Journal of Experimental Psychology: General

ISSN: 0096-3445

2026, Vol. 155, No. 5, 1193–1217
https://doi.org/10.1037/xge0001914

Learning Affect Norms: Implications for Predictions, Experiences, and Social Judgments

Omri Maor[1], Yael Millgram[2], Amit Goldenberg[3, 4, 5], William A. Cunningham[6], and Maya Tamir[1]

[1] Department of Psychology, The Hebrew University of Jerusalem
[2] School of Psychological Sciences, Tel Aviv University
[3] Department of Psychology, Harvard University
[4] Harvard Business School, Harvard University
[5] Digital, Data and Design Institute, Harvard University
[6] Department of Psychology, University of Toronto

Exposure to feelings of another can shape how we feel and think in that context. We propose that exposure to the feelings of multiple others carries broader implications that extend beyond the exposure context. We argue that people identify patterns of affective reactions in the social group, which leads them to infer affect norms. These norms, in turn, serve as benchmarks for subsequent social judgments. In four studies ($N = 418$), three of which were preregistered, we exposed participants to sequential affective reactions of multiple others and manipulated their intensity. We then tested the effects of such exposure on predictions of others' feelings, one's own feelings, and subsequent social judgments. Across studies, we show that people predict feelings of group members based on the norm they were exposed to, that they assimilate their own reactions to the norm, and that they subsequently judge social targets that deviate from the norm more negatively. Our findings demonstrate how exposure to feelings of multiple others enables people to learn affect norms, bridging research on emotions in dyads, groups, and cultures.

Public Significance Statement

With the rise of social media, our exposure to feelings of multiple others is greater than ever. This investigation demonstrates that such exposure can lead people to infer which feelings are common and appropriate. Such inferences can shape how people feel, what they expect others to feel, and how they judge others whose feelings deviate from the norm. This work shows that exposure to feelings of multiple others may carry emotional and social implications beyond the context of exposure.

Keywords: emotions, norms, social judgment, affective learning, culture

Supplemental materials: https://doi.org/10.1037/xge0001914.supp

Throughout human history, from sitting together around campfires to interactions on social media, exposure to feelings of others has shaped how we feel and think (Clément & Dukes, 2017; Goldenberg et al., 2020; Mesquita, 2010; Parkinson, 2011). Research on emotions in dyadic interactions has demonstrated that exposure to an emotional reaction of a person to a stimulus can shape how we appraise, what we express toward, and how we feel about that stimulus. However, exposure to feelings of multiple others may have additional effects that go beyond the context of exposure. For instance, members of social groups learn from interactions with the group which feelings

Jessica D. Payne served as action editor.

Omri Maor https://orcid.org/0000-0002-8576-7352
Yael Millgram https://orcid.org/0000-0001-9665-222X
Amit Goldenberg https://orcid.org/0000-0002-9417-5682
William A. Cunningham https://orcid.org/0000-0002-7063-5859
Maya Tamir https://orcid.org/0000-0002-2675-8042

Data, code, and materials are available on the Open Science Framework at https://osf.io/uaq79/?view_only=1f54ead24bab453ca4053c5fb0495a9f. Data from Study 1 were presented by Omri Maor in a poster at the 2022 annual conference of the Society for Affective Science. This research was funded by the Israel Science Foundation (Grant 2281/20) and supported by the Artery Chair in Personality Studies Endowed by Goldberg, Geller, and Luria to Maya Tamir.

Omri Maor played a lead role in formal analysis and visualization, a supporting role in investigation, and an equal role in writing–original draft. Yael Millgram played a supporting role in data curation, investigation, and writing–review and editing and an equal role in methodology. Amit Goldenberg played a supporting role in conceptualization, formal analysis, methodology, and writing–review and editing. William A. Cunningham played a supporting role in conceptualization and writing–review and editing and an equal role in methodology. Maya Tamir played a lead role in conceptualization, funding acquisition, investigation, project administration, resources, supervision, and writing–review and editing, a supporting role in data curation, and an equal role in methodology and writing–original draft.

Correspondence concerning this article should be addressed to Maya Tamir, Department of Psychology, The Hebrew University of Jerusalem, Mount Scopus, Jerusalem 91905, Israel. Email: tamirm@mail.huji.ac.il

are appropriate to feel or express in future situations (Eid & Diener, 2001; Van Kleef, 2009). Such affect norms are critical for social functioning, yet we know relatively little about how they develop. In this investigation, we test whether people can learn affect norms from exposure to feelings of multiple others. We propose that exposure to affective reactions of others leads observers to form mental representations of feelings in the population, and these mental representations then guide predictions, affective reactions, and subsequent social judgments.

From Exposure to Others' Feelings to Affect Norms

Exposure to feelings of another person can have various effects that are specific to the context of exposure (Clément & Dukes, 2017). For instance, social contagion can occur through three mechanisms. First, in mimicry, emotional expression activates synchronous behavior on the part of the perceiver, which, in turn, activates affective processes (Hatfield et al., 1994; Hess & Fischer, 2014; Parkinson, 2011). Second, in category activation, exposure to emotional expressions primes an emotion category, which, in turn, leads to activation of specific emotional processes (Niedenthal et al., 2009; Peters & Kashima, 2015). Such activation does not necessarily involve behavioral copying. Third, in social appraisal, individuals use the emotions of others as information about the value of stimuli in the environment, which can then guide their own appraisals, leading to similar emotional experiences in the same situation (Clément & Dukes, 2017; Manstead & Fischer, 2001). These mechanisms are tied to the context of exposure, informing the observer how to behave in it, what to feel in it, and how to evaluate it.

People can learn information from exposure to feelings of others that carry implications beyond the context of exposure (Smith & Mackie, 2016). For instance, members of the same culture cultivate shared beliefs about which feelings are appropriate to feel in certain contexts (Eid & Diener, 2001; Hochschild, 1983). Such beliefs are descriptive when they capture what exists in a population and prescriptive when they serve as reference points that signal what should exist (Cialdini et al., 1990). These beliefs also shape subsequent affective reactions and direct people's judgments and behavior (Thoits, 2004). For instance, when individuals believe feelings of members of their group are inappropriate, they may develop negative feelings toward them and change their own affective reactions to differentiate themselves from their group (Goldenberg et al., 2014).

Affect norms vary across groups (Eid & Diener, 2001) and critically inform group processes (Delvaux et al., 2015; Thoits, 2004; Vishkin et al., 2023), but how do they develop? Much like people draw inferences about the likelihood and appropriateness of behavior from exposure to behaviors of multiple others in their group, we propose that people can draw inferences about the likelihood and appropriateness of feelings from exposure to feelings of multiple others in their group.

Inferring Affect Norms From Exposure to Others

Research on social norms suggests that when we are exposed to behaviors or attitudes of multiple others, particularly members of our group, we form representations of the distribution of such behaviors or attitudes in the group and subsequently use them as points of reference that can guide our expectations, behaviors, and social judgments (Biernat & Eidelman, 2007; Miller & Prentice, 1996, 2016). Norms are abstract evaluative inferences that go beyond the specific context of learning and can be incrementally updated over sequential exposures to others, particularly similar others or members of one's group (Prentice & Miller, 1993).

Once formed, norms can be descriptive and help people predict others' behaviors in novel situations (e.g., "Will my neighbor recycle?"; Kahneman & Miller, 1986; Prentice & Miller, 1993). Descriptive norms can also be (but are not automatically) prescriptive, signaling what is desirable or undesirable (Miller & Prentice, 1996). Since prescriptive norms guide us on what is right and wrong in specific contexts, people tend to assimilate their behavior to the norm (Cialdini & Goldstein, 2004). Prescriptive norms can also serve as evaluative criteria for judgments of other people, such that those who deviate from the norm may be judged less favorably (Higgins & Lurie, 1983; Marques et al., 1998; Mussweiler, 2003).

If people infer behavioral norms from exposure to behaviors of others, they might similarly infer both descriptive and prescriptive affect norms from exposure to feelings of others. Accordingly, we define affect norms as shared mentally represented summaries of feelings in a group that can guide predictions in novel situations, shape one's own affective experiences, and serve as benchmarks in social judgments. Like other types of norms (Cialdini & Goldstein, 2004; Miller & Prentice, 1996), and beyond the effects of mimicry, contagion, or social referencing, affect norms should guide behavior, experience, and cognition, in reference to stimuli beyond those pertaining to the learning context.

Based on the above analysis, we predicted that people who are exposed to feelings of multiple others would be able to infer an affect norm, either with or without awareness, representing feelings in the population and their appropriateness. We tested three hypotheses. First, given that exposure to behaviors of multiple others can lead people to infer a norm and expect others to behave in a manner that is consistent with the norm (Bicchieri, 2006; Cialdini et al., 1990; Fehr & Schurtenberger, 2018), we hypothesized that people who are exposed to feelings of others will infer a norm and expect novel others to have affective reactions that are consistent with the norm (Hypothesis 1). Support for this hypothesis would indicate that people can infer descriptive affect norms. Second, given that people tend to behave in ways that are consistent with the normative behavior in their group (Cialdini et al., 1990; Horne & Mollborn, 2020; Sinclair & Agerström, 2023), we hypothesized that people who are exposed to feelings of others will assimilate their own responses to the norm (Hypothesis 2). Third, given that people tend to judge those whose behavior deviates from the norm less favorably (e.g., Fehr & Schurtenberger, 2018; Higgins & Lurie, 1983; Mussweiler, 2003), we hypothesized that people who are exposed to feelings of others will judge new others whose feelings deviate from the norm less favorably (Hypothesis 3). Support for this latter hypothesis would require abstract inferences that go beyond the learning context and would indicate that people can infer prescriptive affect norms.

Given that feelings differ from behaviors, affect norms may differ from behavioral norms (e.g., Vishkin et al., 2023). Feelings are unique in that they serve as salient social signals (Van Kleef, 2009), whose value varies as a function of valence and intensity. For example, people pay more attention to more (vs. less) intense emotional expressions, and this asymmetry may render them more likely to adopt norms associated with more (vs. less) intense

emotions. People are also more likely to attend to unpleasant (vs. pleasant) emotions (e.g., Keltner & Haidt, 1999; Parkinson et al., 2005). Thus, it may be more important to learn norms for unpleasant than for pleasant affect or for more versus less intense feelings. Therefore, in this investigation, we tested how affect norms are learned and whether this varies by valence and intensity.

The Present Research

We tested our hypotheses in four studies, using a novel paradigm—the affect norm learning task. In the task, participants were systematically exposed to affective reactions of others, who were presumably similar to them (i.e., ingroup members). The task includes a learning stage and a social judgment stage. In the learning stage (see Figure 1A), we exposed participants sequentially to feelings of multiple others. In the social judgment stage (see Figure 1B), participants judged novel individuals on evaluative dimensions (e.g., competence, likability), based on the affective reactions of those individuals. To test Hypothesis 1, in half of the learning trials, participants estimated the affective reactions of individuals before exposing them to these reactions. To test Hypothesis 2, in half the learning trials, participants rated their own affective reactions. To test Hypothesis 3, participants judged others based on others' norm-congruent and norm-deviant affective responses.

Participants were assigned to conditions that differed in the affective reactions of others (i.e., above, below, or identical to the population mean). In Study 1, participants were exposed to pleasant and unpleasant images. To ensure participants were learning specific norms rather than anchoring their scale use, affective reactions to pleasant images were identical to the group mean (as determined by a pilot study), whereas affective reactions to unpleasant images were either below or above the mean, depending on the condition participants were assigned to. In Study 2, we tested whether people cultivated an abstract representation of feelings or learned how to use the rating scale by using different rating scales in the learning and social judgment stages of the task. In Study 3, we manipulated the affective reactions of others to both pleasant and unpleasant stimuli to test the generalizability of our effects. In Study 4, we tested whether effects on social judgments are independent of the observer's own affective reactions to stimuli.

This paradigm allowed us to examine each of our key hypotheses separately. First, as evidence of learning, we predicted that exposure to feelings of multiple others would lead participants to accurately estimate feelings of new individuals in response to novel stimuli. Second, we predicted that exposure to feelings of multiple others would lead participants to cultivate feelings that are similar to the norm. Such an assimilation effect would be consistent with participants inferring a prescriptive norm, but may also reflect contagion. Third, we predicted that following exposure to feelings of multiple others, participants would be likely to judge novel individuals whose feelings deviate (vs. not) from the estimated response in the population less favorably. Such an effect on social judgment would be consistent with the inference of a prescriptive affect norm.

Study 1

The goal of Study 1 was to test our three key hypotheses by exposing participants to the feelings of multiple others and measuring the participants' estimations of others' feelings, their own feelings, and their social judgments of novel targets. Unlike Studies 2–4, which were preregistered, Study 1 was exploratory.

Method

Participants

One hundred Jewish Israeli students (72 self-identified as women and 28 self-identified as men; $M_{age} = 23.52$, $SD_{age} = 2.34$) participated in the study. No participants were excluded. An a priori power analysis indicated that a sample size of 90 was required to detect a Condition × Affect interaction with a medium effect size ($\eta_p^2 = 0.07, f = 0.27$), with 80% power and $\alpha = .05$. We oversampled by approximately 10% to account for potential exclusions.

Procedure

Participants were told this was a study on emotional intelligence (the complete set of instructions given to participants is provided verbatim in the Supplemental Materials). They were told that they would watch a series of images and rate how they feel in response to them. They were further told that in prior versions of the study, other students from their university watched the same images and rated their reactions to them. Participants were told that in this study, on some trials, they will rate how they feel in response to the images, and on other trials, they will estimate how previous participants felt in response to the images. Participants then completed a baseline affective reactivity task, followed by the affect norm learning task.

First, in the baseline affective reactivity task, participants viewed 10 unpleasant and 10 pleasant images in a random order. An image appeared on the screen for 2,000 ms, and then a continuous valence rating scale (−100 = *very negative*, 100 = *very positive*) appeared below the image, and participants rated how they felt in response to the image. After rating their reaction, participants' rating was displayed again on the screen.

Second, participants completed the affect norm learning task (see Figure 1). In the learning stage of the task, new pleasant and unpleasant images were presented. On each trial, an image appeared, and after 2,000 ms, the valence rating scale (−100 = *very negative*, 100 = *very positive*) appeared below the image. On half the trials, participants rated their own reaction to the image (i.e., "How do you feel?"), and on half of the trials, participants estimated how another participant, identified by name, who allegedly participated in the study, reacted to the image (e.g., "How did *JANE* feel?"). On each trial, a different name appeared. Participants were presented with 100 names in total, of the same gender as the participant. Names were selected from public name registries, as the most common in the participants' cohort. Others were depicted by name only, and no other information was provided.

The learning stage included 200 trials, with trials varying by image (pleasant, unpleasant) and by the source of the affective reaction (rating one's own reactions, estimating reactions of another), presented in a random order. After rating their reaction, participants saw their own rating on the screen. After estimating another person's reaction, participants saw the purported reaction of that person on the screen. When presented with reactions of others to pleasant images, participants were presented with the mean reaction in the sample population to that specific pleasant picture, as assessed

Figure 1

A Visual Depiction of the Affect Norm Learning Task

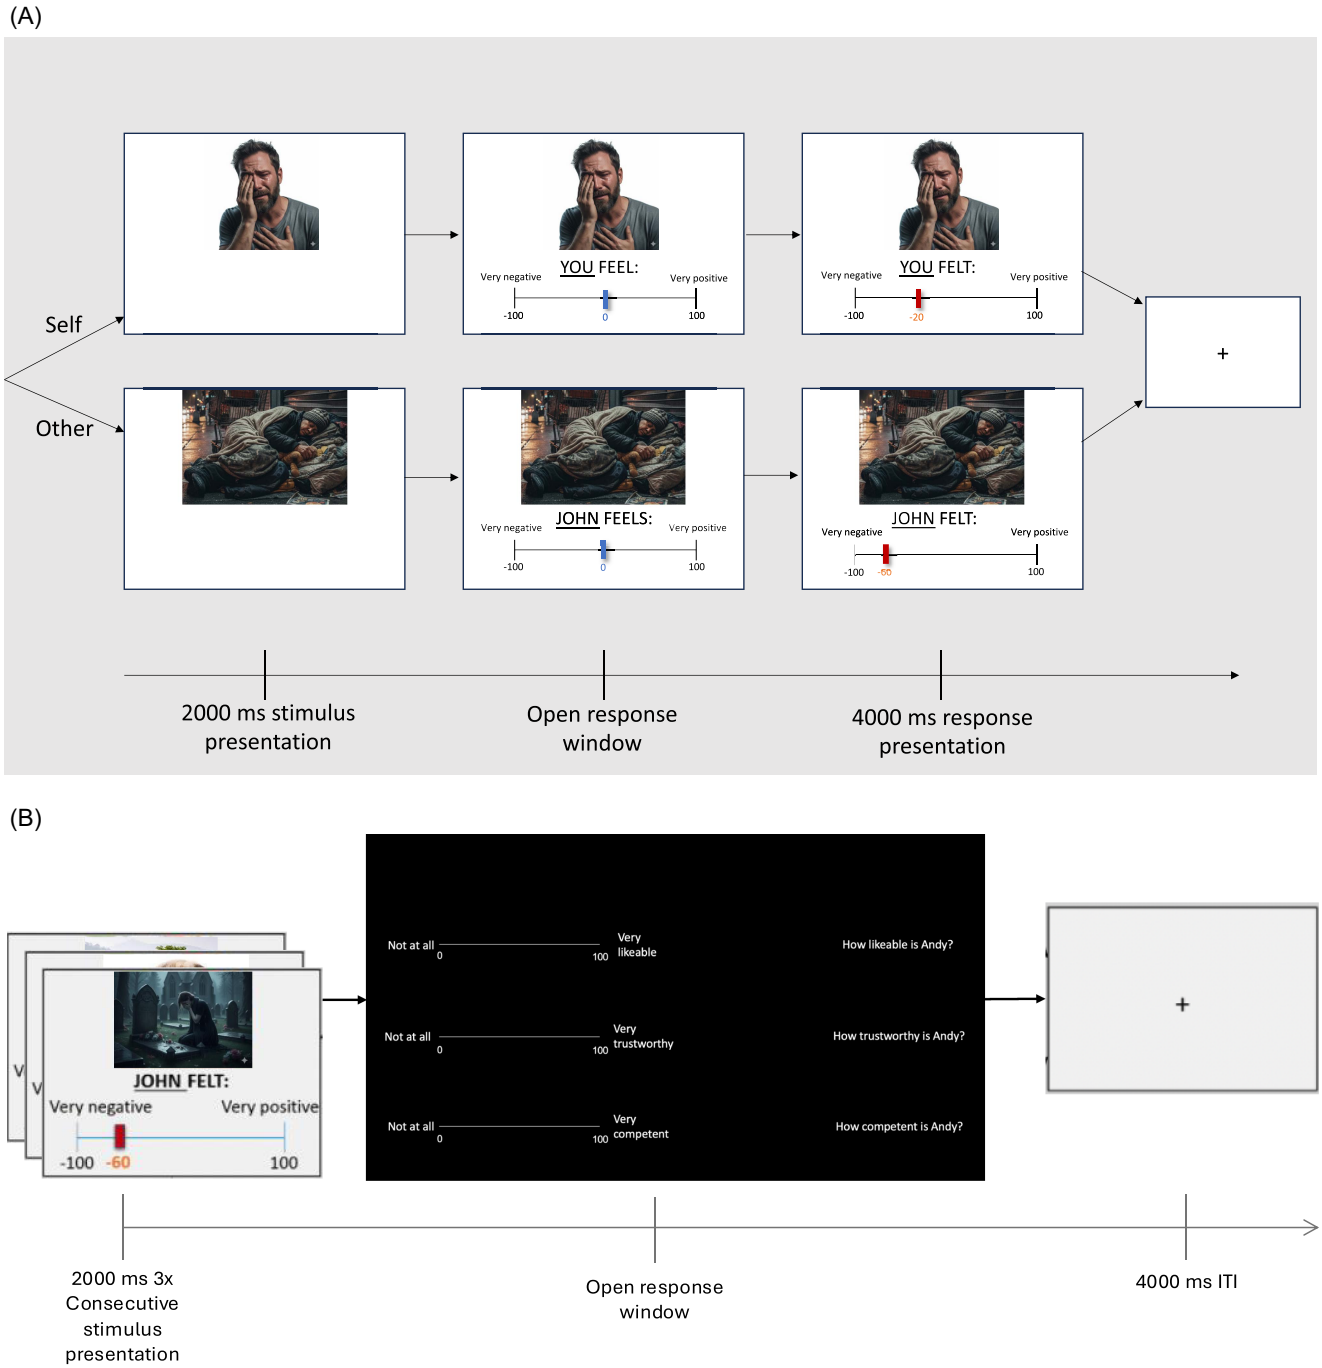

Note. Panel A presents the steps in the learning stage, where participants see an image, indicate how they feel or estimate how another person felt in response to it, and then observe their own rating or the rating of another, which, in effect, varied by condition. Panel B presents the steps in the social judgment stage, where participants see how a target reacted to an image and then judge the target on three evaluative dimensions. ITI = intertrial interval. See the online article for the color version of this figure.

in a pilot test. When presented with reactions of others to unpleasant images, the reactions of others were manipulated, according to the affect norm condition to which the participant was randomly assigned. In one condition, participants saw more extreme negative reactions to unpleasant images (i.e., reactions that were below the population mean). In the second condition, participants saw less extreme negative reactions to unpleasant images (i.e., reactions that were above the population mean). These observed reactions of

others to unpleasant images were determined according to the following formula (Equation 1):

$$R_{it} = \mu_t \pm 0.75\, SD_t + (\mu - \mu_i'), \quad (1)$$

where R_{it} = observed reactions of person i to image t, μ_t = mean reaction in the population to image t, SD_t = standard deviation of the population responses to image t, μ = mean reaction in the population to all unpleasant images, and μ_i' = the participants' mean reaction to unpleasant images in the baseline reactivity task.

When participants were shown more extreme reactions to unpleasant stimuli (reactions below the mean), another's reaction to an unpleasant image was determined by taking the mean response to that image in the population and subtracting .75 of the population *SD* for that image.[1] When participants were shown less extreme reactions to unpleasant stimuli (reactions above the mean), another's reaction to an unpleasant image was determined by taking the mean response to that image in the population and adding .75 of the population *SD* for that image. We adjusted all observed reactions of others to account for participants' own affective reactivity. This was done so that participants' own reactions would not bias their perception of extreme responses, since a more extreme reaction for one individual may not be as extreme for another, depending on their own affective reactivity. This correction also ensured that what is considered "less extreme" or "more extreme" is equivalent across participants and conditions. To adjust for baseline reactivity, we computed a participant's bias score by taking the participant's mean reaction to unpleasant images at the baseline affective reactivity task and subtracting it from the total mean reaction to unpleasant images in the sample population.

Participants then completed the social judgment stage of the task. In this stage, participants observed the affective reactions of three allegedly former participants to novel pleasant and unpleasant images and then judged each of these individuals on competence, likability, and trustworthiness (0 = *not at all*, 100 = *extremely*). Participants observed the reactions of each target to a series of 10 images per target. These reactions were presented one after another on separate slides, with the target's rated affective response (−100 = *very negative*, 100 = *very positive*) appearing underneath each image. Each target responded to different images of the same valence. The reactions of the three target individuals were rigged so that one target responded less extremely to 10 pleasant images (mean reaction to pleasant images −0.75 *SD* + participant bias), one target responded less extremely to 10 unpleasant images (mean reaction to unpleasant images +0.75 *SD* + participant bias), and one target's reactions to 10 pleasant images were consistent with the mean in the sample population for pleasant images (+ participant bias).[2] After all responses of a target were presented, participants evaluated the target on likeability, competence, and trustworthiness on a scale from 0 to 100. A Cronbach's α test indicated high reliability (α = .85), justifying averaging across the three social judgments as our dependent variable.[3] Finally, participants completed several trait measures for exploratory purposes, provided demographic information, and were debriefed.[4]

Materials

Images. We selected images from the International Affective Picture System (Lang et al., 2005). These included 120 unpleasant images with valence ratings under 4 ($M_{valence}$ = 2.93, *SD* = 0.71, on a 1–9 scale) and 130 pleasant images with valence ratings above 6 ($M_{valence}$ = 7.30, *SD* = 0.70, on a 1–9 scale). To validate the affective reactions to the images, participants in a pilot study (N = 57) viewed the images in a random order and rated their affective reactions to them (−100 = *very negative*, 100 = *very positive*). The mean affective reaction for unpleasant images was $M_{valence}$ = −43.11, *SD* = 14.26, and the mean affective reaction for pleasant images was $M_{valence}$ = 37.78, *SD* = 9.48 (see the Supplemental Materials for the full list of images and valence ratings).

Names. We obtained 100 female names and 100 male names from lists of the most popular given Hebrew names in the years 2000–2010 (our participants' cohort), published by the Israeli Central Bureau of Statistics.

Results

Does Exposure to Others' Feelings Shape Estimated Affective Reactions?

To test whether the manipulation affected participants' estimations of the affective reactions of others, we conducted a mixed models analysis, with affect norm condition (exposure to more vs. less extreme reactions), image type (pleasant, unpleasant), and trial number predicting participants' estimations of others' affective reactions, with random intercepts for participant and image. As predicted, we found a significant interaction between trial and condition, b = −0.09, 95% CI [−0.11, −0.07], SE = 0.011, $t(9726)$ = −7.84, $p < .001$, R^2 = .674, indicating that participants' estimations of others' feelings diverged over time, as a function of the affect norm condition (see Figure 2). The bias was in the direction of the observed mean in their condition. Simple slope analyses showed that estimations of reactions to unpleasant images became increasingly less extreme among participants in the less extreme condition, b = 0.063, SE = 0.006, $p < .001$, and more extreme among participants in the more extreme condition, b = −0.026, SE = 0.006, p = .001. For pleasant images, where responses were not manipulated, estimations became less extreme among participants in the more extreme condition, b = −0.032, SE = 0.005, $p < .001$, which may indicate some spillover of norm learning or habituation in the more extreme condition. Estimations did not change significantly among participants in the less extreme condition, b = 0.002, SE = 0.005, p = .778 (see Figure 2).

[1] Based on the distribution of reactions in the population, we selected ±0.75 *SD*, as responses beyond that range could have resulted in floor or ceiling effects.

[2] Due to a programming error, the first 24 participants judged only two targets each that responded to pleasant images. One target responded less extremely to 10 pleasant images, and one target's reactions to 10 pleasant images were consistent with the mean in the population for pleasant images.

[3] Results were largely consistent and did not vary systematically across studies by individual evaluative dimension.

[4] During the debriefing procedure, when asked whether they were suspicious of anything in the study, only 10% expressed suspicion the others were not real. When asked what the study was about, 32% of participants believed the study tested whether and who has the ability to accurately predict others' emotional responses; 18% believed the study tested whether people react similarly to others; 15% offered idiosyncratic accounts (e.g., whether specific names bias expected emotional responses, whether the order of picture presentation impacts emotional responses); 10% believed the study assessed how people differ in emotional reactions; 7% did not know; 1% (one person) correctly identified the purpose of the study.

Figure 2

Participants' Estimated Affective Reactions of Others by Trial Number and Affect Norm Condition (Study 1)

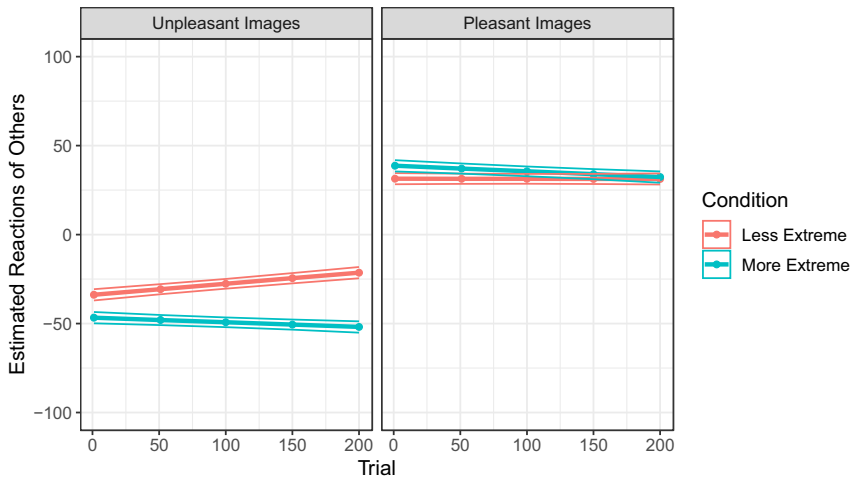

Note. Error bars represent the 95% confidence intervals. See the online article for the color version of this figure.

We further tested whether participants' estimations became closer to the observed responses in their response condition. To do so, we conducted a mixed models analysis testing the effects of affect norm condition (exposure to more vs. less extreme reactions) and trial number. For our predicted variable, we created a difference score between participants' estimations of the affective reactions of others and the actual affective reactions of others, as reflected in the feedback received. The difference score was computed by subtracting the feedback received in the previous trial from the estimation in the current trial, with a random intercept for participant and for image. As predicted, when comparing estimations and actual reactions to unpleasant images, we found a significant main effect for trial, $b = 0.06$, 95% CI [0.05, 0.08], $SE = 0.006$, $t(4895) = 9.80$, $p < .001$, $R^2 = .186$, indicating that participants in both affect norm conditions learned to estimate others' affective reactions over time (see Supplemental Figure S1, left panel). There was no such significant main effect for trial in estimated reactions to pleasant images, $b = 0.00$, 95% CI [−0.01, 0.01], $SE = 0.005$, $t(4820) = -0.36$, $p = .721$, $R^2 = .053$, indicating that manipulating the feedback on unpleasant images affected estimations of reactions to unpleasant images, in particular (see Supplemental Figure S1, right panel).

Participants exposed to more negative reactions of others to unpleasant images subsequently expected others to react more negatively to unpleasant images ($M = -49.01$, $SE = 0.59$), compared to participants who were exposed to less negative reactions of others to unpleasant images ($M = -27.71$, $SE = 0.51$). There was no significant difference between conditions in estimated reactions to pleasant images: less extreme condition, $M = 35.59$, $SE = 0.50$; more extreme condition, $M = 31.30$, $SE = 0.43$.

Does Exposure to Others' Feelings Shape Observers' Own Feelings?

We tested whether the manipulation affected participants' own affective reactions to stimuli. Before testing the learning effects over time, we first tested whether participants' affective reactions differed between the affect norm groups. To do so, we conducted a mixed models analysis, with affect norm condition (exposure to more vs. less extreme reactions) predicting participants' affective reactions, with a random intercept for participant and for image. As predicted, we found a significant interaction effect between condition and image valence, $b = 17.70$, 95% CI [15.64, 19.77], $SE = 1.054$, $t(9717) = 16.79$, $p < .001$, $R^2 = .589$. We then examined the effect of condition by image valence. As predicted, we found a significant effect for condition, $b = -14.87$, 95% CI [−22.37, −7.37], $SE = 3.826$, $t(98) = -3.89$, $p < .001$, $R^2 = .055$, indicating that participants who were exposed to more (vs. less) negative reactions to unpleasant images reacted more negatively to unpleasant images. We found no significant effect for condition on reactions to pleasant images, $b = 2.85$, 95% CI [−4.27, 9.97], $SE = 3.631$, $t(98) = 0.79$, $p = .434$, $R^2 = .003$, indicating that the effect was driven by our manipulation, which targeted unpleasant but not pleasant images (Figure 3B).

As expected, participants exposed to more negative reactions of others to unpleasant images had more negative reactions to unpleasant images ($M = -47.02$, $SE = 0.44$), compared to participants exposed to less negative reactions of others ($M = -28.74$, $SE = 0.39$; Figure 3A). Simple slope analyses showed that affective reactions to unpleasant images became increasingly less extreme among participants in the less extreme condition ($b = 0.052$, $SE = 0.005$, $p < .001$) and increasingly more extreme among participants in the more extreme condition ($b = -0.013$, $SE = 0.005$, $p = .015$; see Supplemental Figure S2). Conditions did not differ in reactions to pleasant images ($M = 34.49$, $SE = 0.390$, in the less extreme condition; $M = 30.92$, $SE = 0.34$, in the more extreme condition).

Does Exposure to Others' Feelings Shape Observers' Social Judgments of Novel Targets?

To assess whether participants' judgments of others varied as a function of previous exposure to feelings of others, we conducted two

Figure 3

Participants' Affective Reactions by Affect Norm Condition and Image Type (Pleasant, Unpleasant; Study 1)

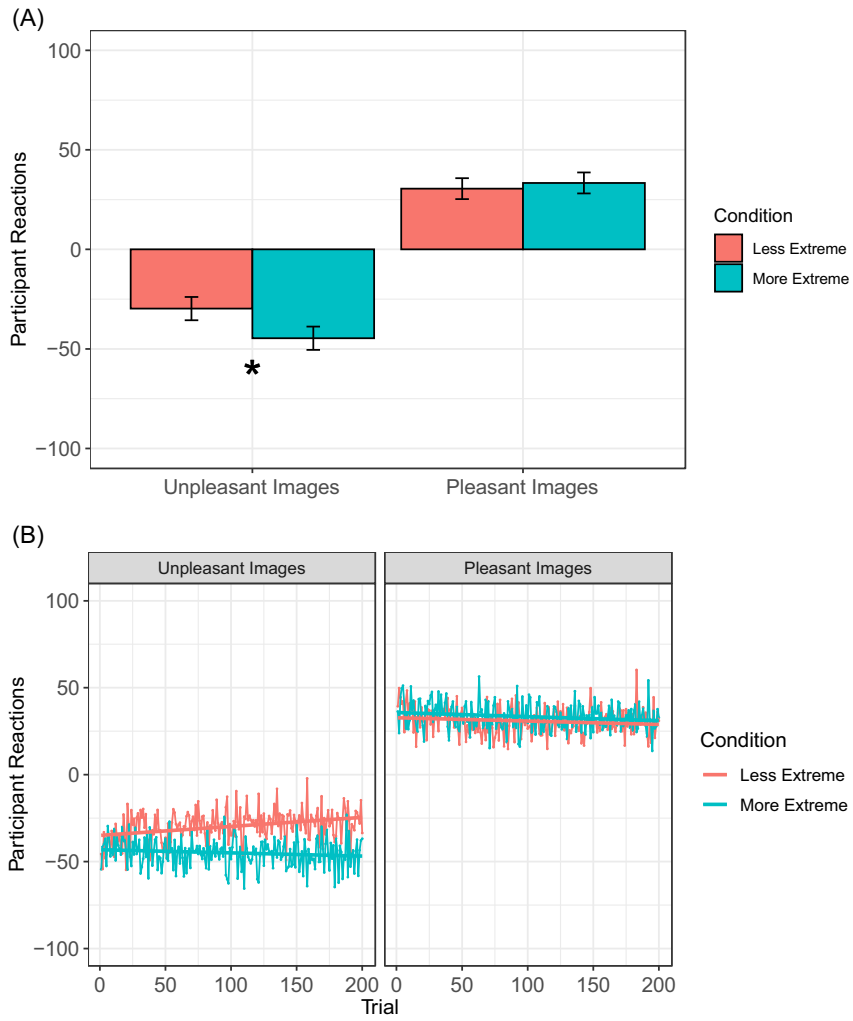

Note. Panel A presents the mean affective reactions for pleasant and unpleasant images by condition. Panel B presents the mean affective reactions by condition for each image type by trial number. Error bars represent the 95% confidence intervals. Participants in both conditions viewed the same reactions to pleasant images. See the online article for the color version of this figure.
* $p < .05$.

linear regression analyses predicting social judgments. First, we tested whether the observed mean in participants' affect norm condition predicted social judgments of targets evaluating unpleasant images.[5] This model included affect norm condition (exposure to more vs. less extreme reactions to unpleasant stimuli) as a between-subjects predictor. For our dependent variable, we used the average of participants' likeability, competence, and trustworthiness judgments. As expected, participants who were exposed to less negative reactions to unpleasant images judged targets reacting less negatively to unpleasant images more favorably ($M = 59.34$, $SE = 3.51$) than did participants who were exposed to more negative reactions ($M = 44.39$, $SE = 3.83$), $b = 14.11$, 95% CI [4.58, 25.32], $SE = 5.204$, $t(73) = 2.87$, $p = .005$, $R^2 = .102$ (Figure 4). As expected, there was no effect of condition (exposure to more vs. less extreme reactions to unpleasant stimuli) on social judgments of targets with less positive reactions to pleasant images ($M = 44.39$, $SE = 3.83$), $b = -6.20$, 95% CI [−14.25, 1.86], $SE = 4.057$, $t(97) = -1.53$, $p = .130$, $R^2 = .023$.

Our second model tested whether participants judged others as a function of deviance from the norm. The data in Study 1 allowed us to compare judgments of targets whose affective reactions to pleasant images were consistent with the true mean in the population to targets whose affective reactions to pleasant images were less extreme (less positive) than the true mean. Our model included target reaction (consistent with vs. less extreme than the population mean) as a

[5] This analysis was conducted on participants who provided judgments of targets reacting to unpleasant images, which excluded the first 24 participants.

Figure 4

Effects of Affect Norm Condition on Social Judgments of Targets With Less Negative Reactions to Unpleasant Images (Study 1)

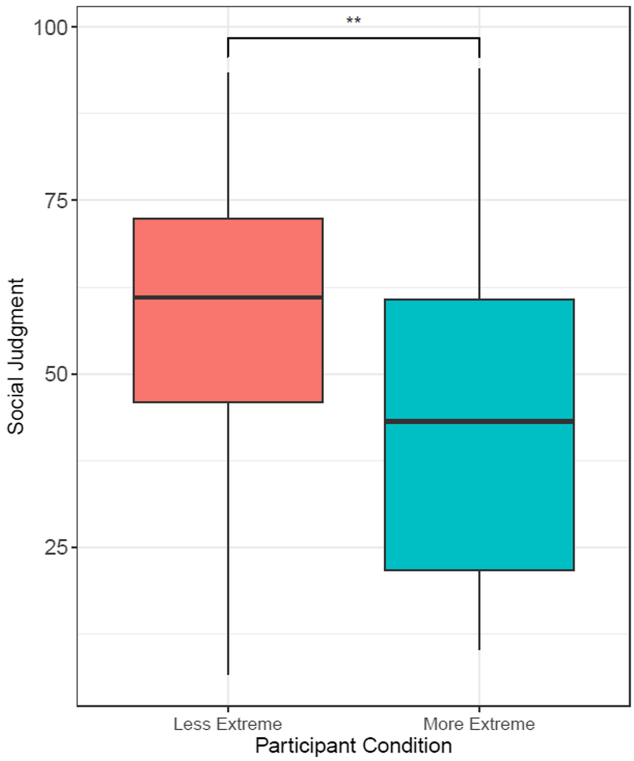

Note. See the online article for the color version of this figure.
** $p < .01$.

within-subject predictor. For our dependent variable, we used the average of participants' likeability, competence, and trustworthiness judgments. As predicted, participants judged targets whose reactions to pleasant images were consistent with the true mean in the population more favorably ($M = 71.46$, $SE = 1.46$) than targets whose reactions to pleasant images were less positive than the true mean in the population ($M = 57.31$, $SE = 2.04$), $b = 14.14$, 95% CI [12.58, 15.72], $SE = 0.798$, $t(494) = 17.73$, $p < .001$, $R^2 = .139$ (Figure 5).

Discussion

Study 1 provided initial support for our predictions. First, exposure to feelings of multiple others led participants to expect feelings of new targets to conform to the estimated response in the sample they have been exposed to. This shows that people infer an expected response in the population and expect new members from that population to assimilate to that response. Second, participants conformed to the estimated response in the population, as their own feelings gradually changed to match the estimated response in the population. Finally, Study 1 offered preliminary support consistent with the hypothesis that individuals use the estimated affective response in the population as a reference point when making subsequent social judgments. When provided with information about the affective reactions of novel targets to novel stimuli, participants tended to judge targets whose affective responses deviated from the estimated response in the population less favorably.

Study 1 had several limitations. First, the same scale was used to indicate others' feelings and to capture participants' own feelings. It is possible, therefore, that participants did not infer the mean in the population, but rather learned to use a specific scale in a consistent manner across rating contexts (i.e., learned how to calibrate responses to the specific rating scale). Another limitation in Study 1 involved the social judgment stage, which failed to include targets whose reactions to unpleasant images were consistent with the mean. We addressed these limitations in Study 2.

Study 2

In Study 2, we sought to replicate the findings of Study 1, while addressing its key limitations. First, one alternative interpretation of the prediction and assimilation effects in Study 1 is that participants were merely learning how to calibrate their use of the scale. To rule out this explanation, in Study 2, whereas others' affective reactions were presented on a horizontal, numeric scale, participants reported their own affective reactions in a separate block, on a vertical, illustrative scale. Second, to address the limitations in the design of the social judgment stage in Study 1, in Study 2, we implemented a 2×2 design in the social judgment stage, such that each participant judged targets whose affective reactions were congruent or incongruent with the mean reaction observed in their condition. Finally, Study 2 was preregistered (https://aspredicted.org/FMR_XSD).[6]

Method

Participants

One hundred and six Jewish Israeli participants (64 self-identified as women, 41 self-identified as men, and one participant did not identify a gender; $M_{age} = 23.84$) took part in the study. This sample size was based on a power analysis, according to which 106 participants are required to detect a medium effect size ($d = .55$) of an interaction in an analysis of variance, with 80% power and $\alpha = .05$. We preregistered this analysis but ultimately decided that multilevel models are more appropriate,[7] but no participants were excluded. One participant terminated her participation after completing the affect norms learning task. We retained this participant in the analyses to be conservative; however, results remained unchanged when this participant was excluded.

Procedure

The procedure was similar to that of Study 1, except for the following modifications. First, we shortened the affect norms learning stage to 100 trials. These included 25 trials in which participants viewed unpleasant images and rated their own affective reaction, 25 trials in which participants viewed pleasant images and rated their own affective reaction, 25 trials in which participants viewed unpleasant images and estimated another participant's affective reaction, and 25 trials in which participants viewed pleasant images and estimated another participant's affective reaction. While we did not have an a priori power analysis supporting the decision

[6] In the Supplemental Materials, we note any discrepancies between the preregistered analyses and the analyses described in the article.

[7] Testing our hypotheses using analyses of variance, as originally preregistered, confirmed our original predictions.

Figure 5

Mean Social Judgments by Targets' Reactions to Pleasant Images (Study 1)

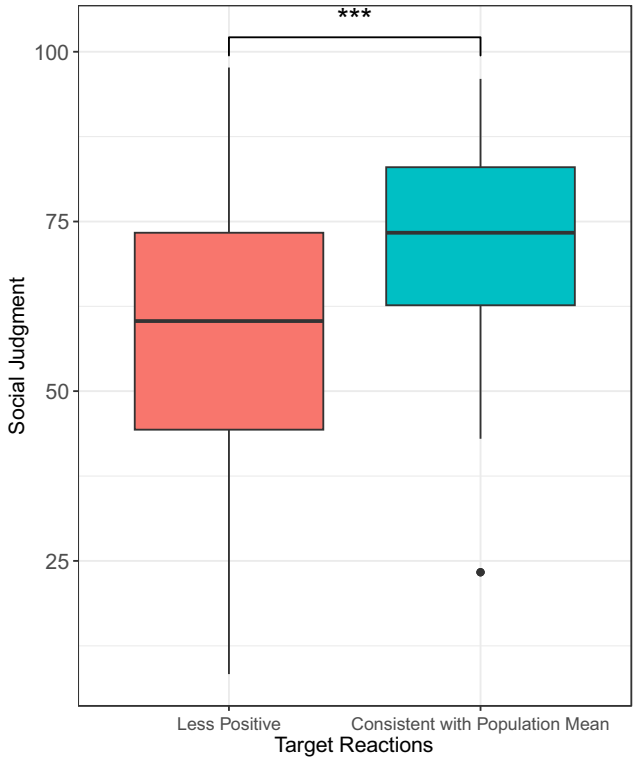

Note. Participants in both conditions viewed the same reactions to pleasant images, which were consistent with the population mean. See the online article for the color version of this figure.
*** $p < .001$.

to shorten the learning stage, an observed power analysis of the effect of norm condition on participants' own reactions in Study 2 showed that 100 trials gave 93.93% power to discover a medium effect size (Green & MacLeod, 2016; Westfall et al., 2014).

Second, after completing the learning stage, participants viewed 10 additional unpleasant images and 10 additional pleasant images. Each image appeared on the screen for 2,000 ms, and then an illustrative rather than a numeric scale appeared. Participants were instructed to use the illustrative scale to indicate their own affective reaction to the image. The illustrative scale was vertical, rather than horizontal, and appeared on the left side of the screen. The anchors of the scale included a smiling face at the top end of the scale, a frowning face at the bottom end of the scale, and a neutral face at the middle of the scale. The scale did not include any numeric values (see Figure 6).

Third, we modified the social judgment stage, such that participants judged eight targets who were allegedly former participants, after viewing five affective reactions of each purported target participant. Participants were shown reactions of two novel individuals (i.e., targets) who reacted more negatively than the population mean to unpleasant images, two targets who reacted less negatively to unpleasant images, two targets who reacted more positively than the mean to pleasant images, and two targets who reacted less positively than the mean to pleasant images. Half of the targets' affective reactions were presented using the numeric scale, and half were presented using the illustrative scale. Participants evaluated each target on likeability, competence, and trustworthiness on a scale from 0 to 100 ($\alpha = .89$). After completing the social judgment stage, participants completed several trait measures for exploratory purposes, provided demographic information, and were debriefed.[8]

Materials

Images. We used a subset of the unpleasant and pleasant images pretested in the pilot study (see the Supplemental Materials for the full list of images used in each study).

Names. We selected the 50 most popular names for females and for males, from the list used in Study 1.

Results

Does Exposure to Others' Feelings Shape Estimated Feelings?

To test whether the manipulation affected estimated feelings of others, we conducted a mixed models analysis as in Study 1. As predicted, we found a significant Trial × Condition interaction, $b = -0.13$, 95% CI [−0.20, −0.07], $SE = 0.035$, $t(5118) = -3.85$, $p < .001$, $R^2 = .673$, indicating that participants' estimations of others' feelings changed over time, as a function of their affect norm condition (see Figure 7). The bias was in the direction of the observed mean in participants' respective condition. Simple slope analyses showed that estimations of reactions to unpleasant images became increasingly less extreme among participants in the less extreme condition, $b = 0.137$, $SE = 0.020$, $p < .001$. Unlike Study 1, we did not find a significant learning trend among participants in the more extreme condition, $b = 0.010$, $SE = 0.020$, $p = .960$. As expected, the slopes for pleasant images were not significant (see Figure 7).

We further tested whether participants' estimations became closer to the observed responses in their respective condition. To do so, we conducted a mixed models analysis as in Study 1. As predicted, regarding unpleasant images, we found a significant main effect for trial, $b = 0.13$, 95% CI [0.09, 0.17], $SE = 0.020$, $t(2517) = 6.35$, $p < .001$, $R^2 = .218$, indicating that participants in both affect norm conditions learned to estimate others' affective reactions to unpleasant images over time (see Supplemental Figure S3, left panel).

Participants exposed to more negative reactions to unpleasant images expected more negative reactions to unpleasant images ($M = -49.84$, $SE = 0.78$), compared to participants who were exposed to less negative reactions ($M = -31.47$, $SE = 0.83$). Participants also differed in the estimations of others' reactions to pleasant images,

[8] Due to a technical error, the debriefing form was recorded for only 61 participants in Study 2. Of those, only 16% suspected others in the study were not real. When asked what the study was about, 21% believed it assessed how people differ in emotional reactions; 15% believed the study assessed how one's feelings are associated with others' feelings, without specifying the direction of this association; 13% believed the study examined whether and who has the ability to accurately predict others' emotional responses; 11% believed the study tested how people's emotional reactions predict their expectations of others' responses; 10% believed the study tested whether people react similarly to others; 10% offered idiosyncratic accounts (e.g., whether specific names impact predicted emotional responses); and 10% did not know. None of the participants correctly guessed the purpose of the study. Also, in this study, we assessed participants' social desirability. Controlling for social desirability did not change the results.

Figure 6

The Illustrative Scale Used in Study 2 (on the Left) and the Numeric Scale Used in Studies 1, 3, and 4 (on the Right)

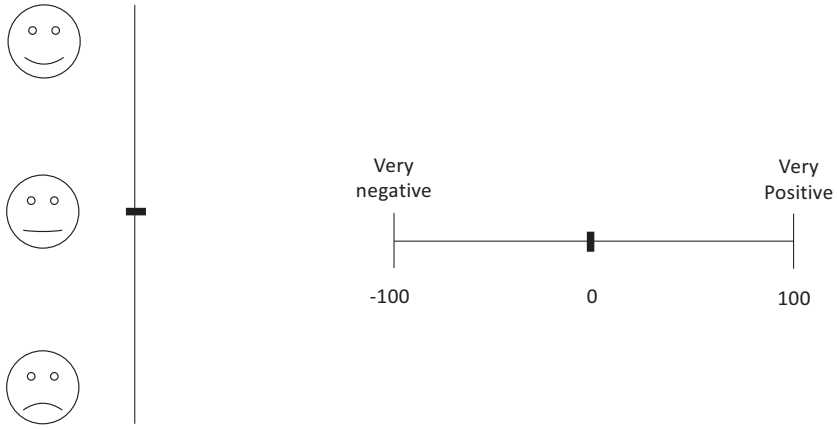

such that participants who observed less extreme reactions to unpleasant images predicted less extreme reactions to pleasant images ($M = 33.18$, $SE = 0.72$) than participants who observed more extreme reactions ($M = 41.61$, $SE = 0.71$). This suggests the norm learned for unpleasant images may have affected some reactions for pleasant images, which were not manipulated. However, estimated reactions of others to pleasant images did not change across trials, $b = -0.02$, 95% CI [−0.05, 0.01], $SE = 0.018$, $t(2509) = -1.19$, $p = .236$, $R^2 = .031$, and so do not appear to be the result of learning from others' reactions to pleasant images (see Supplemental Figure S3, right panel).

Does Exposure to Others' Feelings Shape Observers' Own Feelings?

We tested whether the manipulation affected participants' own reactions to stimuli by conducting a mixed models analysis as in Study 1. As predicted, we found a significant effect for condition, $b = -13.19$, 95% CI [−20.53, −5.84], $SE = 3.744$, $t(104) = -3.52$, $p = .001$, $R^2 = .039$, indicating that participants exposed to more (vs. less) negative feelings of others reported more negative reactions to unpleasant images. We found no significant effect of condition on reactions to pleasant images, $b = 7.67$, 95% CI [−0.06, 15.40], $SE = 3.941$, $t(104) = 1.95$, $p = .054$, $R^2 = .016$ (Figure 8B).

Participants exposed to more negative reactions to unpleasant images reported more negative reactions to the unpleasant images ($M = -48.78$, $SE = 0.57$), compared to participants who were exposed to less negative reactions to unpleasant images ($M = -33.23$, $SE = 0.63$). There was no significant difference across conditions in reactions to pleasant images (Figure 8A). Simple slope analyses showed that affective reactions to unpleasant images became increasingly less extreme among participants in the less extreme

Figure 7

Participants' Estimated Affective Reactions of Others by Trial Number, Affect Norm Condition, and Image Type (Study 2)

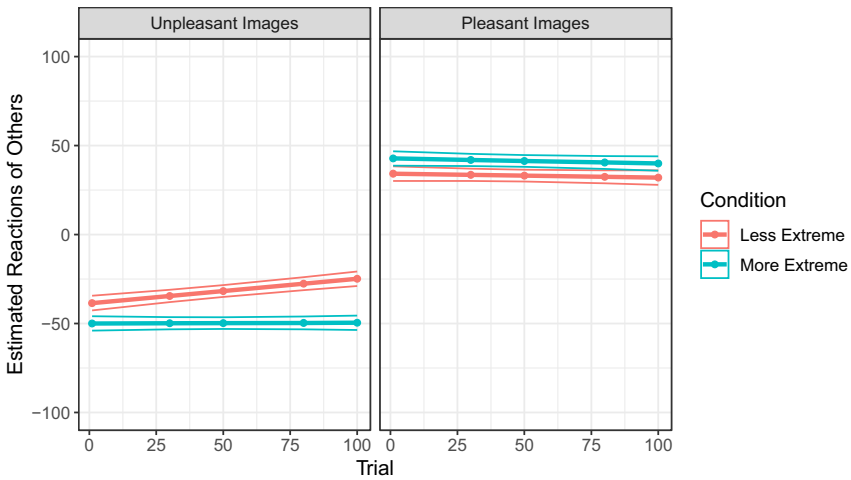

Note. Error bands represent the 95% confidence intervals. Participants in both conditions viewed the same reactions to pleasant images. See the online article for the color version of this figure.

Figure 8

Participants' Affective Reactions by Affect Norm Condition and Image Type (Pleasant, Unpleasant; Study 2)

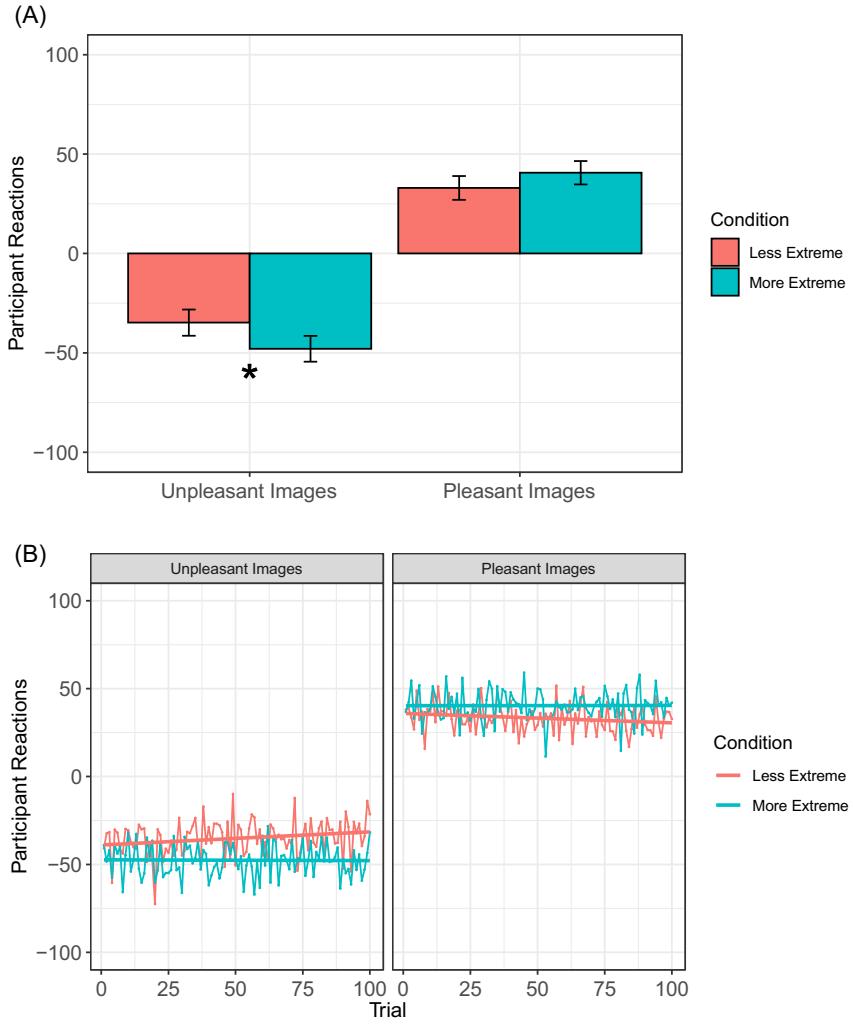

Note. Panel A presents the mean affective reactions for pleasant and unpleasant images per condition. Panel B presents the mean affective reactions per condition for each image type by trial number. Error bars represent the 95% confidence intervals. Participants in both conditions viewed the same reactions to pleasant images. See the online article for the color version of this figure.
* $p < .05$.

condition, $b = 0.084$, $SE = 0.016$, $p < .001$. However, unlike in Study 1, we did not find a significant learning or assimilation trend among participants in the more extreme condition, $b = -0.010$, $SE = 0.016$, $p = .535$ (see Supplemental Figure S4). We found a spillover effect for pleasant images, whereby participants' reactions in the less extreme condition became increasingly less extreme, $b = -0.061$, $SE = 0.015$, $p < .001$.

Does Exposure to Others' Feelings Shape Observers' Social Judgments of Novel Targets?

To test whether participants judged others whose affective reactions were incongruent with the observed mean reaction less favorably, we ran a linear regression model with judgment condition (target of evaluation being congruent or incongruent with the mean reaction observed by the participant) as a between-subjects predictor, predicting social judgments.[9] As we expected, participants judged targets whose reactions to unpleasant images were incongruent (vs. congruent) with the observed mean in their condition less favorably, $b = -11.01$, 95% CI [−17.22, −4.80], $SE = 3.152$, $t(208) = -3.49$, $p < .001$, $R^2 = .055$. There was no significant difference in judgments of targets reacting to

[9] Adding scale type (numeric, illustrative) as a within-subject factor in the linear mixed models analysis did not yield a significant interaction, as expected. Adding this predictor did not significantly change the social judgment results, so for simplicity, we left it out of the model.

pleasant images, $b = -5.17$, 95% CI [−11.02, 0.69], $SE = 2.969$, $t(208) = -1.74$, $p = .083$, $R^2 = .014$. To further explore the interaction, we conducted post hoc pairwise comparisons with Bonferroni p-adjustments. There was a significant effect of participant condition when evaluating targets whose reactions to unpleasant images were less extreme than the population mean. Participants exposed to less extreme reactions judged such targets more favorably ($M = 55.12$, $SE = 2.77$) than participants exposed to more extreme reactions ($M = 40.17$, $SE = 3.17$), $p < .001$. However, there was no significant effect of condition on judgments of targets whose reactions to unpleasant images were more extreme than the population mean, $p = .114$. There were no differences between conditions in judgments of targets who reacted to pleasant images (Figure 9).

Can Anchoring Explain the Assimilation Effects in Observers' Affective Reactions?

It is possible that as people saw ratings of others' affective reactions, they learned how to use the rating scale and calibrated ratings of others' responses as well as their own ratings to those anchors. If so, effects on rated feelings should disappear when using an entirely different scale. To test whether effects persisted when using a different scale, we ran a mixed models analysis, with affect norm condition (exposure to more vs. less extreme reactions) as the predictor, predicting participants' mean affective reactions in the subsequent block, where participants observed new images and rated their reactions using a visually and graphically different scale. As predicted, we found our expected assimilation effects. There was a significant effect for condition, $b = -14.79$, 95% CI [−22.21, −7.36], $SE = 3.783$, $t(103) = -3.91$, $p < .001$, $R^2 = .053$, indicating that participants exposed to less negative reactions to unpleasant images in the learning phase subsequently reacted less negatively to novel unpleasant images ($M = -33.97$, $SE = 1.41$), compared to participants who were exposed to more negative reactions ($M = -48.76$, $SE = 1.29$). We found no significant effect for condition on reactions to pleasant images, $b = 6.36$, 95% CI [−1.27, 14.00], $SE = 3.891$, $t(103) = 1.64$, $p = .105$, $R^2 = .012$ (Figure 10).

We also tested whether participants' affective reactions differed between the numeric and illustrative scales by conducting a mixed models analysis with condition, image valence, and scale type (numeric vs. illustrative), predicting affective reactions. We added a random intercept for participant and for image. As predicted, we found no main effect for scale type, $b = -2.52$, 95% CI [−10.54, 5.51], $SE = 4.092$, $t(123) = -0.62$, $p = .540$, $R^2 = .623$ (Figure 10), nor did we find interaction effects with scale type.

Discussion

By using an illustrative, vertical rating scale, Study 2 helped strengthen the conclusion that participants learned about the mean affective reaction in a group and assimilated their own affective reactions to it. Such evidence of assimilation was found even when participants rated their own affective reactions on a scale that was different than the one used to indicate others' reactions. Furthermore, in Study 2, using a more comprehensive social judgment task, we demonstrated that participants who learned an affect norm subsequently judged people whose reaction was congruent (vs. incongruent) with this norm more favorably.

Unexpectedly, whereas in Study 1, we found evidence for learning among participants who were exposed to less and more extreme unpleasant reactions, in Study 2, we found evidence for learning only among participants who were exposed to less extreme unpleasant reactions. This may have been because Study

Figure 9

Mean Social Judgments of Targets, Based on Targets' Affective Reactions to Pleasant or Unpleasant Images, as a Function of Affect Norm Condition (Study 2)

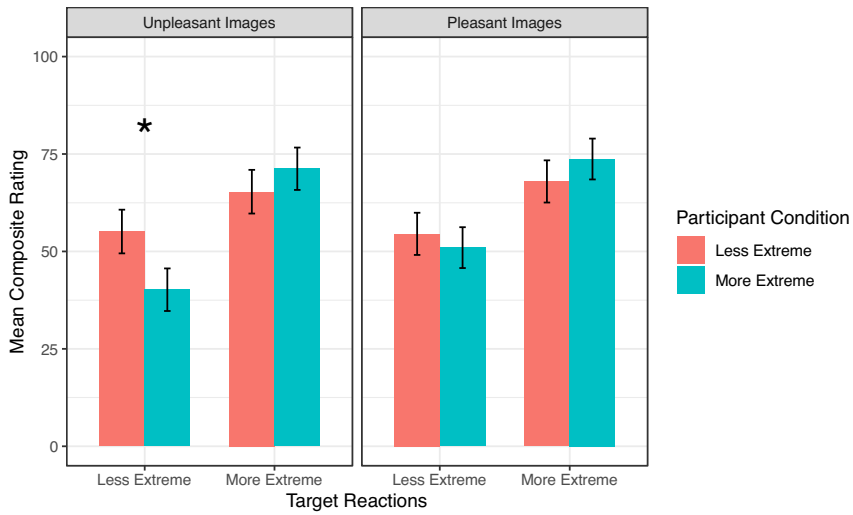

Note. Error bars represent the 95% confidence intervals. See the online article for the color version of this figure.
* $p < .05$.

Figure 10

Participants' Affective Reactions by Affect Norm Condition, Image Type (Pleasant, Unpleasant), and Scale Type (Numeric, Illustrative; Study 2)

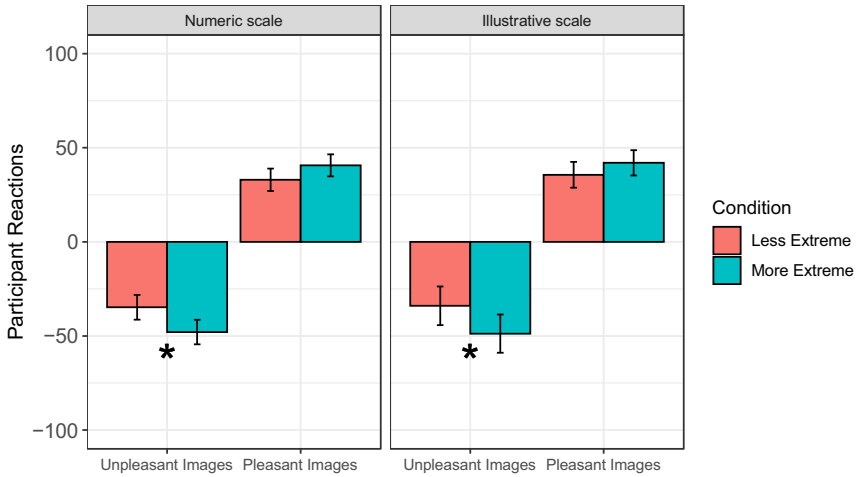

Note. Error bars represent the 95% confidence intervals. Participants in both conditions viewed the same reactions to pleasant images. See the online article for the color version of this figure.

* $p < .05$.

2 included fewer learning trials than Study 1 did, and learning more extreme affect norms may require more exposure. In Study 3, we sought to test whether learning affect norms extends to both less and more extreme reactions of others to both unpleasant and pleasant stimuli.

Study 3

In Studies 1 and 2, we manipulated the emotional reactions of others to unpleasant images and showed the same mean reactions to pleasant images across conditions. This leaves open the possibility that people learn unpleasant, but not necessarily pleasant, affective reactions of multiple others. In Study 3, therefore, we sought to test whether people infer the mean affective response of multiple others to both unpleasant and pleasant stimuli. To this end, in Study 3, we used a 2 (affect norm: more or less extreme reactions) × 2 (affect type manipulated: unpleasant vs. pleasant) design. Half of the participants observed either more or less extreme reactions to unpleasant images and mean reactions to pleasant images (as in Studies 1 and 2), and half of the participants observed either more or less extreme reactions to pleasant images and mean reactions to unpleasant images. We expected participants to infer both unpleasant and pleasant affect norms.

Method

Study 3 was preregistered (https://aspredicted.org/JNH_DTR).[10]

Participants

One hundred and forty Jewish Israeli participants (93 self-identified as women, 46 self-identified as men, and one participant did not identify a gender; $M_{age} = 23.41$) participated in the study. No participants were excluded.

Procedure

The procedure was identical to that of Study 1, except for the following changes. First, participants were randomly assigned to one of four conditions. Half of the participants were exposed to mean affective reactions to pleasant images and either more or less extreme reactions to unpleasant images. Half of the participants were exposed to mean affective reactions to unpleasant images and either more or less extreme reactions to pleasant images. This resulted in an Affect Norm (more vs. less extreme reactions) × Manipulated Affect (pleasant vs. unpleasant) design.

Second, the affect norms learning stage included 160 trials. There were 40 trials in which participants viewed unpleasant images and estimated another participant's affective reaction, 40 trials in which participants viewed pleasant images and rated their own affective reaction, 40 trials in which participants viewed unpleasant images and estimated another participant's affective reaction, and 40 trials in which participants viewed pleasant images and estimated another participant's affective reaction.

Third, in the social judgment stage, participants judged four targets based on five affective reactions per target. Participants judged one target who reacted more negatively to unpleasant images, one target who reacted less negatively to unpleasant images, one target who reacted more positively to pleasant images, and one target who reacted less positively to pleasant images. Participants' affective reactions were presented using a numeric

[10] See the Supplemental Materials for any discrepancies between the preregistered analyses and the analyses described in the article.

scale (−100 = *extremely negative*, 100 = *extremely positive*). Participants evaluated each target on likeability, competence, and trustworthiness on a scale from 0 to 100 (α = .90). At the end of the study, participants completed several trait measures for exploratory purposes, provided demographic information, and were debriefed.[11]

Materials

Images. We used a subset of the unpleasant and pleasant images pretested in the pilot study (see the Supplemental Materials for the full list of images used in each study).

Names. We used a subset of the names we used in Study 1. We selected the 80 most popular names for females and for males.

Results

Does Exposure to Others' Feelings Shape Estimated Feelings?

To test whether the manipulation affected participants' estimations of the affective reactions of others, we conducted a mixed models analysis, as in Studies 1 and 2. As predicted, we found a significant interaction between trial and condition, for unpleasant images, $b = -0.12$, 95% CI [−0.16, −0.08], $SE = 0.020$, $t(5408) = -5.86$, $p < .001$, $R^2 = .702$, as well as for pleasant images, $b = 0.04$, 95% CI [0.00, 0.07], $SE = 0.018$, $t(5413) = 2.18$, $p = .029$, $R^2 = .709$, indicating that participants' estimations of others' feelings diverged over time as a function of affect norm condition (see Figure 11). The bias was in the direction of the observed mean in the participants' respective condition. We then conducted separate analyses for the conditions where we manipulated the affective reactions to unpleasant images and for the conditions where we manipulated the affective reactions to pleasant images.

Manipulated Reactions to Unpleasant Images. Simple slope analyses showed that among participants for whom unpleasant images were manipulated, estimated reactions to unpleasant images became increasingly less extreme among participants in the less extreme condition, $b = 0.074$, $SE = 0.011$, $p < .001$, and more extreme among participants in the more extreme condition, $b = -0.043$, $SE = 0.011$, $p < .001$ (see Figure 11).

Manipulated Reactions to Pleasant Images. Simple slope analyses showed that among participants for whom pleasant images were manipulated, estimations of reactions to pleasant images became increasingly less extreme among participants in the less extreme condition, $b = -0.085$, $SE = 0.008$, $p < .001$, and more extreme among participants in the more extreme condition, $b = 0.056$, $SE = 0.008$, $p < .001$ (see Figure 11).

Do Estimated Affective Reactions Approximate Observed Reactions?

We further tested whether participants' estimations became closer to the observed responses in their condition. To do so, we conducted a mixed models analysis as in Studies 1 and 2. To contrast the effects in each cell of Study 3's 2 × 2 design, we filtered the model by image type (pleasant, unpleasant) and manipulated affect (pleasant, unpleasant), yielding four results.

Manipulating Reactions to Unpleasant Images. First, we compared estimations and feedback on trials with unpleasant images for participants who viewed manipulated reactions to unpleasant images. As predicted, we found a significant main effect for trial, $b = 0.08$, 95% CI [0.06, 0.10], $SE = 0.011$, $t(2690) = 7.12$, $p < .001$, $R^2 = .260$, indicating that participants in both affect norm conditions learned to estimate others' affective reactions over time (see Supplemental Figure S5, bottom left panel). Second, we compared estimations and feedback on trials with unpleasant images for participants who viewed manipulated reactions to pleasant images. There was no such significant main effect for trial, $b = 0.00$, 95% CI [−0.02, 0.02], $SE = 0.010$, $t(2681) = -0.28$, $p = .782$, $R^2 = .051$, indicating that manipulating the feedback on trials with unpleasant images affected estimations with regard to the unpleasant images, in particular, whereas estimations for pleasant images, which were not manipulated for this group of participants, were not affected (see Supplemental Figure S5, top left panel).

Manipulated Reactions to Pleasant Images. Next, we compared estimations and feedback on trials with pleasant images for participants who viewed manipulated reactions to pleasant images. As predicted, we found a significant main effect for trial, $b = -0.08$, 95% CI [−0.10, −0.07], $SE = 0.008$, $t(2688) = -9.96$, $p < .001$, $R^2 = .276$, indicating that participants in both affect norm conditions learned to estimate others' affective reactions over time (see Supplemental Figure S5, top right panel). We then compared estimations and feedback on trials with pleasant images for participants who viewed manipulated reactions to unpleasant images. There was no such significant main effect for trial, $b = -0.02$, 95% CI [−0.03, 0.00], $SE = 0.010$, $t(2679) = -1.64$, $p = .101$, $R^2 = .062$, indicating that manipulating the feedback on trials with pleasant images affected estimations with regard to the pleasant images, in particular, whereas estimations for unpleasant images, which were not manipulated for this group of participants, were not affected (see Supplemental Figure S5, bottom right panel).

Participants exposed to more negative reactions of others to unpleasant images subsequently expected others to react more negatively to unpleasant images ($M = -54.37$, $SE = 0.82$), compared to participants who were exposed to less negative reactions of others to unpleasant images ($M = -33.21$, $SE = 0.75$) and compared to participants who were exposed to more positive reactions of others to pleasant images ($M = -41.01$, $SE = 0.68$). Participants exposed to more positive reactions to pleasant

[11] In the debriefing in Study 3, to be more conservative, we explicitly asked people whether they suspected that "others" in the study were not real. As might be expected, this direct question raised participants' suspicions, but still only 44% of participants said they suspected others were not real. When asked what the study was about, 27% of participants believed the study tested whether and who has the ability to accurately predict others' emotional responses; 22% believed the study assessed how people differ in emotional reactions; 16% believed the study tested whether people react similarly to others; 14% believed the study assessed how one's feelings are associated to the feelings of others, without specifying the direction of this association; 9% believed the study tested how people's emotional reactions predict their expectations of others' responses; 7% offered idiosyncratic accounts; 5% did not know. None of the participants correctly identified the purpose of the study. Also, in this study, we assessed participants' social desirability. Controlling for social desirability did not change the results.

Figure 11

Participants' Estimated Affective Reactions of Others by Trial Number, Affect Norm Condition, Image Type (Pleasant, Unpleasant), and Manipulated Affect (Pleasant vs. Unpleasant; Study 3)

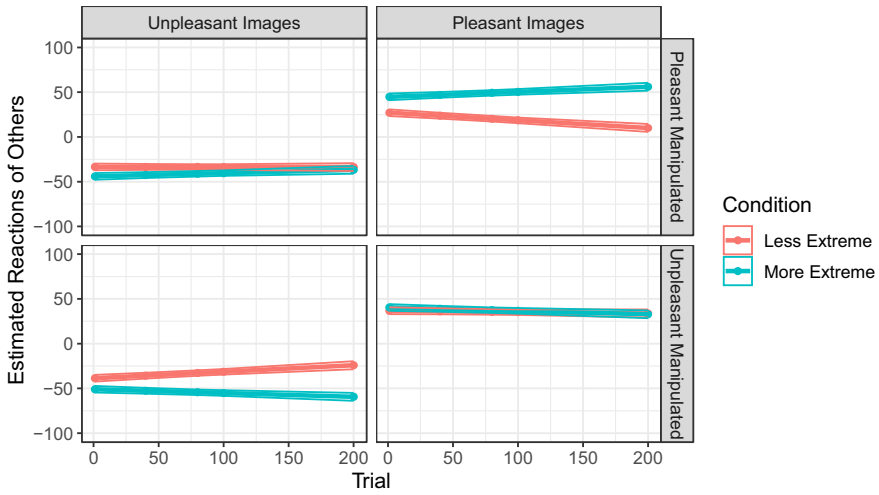

Note. Error bands represent the 95% confidence intervals. See the online article for the color version of this figure.

images subsequently expected others to react more positively to pleasant images ($M = 49.20$, $SE = 0.71$), compared to participants who were exposed to less positive reactions of others to pleasant images ($M = 20.24$, $SE = 0.54$) and compared to participants who were exposed to more negative reactions to unpleasant images ($M = 37.38$, $SE = 0.64$).

Does Exposure to Others' Feelings Shape Observers' Own Feelings?

We tested whether the manipulation affected participants' own reactions to stimuli, conducting a mixed models analysis as in Studies 1 and 2, adjusted for Study 3's 2 × 2 design.

Manipulated Reactions to Unpleasant Images. As predicted, for participants who observed manipulated reactions to unpleasant images, we found a significant effect for condition, $b = -15.82$, 95% CI [−26.12, −5.53], $SE = 5.250$, $t(68) = -3.01$, $p = .004$, $R^2 = .055$, indicating that participants, who observed affective reactions which were more (vs. less) negative than the mean in the population, reported more negative reactions to unpleasant images (Figure 12A). We found no significant effect for condition on these participants' reactions to pleasant images, $b = -0.37$, 95% CI [−8.64, 7.90], $SE = 4.218$, $t(68) = -0.09$, $p = .931$, $R^2 = .000$. This rules out the possibility that observing more negative reactions to unpleasant images leads participants to report more negative reactions overall. Participants exposed to more extreme reactions to unpleasant images had significantly more negative reactions to unpleasant images ($M = -52.98$, $SE = 0.61$), compared to those exposed to less extreme reactions ($M = -34.24$, $SE = 0.56$). These participants did not differ in reactions to pleasant images ($M = 36.39$, $SE = 0.49$; $M = 35.59$, $SE = 0.48$, for more extreme vs. less extreme affect norm conditions, respectively).

Simple slope analyses showed that among participants for whom unpleasant images were manipulated, reactions to unpleasant images became increasingly less extreme among participants in the less extreme condition, $b = 0.087$, $SE = 0.009$, $p < .001$, and more extreme among participants in the more extreme condition, $b = -0.047$, $SE = 0.009$, $p < .001$ (Figure 12B). Reactions to pleasant images became less extreme among participants in both the less extreme condition, $b = -0.047$, $SE = 0.009$, $p < .001$, and the more extreme condition, $b = -0.030$, $SE = 0.008$, $p < .001$, possibly reflecting habituation effects.

Manipulated Reactions to Pleasant Images. For participants who observed manipulated reactions to pleasant images, there was a similar pattern. We found a significant effect for condition, $b = 23.08$, 95% CI [13.83, 32.34], $SE = 4.720$, $t(68) = 4.89$, $p < .001$, $R^2 = .151$, indicating that participants, who observed affective reactions which were more (vs. less) positive than the mean in the population, reported more positive reactions to pleasant images (Figure 12A). We found no significant effect for condition on these participants' reactions to unpleasant images, $b = -3.16$, 95% CI [−10.48, 4.16], $SE = 3.734$, $t(68) = -0.85$, $p = .400$, $R^2 = .003$. Participants exposed to more extreme reactions to pleasant images had significantly more positive reactions to pleasant images ($M = 47.92$, $SE = 0.53$), compared to those exposed to less extreme reactions ($M = 21.84$, $SE = 0.43$). These participants did not differ in reactions to unpleasant images ($M = -39.81$, $SE = 0.50$; $M = -34.37$, $SE = 0.51$, for more extreme vs. less extreme affect norm conditions, respectively).

Simple slope analyses showed that among participants for whom pleasant images were manipulated, reactions to pleasant images became increasingly less extreme among participants in the less extreme condition, $b = -0.061$, $SE = 0.008$, $p < .001$, and more extreme among participants in the more extreme condition, $b = 0.041$, $SE = 0.008$, $p < .001$ (Figure 12B). Reactions to unpleasant

Figure 12

Participants' Affective Reactions by Affect Norm Condition, Image Type (Pleasant, Unpleasant), and Manipulated Affect (Pleasant vs. Unpleasant; Study 3)

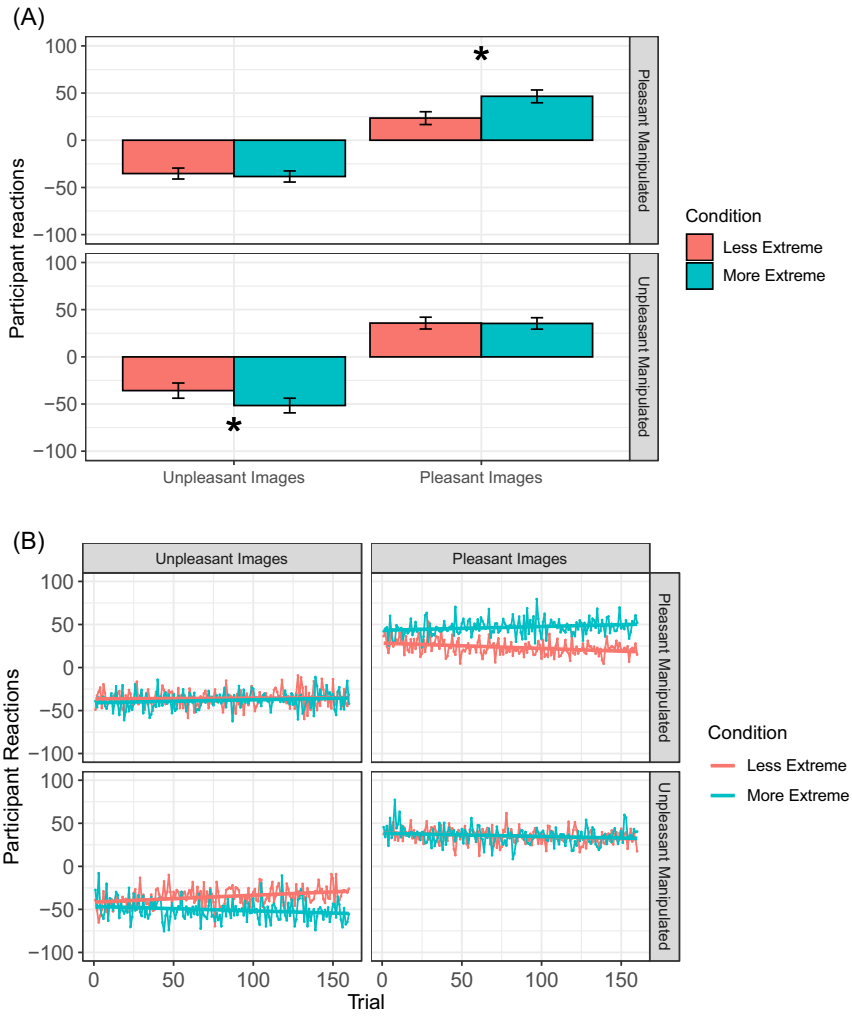

Note. Panel A presents the mean participant reactions per affect norm condition. Panel B plots the mean participant reactions per affect norm condition by trial number. Error bars represent the 95% confidence intervals. See the online article for the color version of this figure.
* $p < .05$.

images were unchanged among participants in the less extreme condition, $b = 0.017$, $SE = 0.009$, $p = .055$, and became more extreme among participants in the more extreme condition, $b = 0.030$, $SE = 0.009$, $p = .001$.

Does Exposure to Others' Feelings Shape Observers' Social Judgments of Novel Targets?

To test whether participants judged others based on the congruency of their affective reactions with the inferred mean of observed reactions, we conducted a linear regression model predicting social judgments, as in Study 2. As predicted, participants judged targets whose reactions were incongruent with the observed mean in their condition less favorably than targets whose reactions were congruent with the observed mean—for pleasant images: $b = -6.02$, 95% CI [−11.04, −1.00], $SE = 2.550$, $t(276) = -2.36$, $p = .019$, $R^2 = .020$; for unpleasant images: $b = -7.40$, 95% CI [−12.65, −2.14], $SE = 2.671$, $t(278) = -2.77$, $p = .006$, $R^2 = .027$.

We found weaker evidence for specificity according to the valence of the images manipulated. As predicted, a pattern of valence specificity emerged for participants who observed manipulated reactions to unpleasant images, such that targets whose reactions were incongruent with the observed mean in their condition were judged less favorably than targets whose reactions were congruent with the observed mean, only for targets who were reacting to unpleasant images, $b = -13.31$, 95% CI [−20.54, −6.07], $SE = 3.660$, $t(138) = -3.64$, $p < .001$, $R^2 = .087$, and not for targets who were reacting to pleasant images, $b = -6.12$, 95% CI [−13.42, 1.17], $SE = 3.688$, $t(138) = -1.66$, $p = .099$, $R^2 = .020$. For participants who observed manipulated reactions to pleasant images, we expected targets whose

reactions were incongruent with the observed mean in their condition to be judged less favorably than targets whose reactions were congruent with the observed mean, only for targets who were reacting to pleasant images and not to unpleasant images. The group difference when judging targets reacting to pleasant images was in the expected direction, but was not significant, $b = -5.88$, 95% CI [−12.87, 1.10], $SE = 3.533$, $t(136) = -1.67$, $p = .098$, $R^2 = .020$. As expected, there was no difference when judging targets reacting to unpleasant images, $b = -1.49$, 95% CI [−9.10, 6.12], $SE = 3.851$, $t(138) = -0.39$, $p = .699$, $R^2 = .001$.

To explore the interaction between participant condition, affect norm condition (less vs. more extreme reactions of others), and manipulated affect (pleasant vs. unpleasant), we conducted post hoc pairwise comparisons with Bonferroni p-adjustment. As shown in Figure 13, among participants who observed manipulated reactions to unpleasant images, participants in the less extreme condition judged targets whose reactions to unpleasant images were less extreme than the population mean more favorably ($M = 59.69$, $SE = 3.06$) than participants in the more extreme condition ($M = 42.19$, $SE = 3.63$), $p < .001$. However, these participants did not differ in their judgments of targets whose reactions to unpleasant images were more extreme than the population mean, $p = .087$. There were no significant differences between these participants in judgments of targets who reacted to pleasant images. Among participants who observed manipulated reactions to pleasant images, the post hoc comparisons were nonsignificant.

Discussion

Study 3 replicated the key results of the first two studies, using a 2 × 2 design in the affect norms learning stage. The findings demonstrate that people infer mean reactions of multiple others to both pleasant and unpleasant stimuli and assimilate their own responses to these means. In Study 1, participants came to expect others' reactions to match both less and more extreme responses, according to what they were exposed to in their condition. In Study 2, participants came to expect others' reactions to match the responses of others when these responses were less extreme, but not when these responses were more extreme. In Study 3, however, similar to Study 1, participants came to expect the reactions of others, both when these reactions were less extreme and when they were more extreme. Although we found evidence for learning and assimilating to the mean response when others' reactions were less or more extreme, effects on social judgments were not similarly symmetric. Instead, as in Study 2, participants generally judged targets with more (vs. less) extreme affective reactions more favorably. Moreover, congruence with mean affective response in the population resulted in more favorable judgments only when targets were exposed to less (but not more) extreme reactions to unpleasant stimuli.

Study 4

Our studies were designed to show that when people are exposed to feelings of multiple others, they infer a norm (i.e., what is common and appropriate to feel). We assumed that if people learn what is appropriate to feel, they would judge people whose reactions are incongruent with the norm more negatively, implying their reactions are inappropriate. In Studies 1–3, in the social judgment stage, participants were informed about the affective reactions of others (i.e., targets) to new images and then judged these targets. Such judgments may have been based on a comparison to the norm that was learned. However, another possibility is that such judgments were based on a comparison to the participant's own affective reactions to the new images, which are likely to be consistent with the norm, given that we found that participants assimilate their own reactions to the norm. If people judge those whose reactions are

Figure 13

Social Judgments by Affect Norm Condition (Study 3)

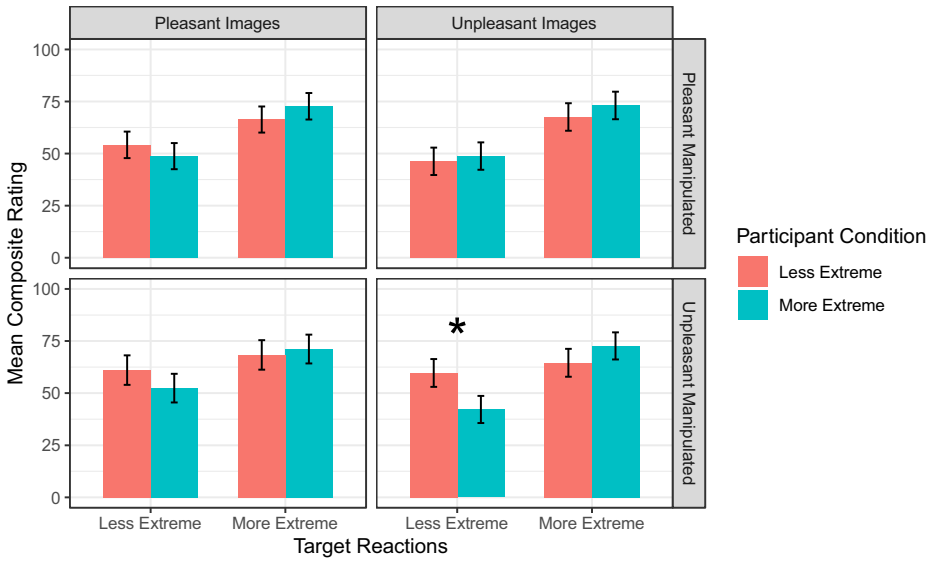

Note. Error bars represent the 95% confidence intervals. See the online article for the color version of this figure.

* $p < .05$.

inconsistent with the norm more negatively, even when they are not exposed to the emotional stimuli to which others are presumably reacting to, they cannot judge others based on their own emotional reactions (as there are no stimuli to react to). To test this possibility in Study 4, in the social judgment stage, participants were informed about "the mean responses" of others to tentatively "similar stimuli," without exposing participants to specific stimuli. Given that in Study 3, effects on judgments were found when reactions to unpleasant stimuli were manipulated in the learning stage, in Study 4, we included only unpleasant stimuli. This allowed us to indicate in the social judgment phase that targets were reacting to stimuli that were similar to those that were previously shown in the study.

Method

Study 4 was preregistered (https://aspredicted.org/JHX_WSL).[12]

Participants

The final sample consisted of 72 Jewish Israeli participants (59 self-identified as women and 13 identified as men; $M_{age} = 22.97$, $SD_{age} = 2.14$). Two additional participants asked to stop their participation during the experiment and were excluded from the analysis.

Procedure

The procedure was identical to that of Study 1, except for the following changes. First, only unpleasant stimuli were included. Accordingly, participants were assigned to one of two conditions (exposure to more vs. less extreme reactions). The affect norms learning stage included 70 trials: 35 trials in which participants viewed unpleasant images and rated their own affective reaction and 35 trials in which participants viewed unpleasant images and estimated another participant's affective reaction. Second, the social judgment stage was modified so that participants were presented with the targets' purported mean affective reactions to images similar to those included in the study up to that point, without viewing any specific images that may elicit affective reactions in the participants (Figure 14). Five targets' reactions were congruent with the less extreme affect norm (less negative than the mean in the population), and five targets' reactions were congruent with the more extreme affect norm (more negative than the mean in the population). As such, each participant judged five targets that were congruent with the mean reaction they observed in the learning stage and five targets whose reactions were incongruent with it. The targets' reactions were nonoverlapping, in that targets whose reactions were congruent with one condition were incongruent with the other condition (fell outside the other condition's distribution of values, as determined in a pilot study).[13] Participants evaluated each target on likeability, competence, and trustworthiness on a scale from 0 to 100 ($\alpha = .94$). Finally, participants completed several trait measures for exploratory purposes, provided demographic information and were debriefed.[14]

Materials

Images. We used a subset of the unpleasant images pretested in the pilot study (see the Supplemental Materials for the full list of images used in each study).

Names. As in Study 1, we selected a subset of the 80 most popular names for females and males.

Results

Does Exposure to Others' Feelings Shape Estimated Feelings?

To test whether the manipulation affected estimated affective reactions of others, we conducted a mixed models analysis, as in Studies 1 and 2. As predicted, we found a significant interaction between trial and condition, $b = -0.25$, 95% CI [−0.32, −0.18], $SE = 0.036$, $t(2369) = -6.91$, $p < .001$, $R^2 = .160$, indicating that participants' estimations of others' feelings diverged over time, as a function of the affect norm condition (Figure 15). The bias was in the direction of the observed mean in their condition. Simple slope analyses showed that estimations of reactions became increasingly less extreme among participants in the less extreme condition, $b = 0.207$, $SE = 0.025$, $p < .001$. Unlike Studies 1 and 3, but like Study 2, we did not find a significant learning trend among participants in the more extreme condition, $b = -0.044$, $SE = 0.026$, $p = .089$.

We further tested whether participants' estimations became increasingly closer to the mean observed response in their condition. To do so, we conducted a mixed models analysis, as in Studies 1 and 2. As predicted, when comparing estimations and feedback per trial, we found a significant main effect for trial, $b = 0.22$, 95% CI [0.17, 0.27], $SE = 0.025$, $t(2424) = 8.56$, $p < .001$, $R^2 = .145$, indicating that participants in both affect norm conditions learned to estimate others' affective reactions over time (see Supplemental Figure S6). Participants who were exposed to less negative reactions of others to unpleasant images expected others to react less negatively ($M = -24.78$, $SE = 0.52$) than did participants exposed to more negative reactions of others ($M = -43.16$, $SE = 0.58$).

Does Exposure to Others' Feelings Shape Observers' Own Feelings?

We tested whether the manipulation affected participants' own reactions by conducting a mixed models analysis, as in Studies 1 and 2. As predicted, we found a significant effect for condition, $b = -15.51$, 95% CI [−22.79, −8.23], $SE = 3.714$, $t(70) = -4.12$, $p < .001$, $R^2 = .065$, indicating that participants who observed affective

[12] See the Supplemental Materials for any discrepancies between the preregistered analyses and the analyses described in the article.

[13] Participants completed the Big Five Inventory (John & Srivastava, 1999).

[14] As in Study 3, when participants in Study 4 were explicitly asked whether they were suspicious that "others" were not real, this direct question raised participants' suspicions, but still only 43% expressed such suspicion. When asked what the study was about, 24% of participants believed the study tested whether and who has the ability to accurately predict others' emotional responses; 16% believed the study tested how people's emotional reactions predict their expectations of others' responses; 15% believed the study assessed how people differ in emotional reactions; 15% offered idiosyncratic accounts (e.g., effects of different names, whether people react differently to images of humans vs. animals); 14% believed the study assessed how one's own feelings are associated to the feelings of others, without specifying the direction of this association; 13% believed the study tested whether people react similarly to others; 3% did not know. None of the participants correctly identified the purpose of the study.

Figure 14

A Depiction of a Trial in the Social Judgment Stage in Study 4

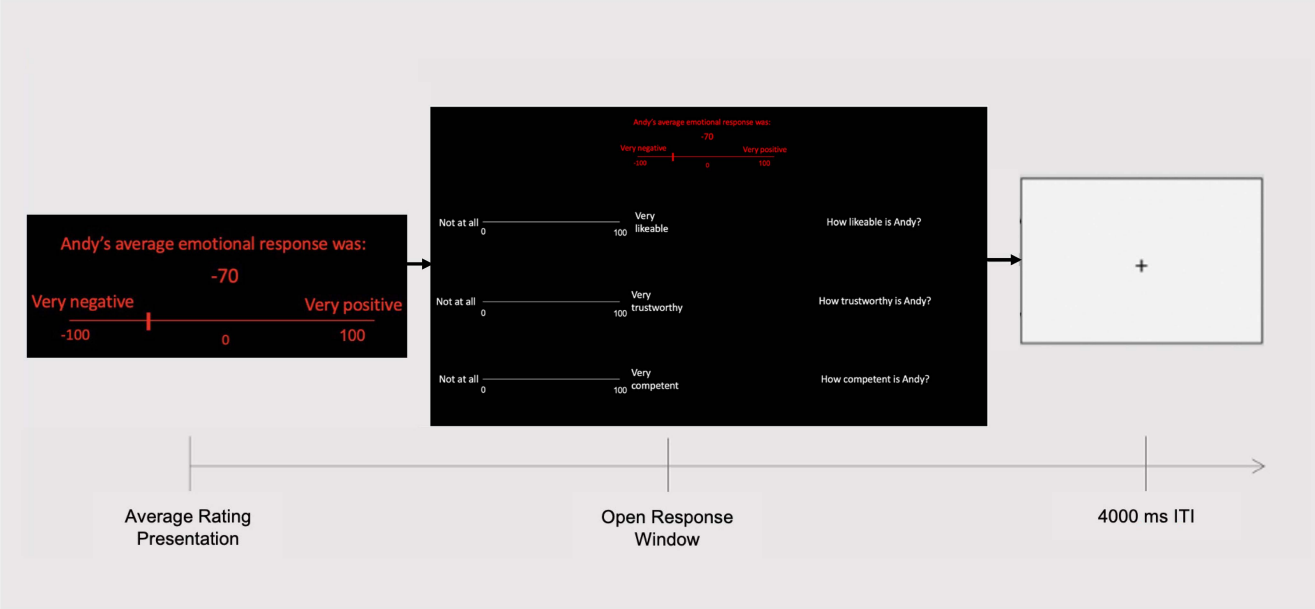

Note. ITI = intertrial interval. See the online article for the color version of this figure.

reactions that were more extreme (more negative) than the mean in the population reported more negative reactions to unpleasant images, compared to participants who observed affective reactions that were less extreme (less negative) than the population mean (Figure 16A). Participants exposed to more extreme reactions had more negative reactions to unpleasant images ($M = -43.16$, $SE = 0.58$), compared to participants exposed to less extreme reactions ($M = -24.78$, $SE = 0.52$).

Simple slopes tests showed that changes in affective reactions were significant in both the less extreme condition, $b = 0.102$, $SE = 0.023$, $p < .001$, and the more extreme condition, $b = -0.126$, $SE = 0.023$, $p = .006$ (see Figure 16B; see also Supplemental Figure S7).

Figure 15

Participants' Estimated Affective Reactions of Others by Trial Number and Affect Norm Condition (Study 4)

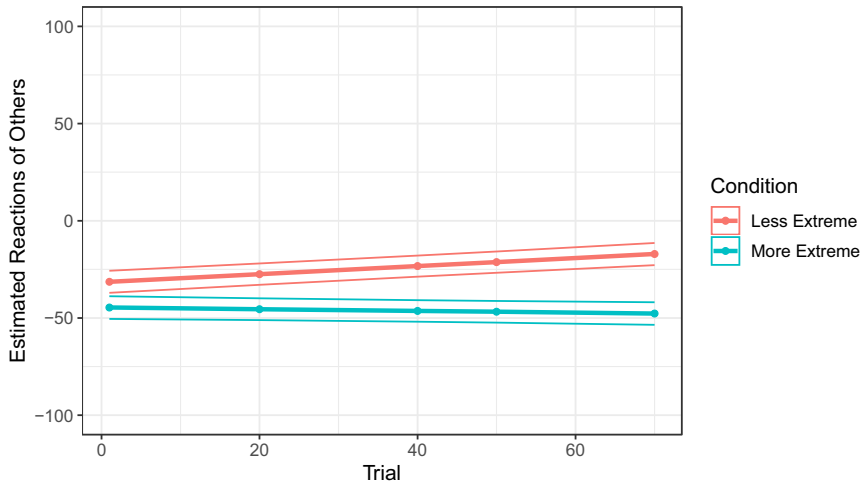

Note. Error bands represent the 95% confidence intervals. See the online article for the color version of this figure.

Figure 16

Participants' Affective Reactions by Affect Norm Condition (Study 4)

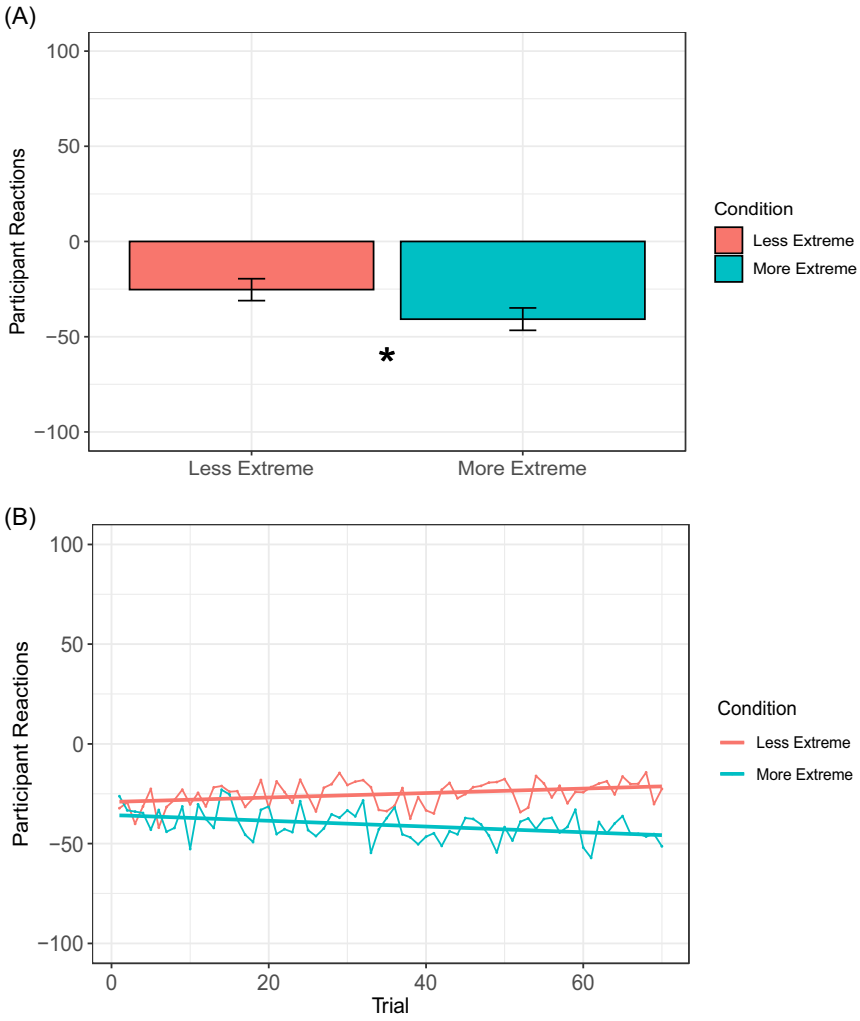

Note. Panel A presents the mean participant reactions per affect norm condition. Panel B plots the mean participant reactions per affect norm condition by trial number. Error bars and bands represent the 95% confidence intervals. See the online article for the color version of this figure.
* $p < .05$.

Does Exposure to Others' Feelings Shape Observers' Social Judgments of Novel Targets?

To test whether participants judged others whose affective reactions were incongruent with the observed mean reaction less favorably, we conducted a linear regression model, predicting social judgments, as in Studies 2 and 3. As predicted, participants judged targets whose mean reaction was incongruent (vs. congruent) with the observed mean in their condition less favorably, $b = -8.29$, 95% CI [−7.29, −5.03], $SE = 1.659$, $t(718) = -5.00$, $p < .001$, $R^2 = .034$ (see Figure 17). As in prior studies, post hoc pairwise comparisons with Bonferroni p-adjustment indicated that effects of affect norm condition on judgments were significant when evaluating targets whose reactions to unpleasant images were less extreme. Participants in the less extreme condition judged such targets more favorably ($M = 55.37$, $SE = 1.51$) than participants in the more extreme condition ($M = 41.80$, $SE = 1.70$), $p < .001$. The effect of condition on judgments of targets whose reactions to unpleasant images were more extreme was not significant, $p = .056$, albeit in the expected direction.

Discussion

Study 4 focused on the potential effects of exposure to feelings of multiple others on social judgments, demonstrating that such effects cannot be attributed to trial-by-trial social comparisons (i.e., comparing a target's reactions to one's own reactions). In particular, we were able to replicate the effects on social judgments found in prior studies. This is despite the fact that participants in Study 4 no longer saw the stimuli to which targets were presumably reacting, and were not exposed to specific affective reactions in the judgment phase, similar to those to which they were exposed during the learning phase. Instead, participants received abstract information

Figure 17

Mean Social Judgments by Affect Norm Condition and Targets' Affective Reactions to Unpleasant Images (Study 4)

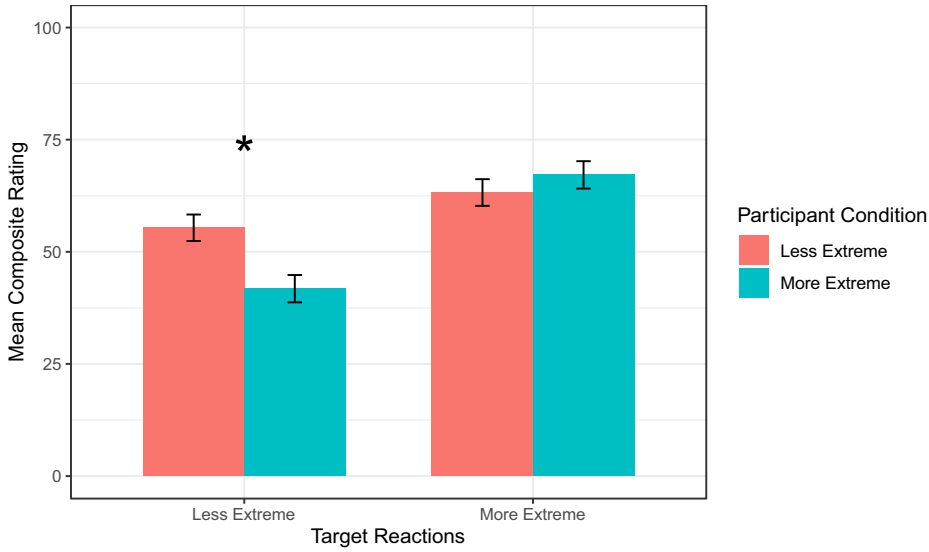

Note. Error bars represent the 95% confidence intervals. See the online article for the color version of this figure.
* $p < .05$.

about targets' mean affective reactions to similar stimuli. Our findings demonstrate that participants may infer a mean affective response and then use it as a reference point when judging the likability, trustworthiness, or competence of others. Nonetheless, given that the study did not include direct measures of perceived norms, the conclusion that a generalized norm representation has been formed should be considered with caution.

Transparency and Openness

We report how we determined our sample size, all data exclusions, all manipulations, and all measures in their respective introduction sections. All data, analysis code, and research materials are available at https://osf.io/uaq79/?view_only=1f54ead24bab453ca4053c5fb0495a9f. Data were analyzed using R, Version 4.2.1. Study 1 was not preregistered. Study 2 (https://aspredicted.org/FMR_XSD), Study 3 (https://aspredicted.org/JNH_DTR), and Study 4 (https://aspredicted.org/JHX_WSL) were preregistered.

General Discussion

Whether virtually or in person, we are surrounded by other people and can often register their affective reactions. Decades of research have documented how exposure to feelings of a single person in a specific context shapes how observers behave, react, and appraise in that context. In this investigation, we demonstrate how people infer additional information from exposure to multiple others. In a series of studies, we demonstrate that when people were exposed to feelings of others, they learned to estimate the reaction in the population and used it to predict novel reactions. They assimilated their own affective reactions to the estimated reaction in the population. Finally, exposure to feelings of others led people to consider the mean reaction in the population as a benchmark for evaluative judgments. This effect could not be explained by social contagion alone, as it required inferential processes. Some targets, whose affective reactions deviated from the norm, were judged less favorably. In particular, deviations from affect norms were judged negatively only when they involved less negative reactions to unpleasant stimuli. Taken together, these findings suggest that exposure to feelings of others may be one process by which people learn affect norms.

Linking Affective Social Learning to Affect Norms

Exposure to affective reactions of others to an object can influence how one behaves toward, feels, and thinks about that object (for a review, see Clément & Dukes, 2017). Such affective social learning is linked to the context of exposure. In this investigation, we demonstrate that what we learn from exposure to feelings of others extends beyond our reactions to specific objects. Just as observers infer norms of behavior from exposure to behaviors of others (Cialdini et al., 1990), observers can infer norms of feelings from exposure to feelings of others. These norms inform observers what to expect in future encounters with novel individuals. These norms can also signal to observers what feelings are appropriate or inappropriate, informing social judgments that are unrelated to the exposure context or to the value of objects per se.

Our findings provide initial support for these ideas by demonstrating how exposing observers to affective reactions of others informs their predictions, their own affective reactions, and their subsequent social judgments of novel targets. Observers judged novel targets more negatively if their negative affective

reactions diverged from the estimated response in the population (Studies 1–3), even when such responses were described without access to the objects that elicited such reactions (Study 4). Our findings, therefore, indicate that what we learn from exposure to feelings of others can extend beyond the value of affect-inducing stimuli to the value of affect itself.

By demonstrating how exposure to the feelings of others can shape subsequent evaluative judgments, our findings bridge research on affective social learning with research on affect and emotion norms. Members of different cultures differ in their evaluations of affective states. For example, European Americans value high-arousal pleasant feelings more and low-arousal pleasant feelings less than Asian Americans and Hong Kong Chinese do (Tsai et al., 2006). These evaluations, in turn, are linked to social judgments. For example, Asian Americans prefer physicians who express calmness, whereas European Americans prefer physicians who express excitement (Sims & Tsai, 2015). Observers judge targets more favorably, therefore, if targets' affective reactions are consistent with observers' valued feelings. There are likely many mechanisms by which cultures transmit the value of affective states (e.g., Mesquita et al., 2014). The current investigation may point to one such mechanism. Exposure to patterns of affective reactions in group members may lead people to infer not only what group members are expected to feel but also what they *should* feel.

Theoretical and Applied Implications

Affect norms define what feelings are likely and appropriate in one's social environment, and impel related experiences, cognitions, and behaviors (Thoits, 2004). Affect norms can be transmitted through processes of socialization, a process by which children learn about feelings from parents and other adults (e.g., through operant conditioning, modeling, and explicit learning; Eisenberg et al., 1998). But how do parents learn affect norms? In this investigation, we demonstrate that people can learn affect norms from exposure to feelings of others. Based on such repeated exposure, people learn what to expect from others, assimilate their response to that expectation, and judge others negatively when they deviate from such norms.

Our findings suggest that the development of affect norms may be similar to the development of other social norms but that affect norms may also have unique characteristics (see also Vishkin & Tamir, 2023). Not all targets whose affective reactions deviated from the expected reactions in the population were judged negatively. Across studies, targets who expressed more intense affective reactions were judged more favorably. In general, these findings show that with respect to affect norms, frequency does not automatically dictate value. Our pattern is consistent with evidence that people who express emotions more intensely, at least in Western cultural contexts, are judged more favorably (Fultz et al., 2024). Furthermore, across studies, we found that targets were judged negatively when diverging from the norm only when they reacted less negatively than expected. Although further research is needed to test why this is the case, it might be because affective reactions serve as important evaluative social signals, even (and perhaps particularly) when they are unexpected. Reactions to unpleasant stimuli that are expected or socially appropriate will be judged more favorably, whether they are less or more intense. Reactions to unpleasant stimuli that are unexpected or socially inappropriate will be judged less favorably when they are less intense, but not when they are more intense. This is possibly because more intense unpleasant reactions could alert group members to potential threats and, therefore, may be beneficial even when they are unexpected or socially inappropriate. Another possibility is that unpleasant reactions that are less intense than expected may signal poor moral character, but this is less likely for reactions that are more intense (Parrott, 2019).

From a more applied perspective, in our modern world, we are exposed to the feelings of others more frequently than ever, on social media. Such exposure may lead to "digital emotion contagion" (Goldenberg & Gross, 2020), whereby the feelings of others observed online influence the feelings of the observer (Fan et al., 2020; Ferrara & Yang, 2015). While it is possible that these effects are tied to the specific context, our findings suggest that such exposure to the feelings of multiple others may extend beyond the exposure context. For instance, when scrolling through a feed of posts, an individual may encounter a news article another user shared, observing how multiple others reacted to it. In addition to influencing the observer's reactions to the article, these reactions may lead the observer to infer which affective reactions are appropriate to new articles, shaping how they react to subsequent articles that are shared on the platform. Furthermore, this may lead observers to judge other users who express incongruent reactions less favorably, regardless of the content to which they are reacting. Although these ideas should be directly tested, understanding how these processes unfold, and the determinants of the occurrence (or co-occurrence) of these two potential effects, might help us understand the network effects that may be driving polarization and other important societal phenomena today (Del Vicario et al., 2016).

Limitations and Future Directions

This investigation used a novel experimental paradigm to study effects of exposure to feelings of multiple others on predictions, observers' affective reactions, and subsequent social judgments. The paradigm's flexibility allowed us to manipulate different aspects of learning over time, while assessing both concurrent and subsequent effects. Our findings offer initial evidence for effects that go beyond the immediate context, but they also raise new questions.

First, across studies, we found that people tended to judge new targets whose affective reactions deviated from the mean observed response in the population less favorably. This pattern, however, did not apply equally to all affective responses. Instead, we found that observers tend to judge targets who express more intense affective reactions more favorably overall. It is judgments of targets who express less intense affective reactions to unpleasant stimuli that vary as a function of prior exposure to affective reactions of others. It is possible that more favorable judgments of more intense affective reactions reflect existing affect norms that participants have cultivated through their own prior experiences. It is also possible that people who deviate from a normatively intense negative reaction are judged harshly because such deviation could be particularly costly, leading people to miss potential danger. Further research is needed to understand why our effects on judgment were more pronounced for some reactions than for others. If people generally penalize others who react less extremely than the affect norm, but do not penalize others who react more extremely, this could encourage

polarization, as less extreme reactions would meet more negative feedback from others.

Second, our studies focused on either unpleasant or pleasant feelings. We found consistent evidence for learning unpleasant affect norms, whereas evidence for learning pleasant affect norms was less consistent. Future research could test whether learning from others' feelings is stronger when such feelings are unpleasant versus pleasant. Future research could also test whether learning affect norms applies to learning norms of discrete emotions.

Third, our paradigm involved exposure to feelings of others. However, to track concurrent effects on prediction and assimilation, we asked participants to estimate others' responses before exposing them to the response. We are unable to determine, therefore, whether our effects arise from the mere exposure to others' feelings or depend on response-contingent reinforcement. Future research could address this question by testing whether our effects replicate when participants are exposed to others' feelings without estimating them and receiving feedback on their estimations.

Fourth, in our studies, participants were exposed to the emotions of others in a sequential manner. It remains to be tested whether people infer affect norms in a similar manner when people are exposed to feelings of multiple others simultaneously (i.e., in a crowd). Previous research has shown that individuals tend to amplify feelings in a group (Goldenberg et al., 2021; Goldenberg, LaFollette, et al., 2022; Goldenberg, Schöne, et al., 2022), whereas in our studies, with additional exposure, people became increasingly more accurate in their estimations of others' feelings. Future research can test whether people infer different information from others when presented sequentially versus simultaneously.

Fifth, in our task, participants learned about the feelings of others via ratings on a scale. However, people typically learn about the feelings of others through facial expressions, text, or emojis. Future tasks could utilize more ecological sources of information. It would also be useful to assess whether certain sources of learning, such as facial expressions, lead to stronger effects than others.

Sixth, taken together, our studies suggest that inferences of affect norms are not equivalent across levels of intensity. In some of our studies (Studies 2 and 4), people exposed to less extreme reactions to unpleasant stimuli learned a new norm, whereas those exposed to more extreme reactions did not. This may be because in these studies, there were relatively fewer learning trials. Thus, it is possible that learning some affect norms requires more learning opportunities. Future studies could explore what underlies these asymmetrical effects.

Finally, people are more likely to assimilate to norms inferred from people with whom they identify (Miller & Prentice, 1996). What people learn and how such learning affects behaviors and judgments may depend on social identification. For instance, people may be more likely to assimilate to an affect norm that was inferred from members of the ingroup than the outgroup. This is consistent with research on mimicry (Hess & Fischer, 2014), showing that individuals are more likely to mimic emotions of ingroup (vs. outgroup) members (e.g., Bourgeois & Hess, 2008). To optimize potential learning, in the current investigation, participants were always exposed to ingroup members of the same gender. Future studies could test whether and to what extent the effects of exposure to feelings of others depend on the extent to which observers identify with them.

Constraints on Generality

Our studies were tested on local student samples of Jewish Israelis, who were presumably exposed to affective reactions of other Jewish Israeli students, like themselves. Thus, our conclusions are constrained to this population and should be replicated in other cultural contexts.

Summary

Similar to the ways in which people infer norms of behavior, by learning about affective reactions of others, people may infer affect norms—norms of feelings. In four studies, we demonstrated how people learn to expect affective reactions in a population, assimilate to them, and judge novel targets negatively on unrelated social dimensions if their affective reactions deviated from the learned norm. Such evidence begins to uncover how people learn the value of feelings from exposure to feelings of others in their social and cultural context.

References

Bicchieri, C. (2006). *The grammar of society: The nature and dynamics of social norms*. Cambridge University Press. https://doi.org/10.1017/CBO9780511616037

Biernat, M., & Eidelman, S. (2007). Standards. In A. W. Kruglanski & E. T. Higgins (Eds.), *Social psychology: Handbook of basic principles* (2nd ed., pp. 308–333). Guilford Press. https://doi.org/10.1111/j.1467-9221.2008.00665.x

Bourgeois, P., & Hess, U. (2008). The impact of social context on mimicry. *Biological Psychology*, *77*(3), 343–352. https://doi.org/10.1016/j.biopsycho.2007.11.008

Cialdini, R. B., & Goldstein, N. J. (2004). Social influence: Compliance and conformity. *Annual Review of Psychology*, *55*(1), 591–621. https://doi.org/10.1146/annurev.psych.55.090902.142015

Cialdini, R. B., Reno, R. R., & Kallgren, C. A. (1990). A focus theory of normative conduct: Recycling the concept of norms to reduce littering in public places. *Journal of Personality and Social Psychology*, *58*(6), 1015–1026. https://doi.org/10.1037/0022-3514.58.6.1015

Clément, F., & Dukes, D. (2017). Social appraisal and social referencing: Two components of affective social learning. *Emotion Review*, *9*(3), 253–261. https://doi.org/10.1177/1754073916661634

Del Vicario, M., Vivaldo, G., Bessi, A., Zollo, F., Scala, A., Caldarelli, G., & Quattrociocchi, W. (2016). Echo chambers: Emotional contagion and group polarization on Facebook. *Scientific Reports*, *6*(1), Article 37825. https://doi.org/10.1038/srep37825

Delvaux, E., Vanbeselaere, N., & Mesquita, B. (2015). Dynamic interplay between norms and experiences of anger and gratitude in groups. *Small Group Research*, *46*(3), 300–323. https://doi.org/10.1177/1046496415576411

Eid, M., & Diener, E. (2001). Norms for experiencing emotions in different cultures: Inter- and intranational differences. *Journal of Personality and Social Psychology*, *81*(5), 869–885. https://doi.org/10.1037/0022-3514.81.5.869

Eisenberg, N., Cumberland, A., & Spinard, T. L., (1998). Parental socialization of emotion. *Psychological Inquiry*, *9*(4), 241–273. https://doi.org/10.1207/s15327965pli0904_1

Fan, R., Xu, K., & Zhao, J. (2020). *Weak ties strengthen anger contagion in social media*. arXiv. https://doi.org/10.48550/arXiv.2005.01924

Fehr, E., & Schurtenberger, I. (2018). Normative foundations of human cooperation. *Nature Human Behavior*, *2*(7), 458–468. https://doi.org/10.1038/s41562-018-0385-5

Ferrara, E., & Yang, Z. (2015). Measuring emotional contagion in social media. *PLOS ONE*, *10*(11), Article e0142390. https://doi.org/10.1371/journal.pone.0142390

Fultz, A. A., Stosic, M. D., & Bernieri, F. J. (2024). Nonverbal expressivity, physical attractiveness, and liking: First impression to established relationship. *Journal of Nonverbal Behavior*, *48*, 189–211. https://doi.org/10.1007/s10919-023-00444-7

Goldenberg, A., Garcia, D., Halperin, E., & Gross, J. J. (2020). Collective emotions. *Current Directions in Psychological Science*, *29*(2), 154–160. https://doi.org/10.1177/0963721420901574

Goldenberg, A., & Gross, J. J. (2020). Digital emotion contagion. *Trends in Cognitive Sciences*, *24*(4), 316–328. https://doi.org/10.1016/j.tics.2020.01.009

Goldenberg, A., LaFollette, K., Huang, Z., Weisz, E., & Cikara, M. (2022). *Judgment of crowds as emotional increases with the proportion of Black faces* [Preprint]. PsyArXiv. https://doi.org/10.31234/osf.io/ys65p

Goldenberg, A., Saguy, T., & Halperin, E. (2014). How group-based emotions are shaped by collective emotions: Evidence for emotional transfer and emotional burden. *Journal of Personality and Social Psychology*, *107*(4), 581–596. https://doi.org/10.1037/a0037462

Goldenberg, A., Schöne, J., Huang, Z., Sweeny, T. D., Ong, D. C., Brady, T. F., Robinson, M. M., Levari, D., Zaki, J., & Gross, J. J. (2022). Amplification in the evaluation of multiple emotional expressions over time. *Nature Human Behaviour*, *6*(10), 1408–1416. https://doi.org/10.1038/s41562-022-01390-y

Goldenberg, A., Weisz, E., Sweeny, T. D., Cikara, M., & Gross, J. J. (2021). The crowd-emotion-amplification effect. *Psychological Science*, *32*(3), 437–450. https://doi.org/10.1177/0956797620970561

Green, P., & MacLeod, C. J. (2016). SIMR: An R package for power analysis of eneralized linear mixed models by simulation. *Methods in Ecology and Evolution*, *7*(4), 493–498. https://doi.org/10.1111/2041-210X.12504

Hatfield, E., Cacioppo, J. T., & Rapson, R. L. (1994). *Emotional contagion*. Cambridge University Press. https://doi.org/10.1017/CBO9781139174138

Hess, U., & Fischer, A. (2014). Emotional mimicry: Why and when we mimic emotions. *Social and Personality Psychology Compass*, *8*(2), 45–57. https://doi.org/10.1111/spc3.12083

Higgins, E. T., & Lurie, L. (1983). Context, categorization, and recall: The "change-of-norm" effect. *Cognitive Psychology*, *15*(4), 525–547. https://doi.org/10.1016/0010-0285(83)90018-X

Hochschild, A. R. (1983). *The managed heart: Commercialization of human feeling*. University of California Press.

Horne, C., & Mollborn, S. (2020). Norms: An integrated framework. *Annual Review of Sociology*, *46*, 467–487. https://doi.org/10.1146/annurev-soc-121919-054658

John, O. P., & Srivastava, S. (1999). The Big-Five trait taxonomy: History, measurement, and theoretical perspectives. In L. A. Pervin & O. P. John (Eds.), *Handbook of personality: Theory and research* (Vol. 2, pp. 102–138). Guilford Press.

Kahneman, D., & Miller, D. T. (1986). Norm theory: Comparing reality to its alternatives. *Psychological Review*, *93*(2), 136–153. https://doi.org/10.1037/0033-295X.93.2.136

Keltner, D., & Haidt, J. (1999). Social functions of emotions at four levels of analysis. *Cognition and Emotion*, *13*(5), 505–521. https://doi.org/10.1080/026999399379168

Lang, P. J., Bradley, M. M., & Cuthbert, B. N. (2005). *International affective picture system (IAPS)* [Database record]. APA PsycTests. https://doi.org/10.1037/t66667-000

Manstead, A. S. R., & Fischer, A. H. (2001). Social appraisal: The social world as object of and influence on appraisal processes. In K. R. Scherer, A. Schorr, & T. Johnstone (Eds.), *Appraisal processes in emotion: Theory, methods, research* (pp. 221–232). Oxford University Press.

Marques, J., Abrams, D., Paez, D., & Martinez-Taboada, C. (1998). The role of categorization and in-group norms in judgments of groups and their members. *Journal of Personality and Social Psychology*, *75*(4), 976–988. https://doi.org/10.1037/0022-3514.75.4.976

Mesquita, B. (2010). Emoting: A contextualized process. In B. Mesquita, L. F. Barrett, & E. R. Smith (Eds.), *The mind in context* (pp. 83–104). Guilford Press.

Mesquita, B., De Leersnyder, J., & Albert, D. (2014). The cultural regulation of emotions. In J. J. Gross (Ed.), *Handbook of emotion regulation* (2nd ed., pp. 284–301). Guilford Press. https://ppw.kuleuven.be/cscp/documents/mesquita/mesquita-de-leersnyder-albert-2014-the-cultural.pdf

Miller, D. T., & Prentice, D. A. (1996). The construction of social norms and standards. In E. T. Higgins & A. W. Kruglanski (Eds.), *Social psychology: Handbook of basic principles* (pp. 799–829). The Guilford Press.

Miller, D. T., & Prentice, D. A. (2016). Changing norms to change behavior. *Annual Review of Psychology*, *67*(1), 339–361. https://doi.org/10.1146/annurev-psych-010814-015013

Mussweiler, T. (2003). Comparison processes in social judgment: Mechanisms and consequences. *Psychological Review*, *110*(3), 472–489. https://doi.org/10.1037/0033-295X.110.3.472

Niedenthal, P. M., Winkielman, P., Mondillon, L., & Vermeulen, N. (2009). Embodiment of emotion concepts. *Journal of Personality and Social Psychology*, *96*(6), 1120–1136. https://doi.org/10.1037/a0015574

Parkinson, B. (2011). Interpersonal emotion transfer: Contagion and social appraisal. *Social and Personality Psychology Compass*, *5*(7), 428–439. https://doi.org/10.1111/j.1751-9004.2011.00365.x

Parkinson, B., Fischer, A. H., & Manstead, A. S. (2005). *Emotion in social relations: Cultural, group, and interpersonal processes*. Psychology Press. https://doi.org/10.4324/9780203644966

Parrott, W. G. (2019). Emotions as signals of moral character. In U. Hess & S. Hareli (Eds.), *The social nature of emotion expression* (pp. 161–177). Springer. https://doi.org/10.1007/978-3-030-32968-6_9

Peters, K., & Kashima, Y. (2015). A multimodal theory of affect diffusion. *Psychological Bulletin*, *141*(5), 966–992. https://doi.org/10.1037/bul0000020

Prentice, D. A., & Miller, D. T. (1993). Pluralistic ignorance and alcohol use on campus: Some consequences of misperceiving the social norm. *Journal of Personality and Social Psychology*, *64*(2), 243–256. https://doi.org/10.1037//0022-3514.64.2.243

Sims, T., & Tsai, J. L. (2015). Patients respond more positively to physicians who focus on their ideal affect. *Emotion*, *15*(3), 303–318. https://doi.org/10.1037/emo0000026

Sinclair, S., & Agerström, J. (2023). Do social norms influence young people's willingness to take the COVID-19 vaccine? *Health Communication*, *38*(1), 152–159. https://doi.org/10.1080/10410236.2021.1937832

Smith, E. R., & Mackie, D. M. (2016). Representation and incorporation of close others' responses: The RICOR model of social influence. *Personality and Social Psychology Review*, *20*(4), 311–331. https://doi.org/10.1177/1088868315598256

Thoits, P. A. (2004). Emotion norms, emotion work and social order. In A. S. R. Manstead, N. Frijda, & A. Fischer (Eds.), *Feelings and emotions: The Amsterdam symposium* (pp. 359–378). Cambridge University Press. https://doi.org/10.1017/CBO9780511806582.021

Tsai, J. L., Knutson, B., & Fung, H. H. (2006). Cultural variation in affect valuation. *Journal of Personality and Social Psychology*, *90*(2), 288–307. https://doi.org/10.1037/0022-3514.90.2.288

Van Kleef, G. A. (2009). How emotions regulate social life: The emotions as social information (EASI) model. *Current Directions in Psychological Science*, *18*(3), 184–188. https://doi.org/10.1111/j.1467-8721.2009.01633.x

Vishkin, A., Kitayama, S., Berg, M. K., Diener, E., Gross-Manos, D., Ben-Arieh, A., & Tamir, M. (2023). Adherence to emotion norms is greater in individualist cultures than in collectivist cultures. *Journal of Personality and Social Psychology*, *124*(6), 1256–1276. https://doi.org/10.1037/pspi0000409

Vishkin, A., & Tamir, M. (2023). Emotions norms are unique. *Affective Science*, *4*, 453–457. https://doi.org/10.1007/s42761-023-00188-z

Westfall, J., Kenny, D. A., & Judd, C. M. (2014). Statistical power and optimal design in experiments in which samples of participants respond to samples of stimuli. *Journal of Experimental Psychology: General*, *143*(5), 2020–2045. https://doi.org/10.1037/xge0000014

Received September 21, 2024
Revision received November 29, 2025
Accepted January 24, 2026 ■