.0 *10.0	[-0.11, 0.09]	[UC.1,2C.U-]
		Sounds	17	0.10	[10.00, 0.0.]	[06.0, 60.0-]	60 E	0.07*	[0.00, 1.09]	[0.08, 1.00]
		Statements Wreads	4 0 4 0	0.33	[0.17, 0.48]	[-0.45, 1.10]	71	0.95	[0.77, 1.08]	[0.15, 1.70]
Number of an anomatic	Internation		00	0.15	[0.02, 0.07]	[06.0, 6C.0-]	1	00.0	[0.40, 0.60]	[-0.13, 1.41]
Number of co-occurrences	Interaction	+1 50		0.10	[170, 000]	[-0.61, 0.94] [0.64 0.01]		0.996	[0.8/, 1.11]	$\begin{bmatrix} 0.21, 1.70 \end{bmatrix}$
			`	0.14	[0.02, 0.20]	[-0.04, 0.91]	°	0.02_{a}	[0.20, 0.74]	[140] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102] [102]
Co-occurrence duration	Interaction	Subliminal	× v	0.28	[-0.10, 0.66]	[-0.60, 1.16]	× 000	2.12a	[1.72, 2.53]	[1.23, 3.01]
		Supranminal	C61	0.14	[22.0, c0.0]	[0.00, 0.95]	677	0.75b	[0.04, 0.81]	[70.1, /0.0–]
type of relational information	Interaction	Causal	5 2 2 2	0.13	[-0.05, 0.51]	[16.0, 0.01]	50 21	0.850	[0.6/, 1.03]	[0.07, 1.03]
			51 5	0.42	[0.14, 0.70]	[-0.59, 1.23]	01 8	0.43_{a}	[1/.0, c1.0]	[-0.58, 1.24]
		Narrative	57	0.31	[0.14, 0.47]	[-0.47, 1.08]	57	1.14 d	[0.97, 1.30]	[0.36, 1.91]
		Other	31	-0.01*	[-0.25, 0.22]	[-0.81, 0.78]	28	0.22_{a}	[-0.02, 0.46]	[-0.57, 1.02]
		Validity	98	0.15^{*}	[0.04, 0.26]	[-0.61, 0.92]	107	$0.73_{ m b}$	[0.62, 0.85]	[-0.03, 1.50]
Source of relational information	Interaction	Language	163	0.16	[0.06, 0.26]	[-0.68, 0.99]	181	$0.78_{ m b}$	[0.68, 0.88]	[-0.06, 1.61]
		Language + observation	22	0.14	[-0.08, 0.36]	[-0.72, 1.00]	33	0.94_{c}	[0.72, 1.15]	[0.08, 1.79]
		Observation	18	0.23	[0.01, 0.46]	[-0.62, 1.09]	23	$0.34_{ m a}$	[0.12, 0.57]	[-0.52, 1.20]
Order of co-occurrence (CO) and	Interaction	CO first, REL	54	0.18^{*}	[0.04, 0.31]	[-0.60, 0.95]	53	1.10 d	[0.96, 1.23]	[0.32, 1.87]
relational (REL) information		immediately after								
		CO first, REL second	39	0.10^{*}	[-0.02, 0.22]	[-0.67, 0.88]	50	$0.53_{ m a}$	[0.41, 0.65]	[-0.25, 1.31]
		Other	4	0.24^{*}	[-0.41, 0.89]	[-0.76, 1.25]	4	1.19^{*}	[0.51, 1.88]	[0.17, 2.22]
		REL first, CO	10	0.20^{*}	[-0.09, 0.50]	[-0.62, 1.03]	16	1.04_{c}	[0.74, 1.34]	[0.22, 1.86]
		immediately after								
		REL first, CO second	23	0.14^{*}	[0.01, 0.28]	[-0.63, 0.92]	38	$0.55_{ m a}$	[0.41, 0.68]	[-0.23, 1.33]
		Simultaneous	73	0.15^{*}	[0.00, 0.30]	[-0.63, 0.93]	76	$0.83_{\rm h}$	[0.69, 0.98]	[0.05, 1.61]
Dependence between co-occurrence	Interaction	Independent	24	$0.05_{ m a}$	[-0.07, 0.17]	[-0.78, 0.88]	24	$1.00_{\rm h}$	[0.88, 1.13]	[0.17, 1.84]
and relational information		Modifies	179	$0.17_{\rm b}$	[0.08, 0.26]	[-0.66, 1.00]	213	$0.74_{\rm a}$	[0.65, 0.82]	[-0.09, 1.57]
<i>Note</i> . Variable = name of moderator v	variable; model = $ $	best-fitting model (measure = 1	nain effect	of measure,	moderator = mai	n effect of moders	ator, main	effects = tw	o main effects, a	nd interaction =
Measure X Moderator Interaction, leve	els = levels of ca	tegorical moderator variable;	$k = num_0$	er of effect s	sizes: esumate =	marginal mean e	stimated	from the bes	st-htting model:	0.000 CI = 30.000 CI

SHIFTING EXPLICIT AND IMPLICIT EVALUATIONS

17

Type of Relational Information

Type of relational information significantly moderated the effect of relational information on implicit evaluations such that the effect was strongest for logical and narrative information and weaker for causal and validity information. Explicit evaluations exhibited a somewhat different pattern: Here, narratives produced the strongest effect, followed by casual information, validity information, and, finally, logical information.

Source of Relational Information

Source of relational information did not significantly moderate the effect of relational information on implicit evaluations. By contrast, the effect of relational information on explicit evaluations was strongest when relational information was acquired from a combination of language and observation, followed by language, and finally by observation alone.

Order of Co-Occurrence and Relational Information

Order of co-occurrence and relational information significantly moderated the effect of relational information on implicit evaluations such that the effect was strongest when relational information was presented first and co-occurrence information was presented immediately after and weakest when co-occurrence information was presented first and relational information was presented second. The remaining orders produced effects of intermediate size. Explicit evaluations exhibited a similar pattern but with cleaner separation such that the strongest effects occurred when the two types of information were presented immediately after each other, followed by simultaneous presentation, and finally by presentation of the two types of information far apart in time.

Dependence Between Co-Occurrence and Relational Information

Dependence between co-occurrence and relational information significantly moderated the effect of relational information on implicit evaluations such that the effect was strongest when relational information modified the meaning of co-occurrence information and considerably weaker when the two were unrelated. By contrast, the effect of relational information on explicit evaluations was stronger when the two types of information were unrelated to each other rather than related.

General Discussion

In the present meta-analysis, we have provided a quantitative synthesis of over 600 effect sizes extracted from 51 research reports to investigate how implicit evaluations shift under conditions where co-occurrence information about an attitude object (e.g., pairings of a vaccine brand with unpleasant symptoms) and relational information about the same attitude object (e.g., the vaccine brand prevents those symptoms) have contradictory evaluative implications. The most important finding emerging from this work is that, much like with explicit (self-reported) evaluations, the dominance of relational information over co-occurrence information is the rule rather than the exception in the acquisition and change of implicit (indirectly revealed) evaluations. These findings

are fundamentally incompatible with the idea that implicit evaluations are the product of an informationally encapsulated cognitive system that reflects only the effects of co-occurrence information, to the full exclusion of relational information. In addition, they suggest that evaluative conditioning effects tend to be subserved by cognitive processes that are penetrable by high-level reasoning involving stimulus relations.

Theoretical Implications

These results may be easiest to reconcile with those accounts that posit that implicit evaluations, similar to their explicit counterparts, are fundamentally propositional in nature (De Houwer, 2014; De Houwer & Hughes, 2016; Kurdi & Dunham, 2020; Mandelbaum, 2016). By contrast, they seem considerably less compatible with those accounts according to which implicit evaluations are thought to originate primarily, if not exclusively, from associative representations (Gawronski & Bodenhausen, 2006; McConnell & Rydell, 2014; Rydell & McConnell, 2006; Smith & DeCoster, 2000; Strack & Deutsch, 2004). After all, in the experiments included in this meta-analysis, participants were exposed to both co-occurrence information and relational information about the same attitude object under conditions where the evaluative implications of the two were clearly in conflict.

Accordingly, the main finding of this meta-analysis seems compatible with the conclusions of a recent narrative review with a similar focus by Kurdi and Dunham (2020). That narrative review concluded that "relational information can have a stronger influence on implicit evaluations than conflicting co-occurrence information" (p. S56) and that "such a pattern [...], although subject to certain boundary conditions, is the rule rather than the exception" (p. S57). In addition, the present findings are also easily reconciled with recent proposals suggesting that propositional processes give rise to evaluative conditioning effects (De Houwer, 2018b).

As mentioned in the introduction, associative and hybrid (dualprocess) accounts of implicit evaluation differ from each other considerably in terms of the degree to which they allow for interactions between implicit and explicit evaluations and, in a similar vein, for relational influences on implicit evaluation. As such, the present results seem most directly incompatible with the SEM (McConnell & Rydell, 2014; Rydell & McConnell, 2006), which posits that implicit and explicit evaluations emerge from two separate (informationally encapsulated) cognitive systems. By contrast, they are easier to reconcile with more permissive theories, such as the APE model (Gawronski & Bodenhausen, 2006), which assume that implicit and explicit evaluations should both emerge from the interplay of learning based on co-occurrences and learning based on more complex relational content.

Theories by Smith and DeCoster (2000) and Strack and Deutsch (2004) posit that implicit evaluations may be influenced by relational information as a result of protracted practice (such as when the sum 3 + 2 is not calculated in a rule-based manner anymore but simply pops into a person's consciousness). However, the present studies involved relatively short learning phases, thus making it unlikely that such overlearning would emerge. As such, it is not easy to see how theories by Smith and DeCoster (2000) and Strack and Deutsch (2004) would account for the present results. In addition, evidence for the rapid revision of implicit evaluations has been provided in a considerable number of studies

(for reviews, see Cone et al., 2017; Ferguson et al., 2019). These experiments, although not eligible for inclusion in the present metaanalysis because they rely purely on verbal narratives and do not use co-occurrence information as a source of evaluative learning, cast further doubt on the accuracy of purely (or predominantly) associative approaches to implicit social cognition.

Notably, to avoid straying too far from experimental procedures and encumbering the reporting of the meta-analytic results with excessive theoretical precommitments, in the main body of this article, we refrained from making any statements about process or mechanism.¹¹ Instead, we focused on the sensitivity of implicit (and explicit) evaluations to two types of input: co-occurrence information and relational information. However, and critically, all relevant theories make claims not only about different sources of information but also about different types of learning processes that unfold in response to these sources of information. As such, it is worth asking what if anything can be concluded from the present evidence with respect to the learning processes giving rise to implicit and explicit evaluations, specifically with respect to the distinction between associative and propositional mechanisms.

We acknowledge that sensitivity to co-occurrence information need not be indicative of the operation of associative processes. For example, co-occurrence information may influence implicit evaluations via the formation of propositional representations about stimulus relations (e.g., "medicine co-occurs with negative symptoms"; e.g., De Houwer, 2014, 2018b). And, conversely, relational information could influence implicit evaluations via associative processes, such as when associative strengths in long-term memory (e.g., MEDICINE–BAD) shift in response to propositional reasoning (e.g., Gawronski & Bodenhausen, 2006) or are assigned validity tags (Petty et al., 2006).¹²

However, although no one-to-one correspondence between sensitivity to input and cognitive processes can be assumed to exist, sensitivity to input is often the primary source of evidence supporting inferences about cognitive process, and this is especially the case under the conditions of interest to the present meta-analysis. Specifically, we believe that it is reasonable to conclude from the present data that propositional processes can (and often do) influence implicit evaluations. The reason for this is that although co-occurrence information can serve as an input to propositional processes, propositional processes, by definition, can and do also incorporate other sources of information, including relational information (e.g., Gawronski & Bodenhausen, 2006; Mitchell et al., 2009; Quilty-Dunn et al., 2021). As such, although it is conceivable to defend a propositional theory under which implicit evaluations emerge from propositional processes but only propositional processes operating over co-occurrence information, we are not aware of any such currently available account. In addition, such accounts would go against the very definition of the nature and scope of propositional reasoning.

Conversely, it is also possible to argue that conceptual associations are embedded in, and modulated by, high-level processes of reasoning. However, even the most permissive dual-process theory of implicit evaluation of which we are aware (the APE model; Gawronski & Bodenhausen, 2006) posits that, at their core, implicit evaluations are associative, even if they can sometimes be affected by propositional processes.¹³ Of course, a theory positing that associative representations are consistently embedded in, and are modulated by, high-level reasoning may be proposed and defended. Under such a theory, implicit evaluations of a target would still be proximally mediated by conceptual associations (such as MEDICINE– BAD) but these conceptual associations would be fully amenable to forming and changing via propositional processes (such as reasoning about the symptoms that the medicine may be able to prevent). However, we are skeptical as to whether and how an account of this kind could be empirically distinguished from already existing propositional accounts (De Houwer, 2018a).

Variability in Relational Influences on Implicit and Explicit Evaluations Due to Type of Comparison

Although the mean meta-analytic effect size is most naturally compatible with a propositional perspective, additional analyses also considering the type of comparison being investigated paint a significantly more complex overall picture. Taken together, these analyses suggest that relational information alone is not sufficient to account for patterns of change in implicit evaluation. Rather, implicit evaluations seem to emerge from the interplay of relational information and co-occurrence information.

The influence of co-occurrence information, over and above relational information, is visible in the fact that (a) relational influences were stronger when co-occurrence information was held constant (A–C and D–B comparisons) rather than manipulated (D–C comparison); (b) even when co-occurrence information was held constant, relational influences were more apparent in the positive (A–C) than in the negative domain (D–B); and (c) in the cases involving both co-occurrence information and contradictory relational information, the contrastive case (D–C) was not simply a mirror image of the assimilative case (A–B) but rather produced a considerably smaller effect. These results were consistently reflected by (a) differences in mean effect size estimates, (b) amounts of associated heterogeneity, and (c) robustness of the effect size estimate across different iterations of the multiverse analysis.

Additionally, and remarkably, explicit evaluations exhibited the very same pattern of comparison differences that implicit evaluations did: Relational information dominated most clearly when the valence of co-occurrence information was held constant across the two targets; when co-occurrence information was held constant, relational effects were larger in the positive than in the negative domain; and when the valence associated with the two targets differed, the difference between them was more than twice as large when relational information was assimilative rather than contrastive. Moreover, an exploratory analysis indicated that explicit and implicit evaluations tended to move in tandem: When one changed,

¹¹ We thank an anonymous reviewer for prompting us to do so and thereby to preserve conceptual clarity across the article.

¹² Along similar lines, it is conceivable that implicit evaluation is sensitive to propositional processes during learning, but that these propositional processes create associative representations, which are then activated via associative processes. We thank an anonymous reviewer for raising this issue.

¹³ For example, Peters and Gawronski (2011) hypothesized an implicit– explicit dissociation in the canonical situation investigated in the present meta-analysis on the basis of the APE model: "Based on dual-process theories that propose two parallel learning mechanisms (e.g., Gawronski & Bodenhausen, 2006; Rydell & McConnell, 2006; Strack & Deutsch, 2004), we expected an evaluative dissociation in this situation because of the simultaneous operation of belief-based and contingency-based learning processes" (p. 558).

This article is intended solely for the personal use of the individual user and is not to be disseminated broadly

This document is copyrighted by the American Psychological Association or one of its allied publishers.

the other did too; and when one did not, the other did not either. As such, these results provide clear evidence for the idea that, much like their implicit counterparts, explicit evaluations are not immune to the effects of co-occurrence information but rather emerge from a joint consideration of relational information and co-occurrence information even when the two are directly in conflict with each other.

We believe that the result that explicit evaluations emerge from a mix of co-occurrence and relational information is among the most important (if unexpected) findings of the present meta-analysis and, as such, will require considerable attention in future theoretical and empirical work. In fact, several investigations outside the field of attitude research have suggested that explicit judgments, which most contemporary theories of attitude acquisition and change assume should be fully propositional, are often influenced by associative processes (e.g., Rehder, 2009, 2015). Importantly, individual investigations in the attitude domain have also yielded similar conclusions (e.g., Moran et al., 2016). The present results seem to be broadly in line with these perspectives, with the above-mentioned caveat that the influence of co-occurrence information need not always map onto associative processes and the influence of relational information need not always map onto propositional processes.

How might one account for this type of result in terms of currently available theories of implicit (and explicit) evaluation? At a first glance, dual-process theories, such as the APE model (Gawronski & Bodenhausen, 2006, 2018), under which implicit evaluation is thought to be a function of the joint operation of associative and propositional processes, may appear to be best equipped to account for the present findings. However, these theories prominently predict a difference between implicit and explicit evaluations in their relative sensitivity to co-occurrence and relational information, with the former thought to be primarily responsive to the former and the latter thought to be primarily responsive to the latter. And yet, in the present meta-analysis, patterns of change were highly similar across the two types of evaluation (other than a main effect of measure type, to which we return below).

That said, dual-process (and associative) theories may be modified to account for the totality of these results. For example, it is conceivable that, as suggested by several accounts of this kind, co-occurrence information may give rise to low-level, stimulusdriven associative processes (Gawronski & Bodenhausen, 2006; McConnell & Rydell, 2014; Rydell & McConnell, 2006; Smith & DeCoster, 2000; Strack & Deutsch, 2004) but that these processes may influence not only implicit evaluations but also their explicit counterparts. These influences may be direct, or may unfold indirectly, as mediated by implicit evaluations (Gawronski & Bodenhausen, 2006; Greenwald & Banaji, 2017). As such, this type of modified account would retain a relatively strict separation between associative and propositional processes but would do away with a one-to-one mapping between associative learning and implicit evaluation and propositional learning and explicit evaluation.

Alternatively, the effects of co-occurrence information on both explicit and implicit evaluations may be mediated by propositional processes, including propositional processes invoked by the experimental context. Specifically, participants in an experiment may treat co-occurrence information and relational information in similar ways, as forming part of the same persuasion attempt (De Houwer & Hughes, 2016). As such, they may reason about the

experimenter's intent in the context of this persuasion attempt and make inferences about why the experimenter chose to say, "It is not true that Sue Ellen is mean" instead of choosing the more straightforward phrasing "Sue Ellen is nice." Is the experimenter trying to create uncertainty? Are they implying that someone else thought that Sue Ellen was mean? We believe that devoting more attention to the pragmatics of the communicative situation between the experimenter and the participant (Grice, 1975) and treating experiments as pedagogical situations (Csibra & Gergely, 2009) may go a long way toward understanding effects of this kind.

Moderators Accounting for Variability in Relational Influences on Implicit Evaluation

Additional complexity in the meta-analytic findings was created by the results of moderator analyses. As we detail below, current theories of implicit evaluation are not particularly well-equipped to account for these effects. As such, we hope that the correlational findings involving moderator variables, to which we now turn, will spur not only attempts at experimental validation but also attempts at new theory development.

In terms of the results that they produced, the moderator variables probed in this meta-analysis can be roughly categorized into three sets: (a) variables that did not moderate the effects of relational information on implicit evaluation; (b) variables that moderated the effects of relational information on both implicit and explicit evaluations in similar ways; and (c) variables that uniquely moderated the effects of relational information on implicit (but not explicit) evaluation.

Notably, the effects of relational information on implicit evaluation were not modulated by six of the 14 moderators included: publication status (published vs. unpublished), study sample (nonstudent vs. student), measured construct (evaluation vs. belief), number of co-occurrences, duration of co-occurrence information (subliminal vs. supraliminal), and the source of relational information (language, language + observation, and observation; for a competing prediction, see De Houwer et al., 2020). As such, these findings provide evidence for the robustness of this type of effect.

The effects of relational information on both implicit and explicit evaluations were similarly affected by the order of the two measures and the order of co-occurrence and relational information. Specifically, the effects of relational information were strongest when implicit evaluations were measured first, weaker when the two measures were counterbalanced, and weakest when explicit evaluations were measured first. This is an effect that we did not expect a priori; in fact, we speculated that completing the self-report measure before the indirect measure might have produced carryover effects such that the relational information retrieved on the former would influence responding on the latter.

A potential post hoc explanation for this effect is that relational information is more difficult to encode, maintain in memory, and retrieve than co-occurrence information is. As such, its effect on implicit evaluation may be stronger when the evaluation is retrieved closer in time to the original learning (without an intervening measure of explicit evaluation). By contrast, because explicit evaluation allows more effortful processing, the effect of relational information on explicit evaluation may be less affected by the order of the measures.¹⁴

¹⁴ We thank an anonymous reviewer for alerting us to this possibility.

Be that as it may, we believe that this pattern of results will require experimental confirmation and, if confirmed, a theoretical explanation.

The effects of the temporal proximity of co-occurrence and relational information on implicit evaluation are more easily explained theoretically (Kurdi & Dunham, 2020): Given processing limits and constraints on long-term memory, the two types of information seem easiest to integrate with each other when they are presented close to each other in time rather than far apart but not simultaneously. Interestingly, these limitations appear to be operating not only in the context of implicit evaluations but even in the context of explicit evaluations, suggesting another way in which the processes giving rise to each may overlap with each other. This result is also roughly in line with predictions about temporal proximity by De Houwer et al. (2020), although these authors predicted that the largest effect should emerge when co-occurrence and relational information are presented at the same time. In fact, it seems that the working memory demands of simultaneous presentation may interfere with learning relative to a design in which the two types of information are presented close to each other in time but not concurrently.

Finally, we found that relational information had a larger effect on implicit (and explicit) evaluations when it was manipulated between, rather than within, participants. This finding is in contradiction with a prediction by De Houwer et al. (2020) who reasoned that relational information should be more impactful in withinparticipant designs because in such designs, participants' attention is focused directly on variation in relational information, whereas in between-participant designs, it is not. Although we did not expect to find the opposite effect, like the effect of temporal order of presentation, it can be attributed to differences in working memory demands. Namely, within-participant but not between-participant designs require participants to hold multiple (contradictory) pieces of relational information simultaneously in mind and to try to commit them into long-term memory.

The effects of relational information on implicit evaluation were uniquely moderated by the domain of the study (known vs. novel targets), the indirect measure (AMP, EPT, and IAT), as well as the type of co-occurrence and relational information. Specifically, implicit evaluations were more likely to be affected by relational information when co-occurrence information was presented verbally (as statements or words) rather than nonverbally (as sounds or images). This effect, although unexpected, is not difficult to explain theoretically given that propositional representations are widely assumed to be language-based; as such, the match between the type of stimulus presentation and type of representation may give rise to stronger effects. When it comes to type of relational information, we believe that the correlational finding emerging from this meta-analysis, according to which implicit evaluations are more responsive to information on logical relations and to narrative information than to causal information and validity information, will require experimental confirmation.

In addition, we found that implicit evaluations were most consistently affected by relational information when they were measured by the AMP and considerably less consistently when they were measured by the IAT or EPT. Crucially, this effect was unique to cases when the evaluative implications of co-occurrence information and relational information diverged from each other; when the two converged, the magnitude of the effects across the three types of measure was in line with previous findings about their general psychometric strength and responsiveness to evaluative information (Bar-Anan & Nosek, 2014, 2017). Notably, we had tentatively predicted that among the three types of measure, the IAT might exhibit the strongest relational effects given that it uniquely involves effortful stimulus categorization. This prediction was clearly not borne out by the data. As such, we recommend that, whenever possible, researchers include both the IAT and the AMP in examinations of relational effects. Moreover, we urge renewed theoretical and empirical focus on the sources of differences in learning effects observed using different indirect measures.

Worryingly with regard to the generality of the overall metaanalytic effect, we found that relational information tended to influence implicit evaluations only for novel and not for known targets. Notably, the opposite effect was observed for explicit evaluations, which tended to be modulated by relational information for known but not for novel targets. Moreover, target type had no significant effect on implicit evaluations when co-occurrence information and relational information had convergent evaluative implications (i.e., the A–B comparison). We believe that this correlational finding should be subjected to experimental testing using a wide variety of designs directly manipulating the type of target (known vs. novel).

For this variable, direct experimental tests seem especially important given that the meta-analytic estimate for known targets was derived from a noticeably smaller set of effect sizes (k = 34) than the estimate for novel targets (k = 169), thus resulting in considerably more uncertainty around the former estimate. At a more general level, the lack of tests involving existing social categories highlights an unfortunate disconnect between theoretically driven experimental work on the basic processes underlying implicit evaluation and translational work that aims to produce change in preexisting social attitudes (most notably, attitudes toward social categories). Given the potential benefits to both fields, it is our hope that future work will be able to bridge this gap.

Sources of Variability to Be Explored in Future Work

We believe that the most important result emerging from the present meta-analysis is that implicit evaluations can be (and, more often than not, are) susceptible to the effects of relational information. Notably, failures of relational information to produce effects in specific studies cannot be taken as evidence against the general idea of such susceptibility. After all, not all relational information is created equal, and it is well-known that some persuasion attempts are more successful than others even with explicit evaluations as the dependent measure.¹⁵ In line with this idea, we found that explicit and implicit evaluations tended to shift in tandem, implying that relevant experiments differed from each other in terms of the intrinsic effectiveness of the information on which they relied, irrespective of the explicit–implicit distinction.

Regarding potential sources of such differences in effectiveness, a large body of past work suggests that persuasion attempts tend to be more effective when participants have the motivation to engage with the material (Cacioppo et al., 1986), when the context is more rather than less conducive to learning (e.g., depending on group size; Harkins & Petty, 1983), and when the message contains strong

¹⁵ We thank an anonymous reviewer for raising this point.

rather than weak arguments (Petty & Wegener, 1991). Given the amount of overlap observed between explicit and implicit evaluations in the present meta-analysis, we believe that these same and many other factors may be fruitfully explored in future empirical work to understand the variability in the strength of relational influences on implicit evaluation. Specifically, we hope that investigators will examine at least three potential moderators whose effects could not be tested in the present meta-analysis due to lack of sufficient variability among the research reports included: (a) a more fine-grained typology of relational information, (b) processing conditions during encoding, and (c) the distinction between immediate and enduring effects.

Types of Relational Information

The categorization of relational information in the present metaanalysis was fairly crude and was not able to fully illuminate theoretical differences across different types of relational manipulations. Notably, in a rare direct experimental test, Mann et al. (2020) found that the effects of negative stimulus pairings were overturned by relational information describing the target as a hero but not by relational information instructing participants to mentally reverse the valence of stimulus pairings. Of course, these two manipulations differ from each other in several ways and, as such, it is impossible to know with certainty why a difference may have emerged between them. An answer to this question, and the more general question of what makes manipulations of relational information successful, would require more systematic investigation.

At a more general level, we believe that, given the state of the evidence, the time is ripe to leave existence proof demonstrations of the effects of relational information on implicit evaluation behind and to instead ask what makes specific types of relational information more or less effective in shifting implicit evaluations. In addition to the factors mentioned above—such as motivation, context, and argument strength—other potentially impactful features may include episodic detail, diagnosticity, as well as the distinction between novel learning and the reinterpretation of old information (see Cone et al., 2017; Kurdi & Dunham, 2020; Mann et al., 2020). However, these proposals are yet to be investigated in sufficient detail. Moreover, given the vastness of the space of relational manipulations, we anticipate that different, and potentially more useful, proposals will be made in the future about how this space should best be carved up in a theoretically meaningful way.

Encoding Conditions

We believe that it would be equally important to systematically investigate whether and how the effects of relational information on implicit evaluations may be modulated by encoding conditions during learning. In the overwhelming majority of studies included in the meta-analytic database, participants were specifically asked to memorize the information to which they were about to be exposed and instructed that they would need to use the information later. In addition, encoding conditions did not tend to be particularly challenging given that participants had the opportunity to fully focus on the relevant evaluative information, without any external distractions. These conditions seem to favor the possibility of effortful propositional processing, thus potentially stacking the deck in favor of finding effects of relational information (although propositional reasoning need not always be effortful; De Neys & Pennycook, 2019; Quilty-Dunn & Mandelbaum, 2018).

In line with this conjecture, what little evidence exists on encoding conditions tends to suggest that these conditions may moderate the effects of relational information on implicit evaluations to a considerable degree. Specifically, Moran et al. (2015) found that implicit evaluations were more likely to reflect relational influences when participants had been instructed to form impressions of the targets rather than to memorize co-occurrences. Similarly, recent results by Fan et al. (2021) show that participants tend to spontaneously incorporate inferences about stimulus relations into implicit evaluations under optimal processing conditions during encoding but not when simultaneously completing a demanding secondary task. As such, we see investigations of encoding conditions as a priority for future work in this domain and as an important step toward bridging the gap between learning effects observed under relatively pristine experimental conditions and under more ecologically realistic conditions often characterized by the presence of competing sources of information.

Long-Term Effects

Equally missing from most relevant research reports are tests of long-term effects (for exceptions, see Cone et al., 2019, 2021; Kurdi & Banaji, 2019). Notably, in one of the few available studies, Kurdi and Banaji (2019) found that the effects of co-occurrence information tended to persist longer than the effects of relational information, even over a relatively short timescale of less than a full hour. However, the scope and generalizability of this difference is unclear given the near-complete absence of relevant empirical work.

In addition to the issue of long-term effects, it would be similarly important to interrogate the relative resistance of co-occurrence information and relational information to counterattitudinal information—a question often investigated in the context of explicit evaluations (e.g., Knowles & Linn, 2004). Remarkably, recent experiments by Kurdi, Mann, et al. (2021) have found that although certain forms of propositional learning can produce strong immediate and long-term effects on implicit evaluation, these same effects are also highly vulnerable to reinstatement via simple and relatively weak forms of co-occurrence information. The issues of long-term persistence and robustness to counterattitudinal information are all the more crucial because they have the potential to inform not only research on the basic nature of implicit evaluation and evaluative learning but also theoretically driven interventions.

Additional Implications for Method and Theory

Finally, although this meta-analysis was not designed to directly test these questions, the present results also speak to two longstanding theoretical controversies: one regarding the malleability of implicit evaluations and the second about the validity of indirect measures.

Broader Theoretical Implications

The present findings are difficult to reconcile with early theoretical ideas about implicit evaluations being generally resistant to change (Bargh, 1999; Devine, 1989; Fiske, 1998; Wilson et al., 2000).

Rather, they are broadly in line with early work that emphasized the flexibility of implicit evaluation (e.g., Blair, 2002; Ferguson & Bargh, 2004; Moskowitz et al., 1999) as well as with recent empirical results (Lai et al., 2014) and theoretical perspectives (Cone et al., 2017; De Houwer, 2014; De Houwer et al., 2020; Kurdi & Dunham, 2020; Mandelbaum, 2016) according to which implicit evaluations should flexibly respond to evaluative information encountered in the environment. The results reported above are also easily reconcilable with approaches positing that implicit and explicit evaluations reflect the same underlying representations (independent of how they were formed) and differ mainly in the conditions under which these representations are expressed in overt behavior (e.g., Fazio, 1990, 2007).

Implications for the Validity of Indirect Measures

The present findings provide solid evidence that indirect measures are capable of reflecting the effects of evaluative information across a wide range of contexts, thus suggesting that they, indeed, index evaluative representations (Kurdi, Ratliff, & Cunningham, 2021). In addition, the convergence between explicit and implicit evaluations both at the level of mean patterns and individual effect sizes suggests that the two reflect similar mental content, thus providing further evidence for construct validity.

At the same time, these results are silent on the validity of indirect measures as measures of individual differences. Moreover, it is not clear whether the same strong results of convergence would generalize to social targets that are well-known to participants, and especially social targets subject to strong social desirability concerns (Nosek, 2005). After all, dissociations between implicit and explicit evaluation are particularly frequent in the context of such targets, and such targets were markedly underrepresented in the present meta-analytic sample. More generally, given that implicit–explicit dissociations are well-documented in the literature, they require a theoretical explanation. Based on the results of the present meta-analysis, differences in sensitivity to co-occurrence information and relational information may not be a particularly strong contender to account for these effects.

Moreover, the meta-analytic effect sizes for implicit evaluations were considerably smaller than the corresponding meta-analytic effect sizes for explicit evaluations. As such, differences observed across the two types of evaluation may be explained not by substantive differences between explicit and implicit social cognition at the level of processes or mechanisms but rather by differences in the psychometric strength of direct and indirect measures (see Footnote 14). In line with this possibility, a difference across measures emerged even when the evaluative implications of cooccurrence information and relational information converged and, as such, the results could not be attributed to differential sensitivity of each type of measure to relational information. Although some theoretical perspectives suggest that the relatively noisy nature of indirect measures is a feature rather than a bug (Dalege & van der Maas, 2020), we believe that improving the psychometric properties of these measures should be a priority in methodologically oriented social cognition research.

It should also be noted that overt responses on any measure of memory or evaluation (Jacoby, 1991), and measures of implicit evaluation in particular (Conrey et al., 2005; Payne et al., 2010), can be shown to emerge from an interplay of different cognitive mechanisms. As such, given that they rely on summary scores rather than formal modeling of constituent processes, the present results are silent as to whether relational information affects the relatively more automatic or relatively more controlled components of implicit evaluation. Some recent work relying on multinomial processing trees (Hütter & Klauer, 2016) has been able to circumvent this limitation of the work relying on direct measures to capture explicit evaluations and indirect measures to capture implicit evaluations by identifying relatively more controlled and relatively more automatic aspects of evaluative responding using the same overt behavioral response (e.g., Heycke & Gawronski, 2020; Kukken et al., 2020). Although this type of work has yet to produce a sufficient number of effect sizes to allow for inclusion in this meta-analysis, we are excited about the potential benefits of this approach and about the insights that it might be able to generate in the future.

We are equally excited about the possibility of developing computational models to more formally capture the different inputs and processes giving rise to evaluative learning. Such computational models may help advance theories with sufficient precision to be able to derive falsifiable, quantitative predictions from them (Smaldino, 2017). Models of this kind may also be helpful in avoiding the types of conceptual ambiguities that we faced in attempting to identify sources of co-occurrence information and relational information when conducting the present meta-analysis. Such ambiguities include whether (a) behavioral statements with minimal relational content should qualify as co-occurrence information or relational information and (b) spatiotemporal features of stimulus pairings are better conceived of as a source of relational information or as part of the co-occurrence information itself. Associative learning has a long history of formal modeling (e.g., Rescorla & Wagner, 1972), and recent years have also seen remarkable advances in the modeling of value-based learning (Dayan & Niv, 2008) and high-level reasoning processes in humans (e.g., Tenenbaum et al., 2011). As such, the field of evaluative learning has a variety of different sources to draw from in the development of formal models.

Finally, some readers may wonder what the present results imply for the validity of the AMP as a measure of implicit evaluations. After all, the validity (and specifically the indirect nature) of the AMP has been questioned in past work (Bar-Anan & Nosek, 2012; Hughes et al., 2022). Moreover, in a multitrait–multimethod investigation of different direct and indirect measures of attitudes, unlike the other indirect measures included, the AMP showed similar loadings on the implicit and explicit latent constructs (Bar-Anan & Vianello, 2018). Seemingly in line with such concerns, the present results indicate that the AMP is more susceptible to propositional influences than other widely used indirect measures (and in that way may appear more similar to direct measures).

We do not share these concerns about the validity of the AMP. We believe that in its totality, available evidence overwhelmingly supports the construct validity of the AMP as a measure of implicit evaluations (for a review, see Payne & Lundberg, 2014), and the specific concerns raised by Bar-Anan and Nosek (2012) and Hughes et al. (2022) have been laid to rest in follow-up work by Payne et al. (2013) and Kurdi, Melnikoff, et al. (2022). Moreover, we are generally of the view that susceptibility to relational information should not be used to judge the construct validity of measures of implicit evaluation. Rather, it seems to be more helpful to ask whether the measure involves automatic (unintentional) retrieval of evaluative information (De Houwer et al., 2009; De Houwer & Moors, 2010) and to leave questions of sensitivity to different types of input and process up to substantive theoretical and empirical work.

Conclusions and Future Directions

The results of the present meta-analysis leave little doubt that implicit evaluations can be influenced by relational information even in the presence of co-occurrence information with conflicting evaluative implications. What is more, it seems that this effect is the rule rather than the exception. Given the strength of the evidence about the possibility of this type of effect, we believe that further proof-of-concept studies will make little if any incremental contribution to the literature at this time. Instead, we urge investigators to reorient their attention and resources toward investigating the mechanism(s) underlying and boundary conditions of these effects. In light of the results of the present meta-analysis, and especially the considerable amounts of both explained and unexplained heterogeneity in effect sizes, such inquiry will be crucial if theoretical progress on the basic nature of implicit (and explicit) evaluations and evaluative learning is to be made in the coming years.

The present results can also help identify specific issues that seem especially ripe for empirical investigation and theorizing. From our perspective, some of the most important open questions include the following: (a) Why and via what mechanism(s) does co-occurrence information influence explicit and implicit evaluations in the presence of contradictory relational information? Are the mechanisms the same for explicit and implicit evaluations? Are they best characterized as associative, propositional, or in some other way? (b) Can relational information affect implicit evaluations of known targets (including social categories), and if so, under what conditions? If the effects are restricted to novel targets, why is this the case? (c) Why does the AMP show stronger effects of relational information than the IAT and the EPT? Is the difference due to a relatively superficial difference in procedures, or a theoretically relevant difference between the two types of indirect measure? (d) What types of relational information are especially powerful in shifting implicit evaluations? What features should be used to carve up the space of relational manipulations in a way that is best suited for implicit social cognition research? (e) Can relational information exert effects on implicit evaluations even under relatively suboptimal encoding conditions? (f) Can the effects of relational information on implicit evaluations endure over time? Under what conditions can they be resistant to reinstatement effects?

In conclusion, we hope that the findings emerging from this metaanalysis will help reorient theoretical debates, spur the development of new accounts of social attitudes and evaluative learning, and shape the agenda of implicit social cognition research in the years ahead.

References

References marked with an asterisk indicate studies included in the metaanalysis.

*Akalis, S. A. (2008). A new spin on losses looming larger than gains: Asymmetric implicit associations from slot machine experience. *Journal* of Behavioral Decision Making, 21(4), 378–398. https://doi.org/10.1002/ bdm.594

- Amodio, D. M., & Ratner, K. G. (2011). A memory systems model of implicit social cognition. *Current Directions in Psychological Science*, 20(3), 143–148. https://doi.org/10.1177/0963721411408562
- APA Publications and Communications Board Working Group on Journal Article Reporting Standards. (2008). Reporting standards for research in psychology: Why do we need them? What might they be? *American Psychologist*, 63(9), 839–851. https://doi.org/10.1037/0003-066X.63.9.839
- *Bading, K., Stahl, C., & Rothermund, K. (2020). Why a standard IAT effect cannot provide evidence for association formation: The role of similarity construction. *Cognition and Emotion*, 34(1), 128–143. https://doi.org/10 .1080/02699931.2019.1604322
- Baeyens, F., Eelen, P., Crombez, G., & Van den Bergh, O. (1992). Human evaluative conditioning: Acquisition trials, presentation schedule, evaluative style and contingency awareness. *Behaviour Research and Therapy*, 30(2), 133–142. https://doi.org/10.1016/0005-7967(92)90136-5
- Bar-Anan, Y., & Nosek, B. A. (2012). Reporting intentional rating of the primes predicts priming effects in the affective misattribution procedure. *Personality and Social Psychology Bulletin*, 38(9), 1194–1208. https:// doi.org/10.1177/0146167212446835
- Bar-Anan, Y., & Nosek, B. A. (2014). A comparative investigation of seven indirect attitude measures. *Behavior Research Methods*, 46(3), 668–688. https://doi.org/10.3758/s13428-013-0410-6
- Bar-Anan, Y., & Nosek, B. A. (2017). A comparison of the sensitivity of four indirect evaluation measures to evaluative information. SocArXiv. https:// doi.org/10.31235/osf.io/edw6z
- Bar-Anan, Y., & Vianello, M. (2018). A multi-method multi-trait test of the dual-attitude perspective. *Journal of Experimental Psychology: General*, 147(8), 1264–1272. https://doi.org/10.1037/xge0000383
- Bargh, J. A. (1999). The cognitive monster: The case against the controllability of automatic stereotype effects. In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology* (pp. 361–382). Guilford Press.
- *Berry, T. R. (2020). You do (not?) have to go to the gym: Effects of negations in exercise messages. Sport Exercise and Performance Psychology, 9(3), 437–449. https://doi.org/10.1037/spy0000185
- Blair, I. V. (2002). The malleability of automatic stereotypes and prejudice. *Personality and Social Psychology Review*, 6(3), 242–261. https://doi.org/ 10.1207/S15327957PSPR0603_8
- Borenstein, M. (2009). Effect sizes for continuous data. In H. Cooper, L. V. Hedges, & J. C. Valentine (Eds.), *The handbook of research synthesis and meta-analysis* (2nd ed., pp. 221–235). Russell Sage Foundation.
- Borenstein, M., Higgins, J. P. T., Hedges, L. V., & Rothstein, H. R. (2017). Basics of meta-analysis: 12 is not an absolute measure of heterogeneity. *Research Synthesis Methods*, 8(1), 5–18. https://doi.org/10.1002/jrsm.1230
- *Boucher, K. L., & Rydell, R. J. (2012). Impact of negation salience and cognitive resources on negation during attitude formation. *Personality and Social Psychology Bulletin*, 38(10), 1329–1342. https://doi.org/10.1177/ 0146167212450464
- Brannon, S. M., & Gawronski, B. (2017). A second chance for first impressions? Exploring the context-(in)dependent updating of implicit evaluations. *Social Psychological and Personality Science*, 8(3), 275–283. https://doi.org/10.1177/1948550616673875
- Cacioppo, J. T., Petty, R. E., Chuan, F. K., & Rodriguez, R. (1986). Central and peripheral routes to persuasion: An individual difference perspective. *Journal of Personality and Social Psychology*, 51(5), 1032–1043. https:// doi.org/10.1037/0022-3514.51.5.1032
- Calanchini, J., Gonsalkorale, K., Sherman, J. W., & Klauer, K. C. (2013). Counter-prejudicial training reduces activation of biased associations and enhances response monitoring. *European Journal of Social Psychology*, 43(5), 321–325. https://doi.org/10.1002/ejsp.1941
- Carey, S. (2011). The origin of concepts. Oxford University Press.
- *Carraro, L., Gawronski, B., & Castelli, L. (2010). Losing on all fronts: The effects of negative versus positive person-based campaigns on implicit and

explicit evaluations of political candidates. *British Journal of Social Psychology*, 49(3), 453–470. https://doi.org/10.1348/014466609X468042

- Chaiken, S. (1980). Heuristic versus systematic information processing and the use of source versus message cues in persuasion. *Journal of Personality and Social Psychology*, 39(5), 752–766. https://doi.org/10.1037/ 0022-3514.39.5.752
- *Chang, Y. (2018). Exploring the patterns of dual attitude changes in the context of athlete endorsement: The interplay of fit, evaluative conditioning, and introspection focus. *Journal of Sport Management*, 32(5), 413– 425. https://doi.org/10.1123/jsm.2017-0149
- Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, 82(6), 407–428. https:// doi.org/10.1037/0033-295X.82.6.407
- Cone, J., & Ferguson, M. J. (2015). He did what? The role of diagnosticity in revising implicit evaluations. *Journal of Personality and Social Psychol*ogy, 108(1), 37–57. https://doi.org/10.1037/pspa0000014
- *Cone, J., Flaharty, K., & Ferguson, M. J. (2019). Believability of evidence matters for correcting social impressions. *Proceedings of the National Academy of Sciences of the United States of America*, 116(20), 9802– 9807. https://doi.org/10.1073/pnas.1903222116
- *Cone, J., Flaharty, K., & Ferguson, M. J. (2021). The long-term effects of new evidence on implicit impressions of other people. *Psychological Science*, 32(2), 173–188. https://doi.org/10.1177/0956797620963559
- 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. https:// doi.org/10.1016/bs.aesp.2017.03.001
- Conrey, F. R., Sherman, J. W., Gawronski, B., Hugenberg, K., & Groom, C. J. (2005). Separating multiple processes in implicit social cognition: The quad model of implicit task performance. *Journal of Personality and Social Psychology*, 89(4), 469–487. https://doi.org/10.1037/0022-3514.89 .4.469
- Crano, W. D., & Prislin, R. (2006). Attitudes and persuasion. Annual Review of Psychology, 57(1), 345–374. https://doi.org/10.1146/annurev.psych.57 .102904.190034
- Csibra, G., & Gergely, G. (2009). Natural pedagogy. *Trends in Cognitive Sciences*, 13(4), 148–153. https://doi.org/10.1016/j.tics.2009.01.005
- Cunningham, W. A., Zelazo, P. D., Packer, D. J., & Van Bavel, J. J. (2007). The iterative reprocessing model: A multilevel framework for attitudes and evaluation. *Social Cognition*, 25(5), 736–760. https://doi.org/10.1521/ soco.2007.25.5.736
- Dalege, J., & van der Maas, H. L. J. (2020). Accurate by being noisy: A formal network model of implicit measures of attitudes. *Social Cognition*, 38(Suppl.), S26–S41. https://doi.org/10.1521/soco.2020.38.supp.s26
- Dayan, P., & Niv, Y. (2008). Reinforcement learning: The good, the bad and the ugly. *Current Opinion in Neurobiology*, 18(2), 185–196. https:// doi.org/10.1016/j.conb.2008.08.003
- De Houwer, J. (2014). A propositional model of implicit evaluation. *Social* and Personality Psychology Compass, 8(7), 342–353. https://doi.org/10 .1111/spc3.12111
- De Houwer, J. (2018a). A functional-cognitive perspective on the relation between conditioning and placebo research. *International Review of Neurobiology*, 138, 95–111. https://doi.org/10.1016/bs.irn.2018.01.007
- De Houwer, J. (2018b). Propositional models of evaluative conditioning. *Social Psychological Bulletin*, *13*(3), 49–21. https://doi.org/10.5964/spb .v13i3.28046
- De Houwer, J., & Hughes, S. (2016). Evaluative conditioning as a symbolic phenomenon: On the relation between evaluative conditioning, evaluative conditioning via instructions, and persuasion. *Social Cognition*, 34(5), 480–494. https://doi.org/10.1521/soco.2016.34.5.480
- De Houwer, J., & Moors, A. (2010). Implicit measures: Similarities and differences. In B. Gawronski & B. K. Payne (Eds.), *Handbook of implicit* social cognition: Measurement, theory, and applications (pp. 176–196). Guilford Press.

- De Houwer, J., Teige-Mocigemba, S., Spruyt, A., & Moors, A. (2009). Implicit measures: A normative analysis and review. *Psychological Bulletin*, 135(3), 347–368. https://doi.org/10.1037/a0014211
- De Houwer, J., Van Dessel, P., & Moran, T. (2020). Attitudes beyond associations: On the role of propositional representations in stimulus evaluation. Advances in Experimental Social Psychology, 61, 127–183. https://doi.org/10.1016/bs.aesp.2019.09.004
- *De Houwer, J., & Vandorpe, S. (2010). Using the Implicit Association Test as a measure of causal learning does not eliminate effects of rule learning. *Experimental Psychology*, 57(1), 61–67. https://doi.org/10.1027/1618-3169/a000008
- De Neys, W., & Pennycook, G. (2019). Logic, fast and slow: Advances in dual-process theorizing. *Current Directions in Psychological Science*, 28(5), 503–509. https://doi.org/10.1177/0963721419855658
- *DeCoster, J., Banner, M. J., Smith, E. R., & Semin, G. R. (2006). On the inexplicability of the implicit: Differences in the information provided by implicit and explicit tests. *Social Cognition*, 24(1), 5–21. https://doi.org/10 .1521/soco.2006.24.1.5
- *Deutsch, R., Gawronski, B., & Strack, F. (2006). At the boundaries of automaticity: Negation as reflective operation. *Journal of Personality* and Social Psychology, 91(3), 385–405. https://doi.org/10.1037/0022-3514.91.3.385
- *Deutsch, R., Kordts-Freudinger, R., Gawronski, B., & Strack, F. (2009). Fast and fragile: A new look at the automaticity of negation processing. *Experimental Psychology*, 56(6), 434–446. https://doi.org/10.1027/1618-3169.56.6.434
- Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology*, 56(1), 5–18. https://doi.org/10.1037/0022-3514.56.1.5
- Dickinson, A. (2012). Associative learning and animal cognition. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1603), 2733–2742. https://doi.org/10.1098/rstb.2012.0220
- Eagly, A. H., & Chaiken, S. (1993). *The psychology of attitudes*. Harcourt Brace Jovanovich College Publishers.
- *Evans, C. D. (2008). The effect of implicit theories of judgment on attitudes and evaluative outcomes [Doctoral dissertation, Miami University]. Ohio-LINK Electronic Theses and Dissertations Center. http://rave.ohiolink .edu/etdc/view?acc_num=miami1214597292
- Fan, X., Bodenhausen, G. V., & Lee, A. Y. (2021). Acquiring favorable attitudes based on aversive affective cues: Examining the spontaneity and efficiency of propositional evaluative conditioning. *Journal of Experimental Social Psychology*, 95, Article 104139. https://doi.org/10.1016/j.je sp.2021.104139
- Fazio, R. H. (1990). Multiple processes by which attitudes guide behavior: The MODE model as an integrative framework. *Advances in Experimental Social Psychology*, 23, 75–109. https://doi.org/10.1016/S0065-2601(08) 60318-4
- Fazio, R. H. (2007). Attitudes as object–evaluation associations of varying strength. *Social Cognition*, 25(5), 603–637. https://doi.org/10.1521/soco .2007.25.5.603
- Fazio, R. H., & Olson, M. A. (2003). Implicit measures in social cognition research: Their meaning and use. *Annual Review of Psychology*, 54(1), 297–327. https://doi.org/10.1146/annurev.psych.54.101601.145225
- 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(2), 229–238. https://doi.org/10.1037/0022-3514.50.2.229
- Ferguson, M. J., & Bargh, J. A. (2004). Liking is for doing: The effects of goal pursuit on automatic evaluation. *Journal of Personality and Social Psychology*, 87(5), 557–572. https://doi.org/10.1037/0022-3514.87.5.557
- Ferguson, M. J., & Bargh, J. A. (2008). Evaluative readiness: The motivational nature of automatic evaluation. In A. J. Elliot (Ed.), *Handbook of* approach and avoidance motivation (pp. 287–304). Psychology Press.
- Ferguson, M. J., Mann, T. C., Cone, J., & Shen, X. (2019). When and how implicit first impressions can be updated. *Current Directions in*

Psychological Science, 28(4), 331–336. https://doi.org/10.1177/096372 1419835206

- Fisher, Z., & Tipton, E. (2015). robumeta: An R-package for robust variance estimation in meta-analysis. arXiv. https://doi.org/10.48550/arXiv.1503. 02220
- Fiske, S. T. (1998). Stereotyping, prejudice, and discrimination. In D. Gilbert, S. Fiske, & G. Lindzey (Eds.), *The handbook of social psychology* (4th ed., Vol. 2, pp. 357–411). McGraw-Hill.

Fodor, J. A. (1975). The language of thought. Harvard University Press.

- Förderer, S., & Unkelbach, C. (2012). Hating the cute kitten or loving the aggressive pit-bull: EC effects depend on CS–US relations. *Cognition* and Emotion, 26(3), 534–540. https://doi.org/10.1080/02699931.2011 .588687
- *Förderer, S., & Unkelbach, C. (2016). Changing US attributes after CS–US pairings changes CS-attribute-assessments. *Personality and Social Psychology Bulletin*, 42(3), 350–365. https://doi.org/10.1177/0146167215 626705
- Forscher, P. S., Mitamura, C., Dix, E. L., Cox, W. T. L., & Devine, P. G. (2017). Breaking the prejudice habit: Mechanisms, timecourse, and longevity. *Journal of Experimental Social Psychology*, 72, 133–146. https://doi.org/10.1016/j.jesp.2017.04.009

Gallistel, C. R. (1990). The organization of learning. MIT Press.

- *Gast, A., & De Houwer, J. (2013). The influence of extinction and counterconditioning instructions on evaluative conditioning effects. *Learning and Motivation*, 44(4), 312–325. https://doi.org/10.1016/j .lmot.2013.03.003
- 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. https://doi.org/ 10.1037/0033-2909.132.5.692
- Gawronski, B., & Bodenhausen, G. V. (2018). Evaluative conditioning from the perspective of the associative–propositional evaluation model. *Social Psychological Bulletin*, *13*(3), Article e28024. https://doi.org/10.5964/spb .v13i3.28024
- *Gawronski, B., & Walther, E. (2008). The TAR effect: When the ones who dislike become the ones who are disliked. *Personality and Social Psychology Bulletin*, 34(9), 1276–1289. https://doi.org/10.1177/014616 7208318952
- *Gawronski, B., Walther, E., & Blank, H. (2005). Cognitive consistency and the formation of interpersonal attitudes: Cognitive balance affects the encoding of social information. *Journal of Experimental Social Psychology*, 41(6), 618–626. https://doi.org/10.1016/j.jesp.2004 .10.005
- *Green, L. J. S., Luck, C. C., Gawronski, B., & Lipp, O. V. (2021). Contrast effects in backward evaluative conditioning: Exploring effects of affective relief/disappointment versus instructional information. *Emotion*, 21(2), 350–359. https://doi.org/10.1037/emo0000701
- Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review*, 102(1), 4–27. https://doi.org/10.1037/0033-295X.102.1.4
- Greenwald, A. G., & Banaji, M. R. (2017). The implicit revolution: Reconceiving the relation between conscious and unconscious. *American Psychologist*, 72(9), 861–871. https://doi.org/10.1037/amp0000238
- Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The implicit association test. *Journal of Personality and Social Psychology*, 74(6), 1464–1480. https:// doi.org/10.1037/0022-3514.74.6.1464
- Grice, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and semantics, Vol. 3: Speech acts* (pp. 41–58). Academic Press.
- Gwet, K. L. (2010). Computing inter-rater reliability and its variance in the presence of high agreement. *British Journal of Mathematical & Statistical Psychology*, 61(1), 29–48. https://doi.org/10.1348/00071100 6X126600

- Harkins, S. G., & Petty, R. E. (1983). Social context effects in persuasion: The effects of multiple sources and multiple targets. In P. B. Paulus (Ed.), *Basic group processes* (pp. 149–175). Springer. https://doi.org/10.1007/ 978-1-4612-5578-9_7
- Harrer, M., Cuijpers, P., Furukawa, T. A., & Ebert, D. D. (2021). Doing meta-analysis with R: A hands-on guide. Chapman & Hall. https://doi.org/ 10.1201/9781003107347
- Hebb, D. O. (1949). The organization of behavior: A neuropsychological theory. Wiley.
- Heycke, T., & Gawronski, B. (2020). Co-occurrence and relational information in evaluative learning: A multinomial modeling approach. *Journal* of Experimental Psychology: General, 149(1), 104–124. https://doi.org/10 .1037/xge0000620
- *Heycke, T., Gehrmann, S., Haaf, J. M., & Stahl, C. (2018). Of two minds or one? A registered replication of Rydell et al. (2006). *Cognition and Emotion*, 32(8), 1708–1727. https://doi.org/10.1080/02699931.2018.14 29389
- Hofmann, W., De Houwer, J., Perugini, M., Baeyens, F., & Crombez, G. (2010). Evaluative conditioning in humans: A meta-analysis. *Psychological Bulletin*, 136(3), 390–421. https://doi.org/10.1037/a0018916
- *Högden, F., & Unkelbach, C. (2021). The role of relational qualifiers in attribute conditioning: Does disliking an athletic person make you unathletic? *Personality and Social Psychology Bulletin*, 47(4), 643–656. https:// doi.org/10.1177/0146167220945538
- *Hu, X., Gawronski, B., & Balas, R. (2017a). Propositional versus dualprocess accounts of evaluative conditioning: I. The effects of cooccurrence and relational information on implicit and explicit evaluations. *Personality and Social Psychology Bulletin*, 43(1), 17–32. https://doi.org/ 10.1177/0146167216673351
- *Hu, X., Gawronski, B., & Balas, R. (2017b). Propositional versus dualprocess accounts of evaluative conditioning: II. The effectiveness of counter-conditioning and counter-instructions in changing implicit and explicit evaluations. *Social Psychological and Personality Science*, 8(8), 858–866. https://doi.org/10.1177/1948550617691094
- Hughes, S., Barnes-Holmes, D., Van Dessel, P., de Almeida, J. H., Stewart, I., & De Houwer, J. (2018). On the symbolic generalization of likes and dislikes. *Journal of Experimental Social Psychology*, 79, 365–377. https:// doi.org/10.1016/j.jesp.2018.09.002
- Hughes, S., Cummins, J., & Hussey, I. (2022). Effects on the affect misattribution procedure are strongly moderated by influence awareness. *Behavior Research Methods*. Advance online publication. https://doi.org/ 10.3758/s13428-022-01879-4
- *Hughes, S., Mattavelli, S., & De Houwer, J. (2018). Examining the impact of distance as a contextual cue in evaluative conditioning. *PLOS ONE*, *13*(10), Article e0204855. https://doi.org/10.1371/journal.pone .0204855
- *Hughes, S., Ye, Y., & De Houwer, J. (2019). Evaluative conditioning effects are modulated by the nature of contextual pairings. *Cognition* and Emotion, 33(5), 871–884. https://doi.org/10.1080/02699931.2018 .1500882
- *Hughes, S., Ye, Y., Van Dessel, P., & De Houwer, J. (2019). When people co-occur with good or bad events: Graded effects of relational qualifiers on evaluative conditioning. *Personality and Social Psychology Bulletin*, 45(2), 196–208. https://doi.org/10.1177/0146167218781340
- Hütter, M., & Klauer, K. C. (2016). Applying processing trees in social psychology. *European Review of Social Psychology*, 27(1), 116–159. https://doi.org/10.1080/10463283.2016.1212966
- Jacoby, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language*, 30(5), 513–541. https://doi.org/10.1016/0749-596X(91)90025-F
- *Johnson, I. R., Kopp, B. M., & Petty, R. E. (2018). Just say no! (and mean it): Meaningful negation as a tool to modify automatic racial attitudes. *Group Processes & Intergroup Relations*, 21(1), 88–110. https://doi.org/ 10.1177/1368430216647189

- Jones, C. R., Fazio, R. H., & Olson, M. A. (2009). Implicit misattribution as a mechanism underlying evaluative conditioning. *Journal of Personality* and Social Psychology, 96(5), 933–948. https://doi.org/10.1037/a00 14747
- *Kasran, S., Hughes, S., & De Houwer, J. (2022). Observational evaluative conditioning is sensitive to relational information. *Quarterly Journal of Experimental Psychology*, 75(11), 2043–2063. https://doi.org/10.1177/ 17470218221080471
- Kawakami, K., Dovidio, J. F., Moll, J., Hermsen, S., & Russin, A. (2000). Just say no (to stereotyping): Effects of training in the negation of stereotypic associations on stereotype activation. *Journal of Personality* and Social Psychology, 78(5), 871–888. https://doi.org/10.1037/0022-3514.78.5.871
- Knowles, E. S., & Linn, J. A. (2004). Resistance and persuasion. Psychology Press. https://doi.org/10.4324/9781410609816
- *Koranyi, N., Gast, A., & Rothermund, K. (2013). "Although quite nice, I was somehow not attracted by that person": Attitudes toward romantically committed opposite-sex others are immune to positive evaluative conditioning. *Social Psychological and Personality Science*, 4(4), 403–410. https://doi.org/10.1177/1948550612467037
- Kukken, N., Hütter, M., & Holland, R. W. (2020). Are there two independent evaluative conditioning effects in relational paradigms? Dissociating the effects of CS–US pairings and their meaning. *Cognition and Emotion*, 34(1), 170–187. https://doi.org/10.1080/02699931.2019 .1617112
- Kurdi, B., & Banaji, M. R. (2017). Repeated evaluative pairings and evaluative statements: How effectively do they shift implicit attitudes? *Journal of Experimental Psychology: General*, 146(2), 194–213. https:// doi.org/10.1037/xge0000239
- *Kurdi, B., & Banaji, M. R. (2019). Attitude change via repeated evaluative pairings versus evaluative statements: Shared and unique features. *Journal* of Personality and Social Psychology, 116(5), 681–703. https://doi.org/10 .1037/pspa0000151
- Kurdi, B., & Banaji, M. R. (2022). Implicit person memory: Domain-general and domain-specific processes of learning and change. In G. B. Moskowitz & E. Balcetis (Eds.), *The handbook of impression formation: A social psychological approach* (pp. 459–488). Routledge. https://doi.org/10.4324/9781003045687-25
- Kurdi, B., & Dunham, Y. (2020). Propositional accounts of implicit evaluation: Taking stock and looking ahead. *Social Cognition*, 38(Suppl.), S42–S67. https://doi.org/10.1521/soco.2020.38.supp.s42
- Kurdi, B., & Dunham, Y. (2021). Sensitivity of implicit evaluations to accurate and erroneous propositional inferences. *Cognition*, 214, Article 104792. https://doi.org/10.1016/j.cognition.2021.104792
- Kurdi, B., Mann, T. C., Axt, J. R., & Ferguson, M. J. (2021). The prospect of long-term implicit attitude change: Tests of spontaneous recovery and reinstatement [Unpublished manuscript]. Department of Psychology, Yale University.
- Kurdi, B., Mann, T. C., Charlesworth, T. E. S., & Banaji, M. R. (2019). The relationship between implicit intergroup attitudes and beliefs. *Proceedings* of the National Academy of Sciences of the United States of America, 116(13), 5862–5871. https://doi.org/10.1073/pnas.1820240116
- Kurdi, B., Melnikoff, D. E., Hannay, J. W., Korkmaz, A., Lee, K. M., Ritchie, E., Surdel, N., Vuletich, H. A., Yang, X., Payne, B. K., & Ferguson, M. J. (2022). Probing the automaticity features of the affect misattribution procedure: The roles of awareness and intentionality. PsyArXiv. https://doi.org/10.31234/osf.io/akz5v
- *Kurdi, B., Morris, A., & Cushman, F. A. (2022). The role of causal structure in implicit evaluation. *Cognition*, 225, Article 105116. https://doi.org/10 .1016/j.cognition.2022.105116
- Kurdi, B., Ratliff, K. A., & Cunningham, W. A. (2021). Can the Implicit Association Test serve as a valid measure of automatic cognition? A response to Schimmack (2021). *Perspectives on Psychological Science*, 16(2), 422–434. https://doi.org/10.1177/1745691620904080

- Lai, C. K., Marini, M., Lehr, S. A., Cerruti, C., Shin, J.-E. L., Joy-Gaba, J. A., Ho, A. K., Teachman, B. A., Wojcik, S. P., Koleva, S. P., Frazier, R. S., Heiphetz, L., Chen, E. E., Turner, R. N., Haidt, J., Kesebir, S., Hawkins, C. B., Schaefer, H. S., Rubichi, S., ... Nosek, B. A. (2014). Reducing implicit racial preferences: I. A comparative investigation of 17 interventions. *Journal of Experimental Psychology: General*, 143(4), 1765–1785. https://doi.org/10.1037/a0036260
- Lai, C. K., Skinner, A. L., Cooley, E., Murrar, S., Brauer, M., Devos, T., Calanchini, J., Xiao, Y. J., Pedram, C., Marshburn, C. K., Simon, S., Blanchar, J. C., Joy-Gaba, J. A., Conway, J., Redford, L., Klein, R. A., Roussos, G., Schellhaas, F. M. H., Burns, M., ... Nosek, B. A. (2016). Reducing implicit racial preferences: II. Intervention effectiveness across time. *Journal of Experimental Psychology: General*, *145*(8), 1001–1016. https://doi.org/10.1037/xge0000179
- Levey, A. B., & Martin, I. (1975). Classical conditioning of human 'evaluative' responses. *Behaviour Research and Therapy*, 13(4), 221–226. https:// doi.org/10.1016/0005-7967(75)90026-1
- Levey, A. B., & Martin, I. (1990). Evaluative conditioning: Overview and further options. *Cognition and Emotion*, 4(1), 31–37. https://doi.org/10 .1080/02699939008406762
- Levine, T. R., Asada, K. J., & Carpenter, C. (2009). Sample sizes and effect sizes are negatively correlated in meta-analyses: Evidence and implications of a publication bias against nonsignificant findings. *Communication Monographs*, 76(3), 286–302. https://doi.org/10.1080/036377509030 74685
- Lowery, B. S., Hardin, C. D., & Sinclair, S. (2001). Social influence effects on automatic racial prejudice. *Journal of Personality and Social Psychol*ogy, 81(5), 842–855. https://doi.org/10.1037/0022-3514.81.5.842
- *Luck, C. C., & Lipp, O. V. (2018). Verbal instructions targeting valence alter negative conditional stimulus evaluations (but do not affect reinstatement rates). *Cognition and Emotion*, 32(1), 61–80. https://doi.org/10 .1080/02699931.2017.1280449
- Mandelbaum, E. (2016). Attitude, inference, association: On the propositional structure of implicit bias. *Noûs*, 50(3), 629–658. https://doi.org/10 .1111/nous.12089
- Mann, T. C., & Ferguson, M. J. (2015). Can we undo our first impressions? The role of reinterpretation in reversing implicit evaluations. *Journal of Personality and Social Psychology*, 108(6), 823–849. https://doi.org/10 .1037/pspa0000021
- *Mann, T. C., Kurdi, B., & Banaji, M. R. (2020). How effectively can implicit evaluations be updated? Using evaluative statements after aversive repeated evaluative pairings. *Journal of Experimental Psychology: General*, 149(6), 1169–1192. https://doi.org/10.1037/xge0000701
- McConnell, A. R., & Rydell, R. J. (2014). The systems of evaluation model: A dual-systems approach to attitudes. In J. W. Sherman, B. Gawronski, & Y. Trope (Eds.), *Dual-process theories of the social mind* (pp. 204–217). Guilford Press.
- Meyer, D. E., & Schvaneveldt, R. W. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology*, 90(2), 227–234. https://doi.org/10 .1037/h0031564
- Mitchell, C. J., De Houwer, J., & Lovibond, P. F. (2009). The propositional nature of human associative learning. *Behavioral and Brain Sciences*, 32(2), 183–198. https://doi.org/10.1017/S0140525X09000855
- *Moran, T., & Bar-Anan, Y. (2013). The effect of object-valence relations on automatic evaluation. *Cognition and Emotion*, 27(4), 743–752. https:// doi.org/10.1080/02699931.2012.732040
- *Moran, T., & Bar-Anan, Y. (2018). How actions change liking: The effect of an action's outcome on the evaluation of the action's object. *Journal of Experimental Psychology: General*, 147(11), 1597–1618. https://doi.org/ 10.1037/xge0000419
- *Moran, T., & Bar-Anan, Y. (2020). The effect of co-occurrence and relational information on speeded evaluation. *Cognition and Emotion*, 34(1), 144–155. https://doi.org/10.1080/02699931.2019.1604321

- *Moran, T., Bar-Anan, Y., & Nosek, B. A. (2015). Processing goals moderate the effect of co-occurrence on automatic evaluation. *Journal* of Experimental Social Psychology, 60, 157–162. https://doi.org/10.1016/ j.jesp.2015.05.009
- *Moran, T., Bar-Anan, Y., & Nosek, B. A. (2016). The assimilative effect of co-occurrence on evaluation above and beyond the effect of relational qualifiers. *Social Cognition*, 34(5), 435–461. https://doi.org/10.1521/soco .2016.34.5.435
- *Moran, T., Bar-Anan, Y., & Nosek, B. A. (2017). The effect of the validity of co-occurrence on automatic and deliberate evaluations. *European Journal of Social Psychology*, 46(6), 1101–1116. https://doi.org/10.10 02/ejsp.2266
- *Moran, T., Cummins, J., & De Houwer, J. (2022). Examining automatic stereotyping from a propositional perspective: Is automatic stereotyping sensitive to relational and validity information? *Personality and Social Psychology Bulletin*, 48(7), 1024–1038. https://doi.org/10.1177/014616 72211024121
- Moskowitz, G. B. (2014). The implicit volition model: The unconscious nature of goal pursuit. In J. W. Sherman, B. Gawronski, & Y. Trope (Eds.), *Dual-process theories of the social mind* (pp. 400–422). Guilford Press.
- Moskowitz, G. B., Gollwitzer, P. M., Wasel, W., & Schaal, B. (1999). Preconscious control of stereotype activation through chronic egalitarian goals. *Journal of Personality and Social Psychology*, 77(1), 167–184. https://doi.org/10.1037/0022-3514.77.1.167
- Neely, J. H. (1976). Semantic priming and retrieval from lexical memory: Evidence for facilitatory and inhibitory processes. *Memory & Cognition*, 4(5), 648–654. https://doi.org/10.3758/BF03213230
- Nosek, B. A. (2005). Moderators of the relationship between implicit and explicit evaluation. *Journal of Experimental Psychology: General*, 134(4), 565–584. https://doi.org/10.1037/0096-3445.134.4.565
- Papineau, D. (2003). Human minds. In A. O'Hear (Ed.), *Minds and persons* (pp. 159–184). Cambridge University Press. https://doi.org/10.1017/ CBO9780511550294.010
- Payne, B. K., Brown-Iannuzzi, J., Burkley, M., Arbuckle, N. L., Cooley, E., Cameron, C. D., & Lundberg, K. B. (2013). Intention invention and the affect misattribution procedure: Reply to Bar-Anan and Nosek (2012). *Personality and Social Psychology Bulletin*, 39(3), 375–386. https:// doi.org/10.1177/0146167212475225
- Payne, B. K., Cheng, C. M., Govorun, O., & Stewart, B. D. (2005). An inkblot for attitudes: Affect misattribution as implicit measurement. *Journal of Personality and Social Psychology*, 89(3), 277–293. https:// doi.org/10.1037/0022-3514.89.3.277
- Payne, B. K., Hall, D. L., Cameron, C. D., & Bishara, A. J. (2010). A process model of affect misattribution. *Personality and Social Psychology Bulletin*, 36(10), 1397–1408. https://doi.org/10.1177/0146167210383440
- Payne, B. K., & Lundberg, K. (2014). The affect misattribution procedure: Ten years of evidence on reliability, validity, and mechanisms. *Social and Personality Psychology Compass*, 8(12), 672–686. https://doi.org/10 .1111/spc3.12148
- *Perez, W. F., de Almeida, J. H., & de Rose, J. C. (2015). Transformation of meaning through relations of sameness and opposition. *The Psychological Record*, 65(4), 679–689. https://doi.org/10.1007/s40732-015-0138-z
- *Peters, K. R., & Gawronski, B. (2011). Are we puppets on a string? Comparing the impact of contingency and validity on implicit and explicit evaluations. *Personality and Social Psychology Bulletin*, 37(4), 557–569. https://doi.org/10.1177/0146167211400423
- Petty, R. E., & Cacioppo, J. T. (1981). Attitudes and persuasion: Classic and contemporary approaches. Routledge.
- Petty, R. E., & Cacioppo, J. T. (1986). The elaboration likelihood model of persuasion. *Communication and persuasion. Springer Series in Social Psychology*. Springer. https://doi.org/10.1007/978-1-4612-4964-1_1
- Petty, R. E., Tormala, Z. L., Briñol, P., & Jarvis, W. B. G. (2006). Implicit ambivalence from attitude change: An exploration of the PAST model.

Journal of Personality and Social Psychology, 90(1), 21-41. https://doi.org/10.1037/0022-3514.90.1.21

- Petty, R. E., & Wegener, D. T. (1991). Thought systems, argument quality, and persuasion. In R. S. Wyer Jr., & T. K. Srull (Eds.), *Advances in social cognition* (Vol. IV, pp. 147–162). Lawrence Erlbaum.
- Phills, C. E., Hahn, A., & Gawronski, B. (2020). The bidirectional causal relation between implicit stereotypes and implicit prejudice. *Personality* and Social Psychology Bulletin, 46(9), 1318–1330. https://doi.org/10 .1177/0146167219899234
- Quilty-Dunn, J., & Mandelbaum, E. (2018). Inferential transitions. Australasian Journal of Philosophy, 96(3), 532–547. https://doi.org/10.1080/ 00048402.2017.1358754
- Quilty-Dunn, J., Porot, N., & Mandelbaum, E. (2021). The best game in town: The re-emergence of the language of thought hypothesis across the cognitive sciences [Unpublished manuscript]. Department of Philosophy, Washington University in St. Louis.
- Rehder, B. (2009). Causal-based property generalization. *Cognitive Science*, 33(3), 301–344. https://doi.org/10.1111/j.1551-6709.2009.01015.x
- Rehder, B. (2015). The role of functional form in causal-based categorization. Journal of Experimental Psychology: Learning, Memory, and Cognition, 41(3), 670–692. https://doi.org/10.1037/xlm0000048
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64–99). Appleton-Century-Crofts.
- Richeson, J. A., & Ambady, N. (2001). Who's in charge? Effects of situational roles on automatic gender bias. *Sex Roles*, 44(9–10), 493– 512. https://doi.org/10.1023/A:1012242123824
- Rydell, R. J., & McConnell, A. R. (2006). Understanding implicit and explicit attitude change: A systems of reasoning analysis. *Journal of Personality and Social Psychology*, 91(6), 995–1008. https://doi.org/10 .1037/0022-3514.91.6.995
- *Rydell, R. J., McConnell, A. R., Mackie, D. M., & Strain, L. M. (2006). Of two minds: Forming and changing valence-inconsistent implicit and explicit attitudes. *Psychological Science*, *17*(11), 954–958. https:// doi.org/10.1111/j.1467-9280.2006.01811.x
- Rydell, R. J., McConnell, A. R., Strain, L. M., Claypool, H. M., & Hugenberg, K. (2007). Implicit and explicit attitudes respond differently to increasing amounts of counterattitudinal information. *European Journal of Social Psychology*, 37(5), 867–878. https://doi.org/10.1002/ej sp.393
- *Siegel, E., Sigall, H., & Huber, D. E. (2012). The IAT is sensitive to the perceived accuracy of newly learned associations. *European Journal of Social Psychology*, 42(2), 189–199. https://doi.org/10.1002/ejsp.859
- Smaldino, P. E. (2017). Models are stupid, and we need more of them. In R. R. Vallacher, A. Nowak, & S. J. Read (Eds.), *Computational models in social psychology* (pp. 311–331). Routledge. https://doi.org/10.4324/ 9781315173726-14
- Smith, E. R., & DeCoster, J. (2000). Dual-process models in social and cognitive psychology: Conceptual integration and links to underlying memory systems. *Personality and Social Psychology Review*, 4(2), 108– 131. https://doi.org/10.1207/S15327957PSPR0402_01
- Steegen, S., Tuerlinckx, F., Gelman, A., & Vanpaemel, W. (2016). Increasing transparency through a multiverse analysis. *Perspectives on Psychological Science*, 11(5), 702–712. https://doi.org/10.1177/1745691616 658637
- Strack, F., & Deutsch, R. (2004). Reflective and impulsive determinants of social behavior. *Personality and Social Psychology Review*, 8(3), 220–247. https://doi.org/10.1207/s15327957pspr0803_1
- *Sweldens, S., Van Osselaer, S. M. J., & Janiszewski, C. (2010). Evaluative conditioning procedures and the resilience of conditioned brand attitudes. *Journal of Consumer Research*, 37(3), 473–489. https://doi.org/10.1086/ 653656

- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, 331(6022), 1279–1285. https://doi.org/10.1126/science.1192788
- Tolman, E. C. (1949). There is more than one kind of learning. *Psychological Review*, 56(3), 144–155. https://doi.org/10.1037/h0055304
- Uhlmann, E. L., Brescoll, V. L., & Paluck, E. L. (2006). Are members of low status groups perceived as bad, or badly off? Egalitarian negative associations and automatic prejudice. *Journal of Experimental Social Psychology*, 42(4), 491–499. https://doi.org/10.1016/j.jesp.2004.10.003
- Viechtbauer, W. (2010). Conducting meta-analyses in R with the metafor package. *Journal of Statistical Software*, 36(3), 1–48. https://doi.org/10 .18637/jss.v036.i03
- *Whitfield, M., & Jordan, C. H. (2009). Mutual influence of implicit and explicit attitudes. *Journal of Experimental Social Psychology*, 45(4), 748–759. https://doi.org/10.1016/j.jesp.2009.04.006

- Wilson, T. D., Lindsey, S., & Schooler, T. Y. (2000). A model of dual attitudes. *Psychological Review*, 107(1), 101–126. https://doi.org/10 .1037/0033-295X.107.1.101
- *Zanon, R., De Houwer, J., & Gast, A. (2012). Context effects in evaluative conditioning of implicit evaluations. *Learning and Motivation*, 43(3), 155–165. https://doi.org/10.1016/j.lmot.2012.02.003
- *Zanon, R., De Houwer, J., Gast, A., & Smith, C. T. (2014). When does relational information influence evaluative conditioning? *Quarterly Journal of Experimental Psychology*, 67(11), 2105–2122. https://doi.org/10 .1080/17470218.2014.907324

Received January 6, 2022

Revision received September 8, 2022

Accepted September 28, 2022