

People Update Their Injunctive Norm and Moral Beliefs After Receiving Descriptive Norm Information

Paul Deutchman^{1, 2}, Gordon Kraft-Todd¹, Liane Young¹, and Katherine McAuliffe¹

¹ Department of Psychology and Neuroscience, Boston College

² Center for Social Norms and Behavioral Dynamics, University of Pennsylvania

How do descriptive norms shape injunctive norm beliefs, and what does this tell us about the cognitive processes underlying social norm cognition? Across six studies ($N = 2,671$), we examined whether people update their injunctive norm beliefs—as well as their moral judgments and behavioral intentions—after receiving descriptive norm information about how common (or uncommon) a behavior is. Specifically, we manipulated the descriptive normativity of behaviors, describing behaviors as uncommon (20% of people were doing the behavior) or common (80% of people were doing the behavior), and the type of behavior across studies (fairness, conventional, harm, preference). To measure belief updating, we assessed beliefs prior to and after receiving information about the descriptive norm. We had three main findings: First, participants positively updated their prior injunctive norm beliefs, moral judgments, and behavioral intentions (i.e., rated behaviors more injunctively normative and moral) after receiving a common descriptive norm and negatively updated their beliefs (i.e., rated behaviors less injunctive and moral) after receiving an uncommon descriptive norm, and updated to a larger extent for the common than uncommon descriptive norm. Second, participants were more likely to update their beliefs about what is moral for *others* compared to what is moral for the *self*. Third, participants updated their beliefs to a greater extent for fairness and conventional behaviors compared to harm behaviors and preferences. Together, our findings suggest that descriptive norms shape our injunctive norm beliefs and moral judgments and help to paint a fuller picture of the social cognition of social norms.

Statement of Limitations

The present work has several potential limitations. First, our data were collected via online convenience samples, which, while generally more representative than a sample of college undergraduates, potentially limits the generalizability of our findings to people more broadly, particularly those from non-Western, educated, industrialized, rich, and democratic societies. Second, we used a limited number of vignettes as stimuli in our study. While these were stringently normed on several dimensions preinclusion, it is possible that our findings might not generalize to other normative behaviors. Third, because we used real (rather than novel) behaviors, participants likely entered our study with prior beliefs about those behaviors, potentially influencing the extent of their belief updating. However, because we were interested primarily in between-condition comparisons, this is unlikely to explain our results. Fourth, while our dependent measures are interrelated, our modeling approach treated them as independent, including each dependent measure in a different model. To account for interdependence between measures, we conducted exploratory multivariate regression models which support our initial findings. Lastly, because we assessed self-report behavioral intentions, it is possible that participants' actual behavior would diverge from their intentions—consequently, we cannot make claims about the effect of descriptive norms on behavior *per se*.

Keywords: social norms, descriptive norms, injunctive norms, moral judgment, social cognition

Supplemental materials: <https://doi.org/10.1037/pspa0000416.supp>

Andrew Todd served as action editor.

Paul Deutchman  <https://orcid.org/0000-0001-7612-653X>

All of the code, materials, and data are publicly available online on the Open Science Framework at https://osf.io/sc842/?view_only=83662b719da3479f8241eefce2c4ade4. All of the studies in this project were preregistered prior to data collection: Pilot Study, <https://aspredicted.org/jp2b7.pdf>; Study 1, <https://osf.io/e3c6n>; Study 2, <https://osf.io/gua3q>; Study 3, <https://osf.io/c6r8p>; Study 4, <https://osf.io/5pz8m>; and Study 5,

<https://osf.io/pdnyx>. All data, analysis script, and research materials are available on the Open Science Framework (<https://osf.io/sc842/>).

This project received funding from the John Templeton Foundation Grant 61061 awarded to Katherine McAuliffe and Liane Young. The authors thank the Cooperation Lab for the helpful feedback on all stages of the research process and Abby McLaughlin and Sophie Riddick for their help proofreading and editing the article.

Paul Deutchman played a lead role in data curation, formal analysis, investigation, methodology, project administration, validation, visualization,

continued

Why is it appropriate to serve chocolate muffins for breakfast in the United States while serving chocolate cupcakes seems odd? Why are we expected to tip food servers and hair dressers but not teachers or dentists? That we accept these unspoken rules, often without thinking or questioning them, demonstrates the ubiquity of social norms in human life. Social norms are a foundational part of human societies and pervade nearly every aspect of human social life—from what we eat for breakfast to how we share resources. However, while social norms are a foundational part of social interactions and have been a central focus of study in social psychology for over 40 years (Axelrod, 1986; Buffalo & Rodgers, 1971), we still have much to learn about the cognitive processes that underlie norm formation and change. Given the importance of norms in human social life, it is critical to develop a better understanding of the factors that shape normative beliefs.

Do People Update Their Beliefs After Receiving Descriptive Norm Information?

Much recent research on social norms has examined two related types of normative information: descriptive norms—what we think other people are actually doing (norms of *is*)—and injunctive norms, what we think other people approve or disapprove of (norms of *ought*; Cialdini et al., 1990). For example, that most people speak quietly in a library is a descriptive norm (it is what people do), while the belief that most people approve of talking quietly is an injunctive norm (it is what people approve of and expect). Both descriptive and injunctive norms can influence behavior in important ways (e.g., getting people to conserve energy by reusing towels, increasing gambling behavior, or promoting college drinking; Goldstein et al., 2008; Meisel & Goodie, 2014; Neighbors et al., 2008).

While these norms are often congruent—that is, most people generally approve of what is commonly done (Bear & Knobe, 2017; Blanton et al., 2008)—descriptive and injunctive norms can dissociate and differentially influence behavior (e.g., people might approve of sustainable consumption but perceive it as uncommonly done in practice; Christensen et al., 2004; Ge et al., 2020; Jacobson et al., 2011). Furthermore, while less common, descriptive and injunctive norms can also come into conflict (e.g., people are more likely to litter when littering is the descriptive norm, and there is an antilitter injunctive message than when there is only the descriptive norm; Cestac et al., 2014; Keizer et al., 2011; Ostermann & de Barcellos, 2021; Smith et al., 2012; Staunton et al., 2014). That descriptive and injunctive norms can occasionally dissociate has led some researchers to conceptualize them as distinct constructs (Cialdini et al., 2006; Reno et al., 1993).

Recent work, however, highlights the ways in which descriptive and injunctive norms are in fact highly interrelated. For example, using the Implicit Association Task, Eriksson et al. found that people show an automatic association between concepts that are descriptive

and concepts that are injunctive (2015). This association is not just implicit; people also made explicit bidirectional inferences between descriptive and injunctive norms—when told that a behavior is common (or uncommon), people infer that it is injunctive (or not injunctive). Other work finds that people tend to infer what *ought to be* (injunctive inferences) from what *is common* (descriptive norms; Tworek & Cimpian, 2016) and that this “is-to-ought” relationship is present early in development, emerging by 4 years of age (Roberts et al., 2017, 2018). By age 6, children’s injunctive norm beliefs are influenced by descriptive norm information that a behavior is common (Deutchman et al., 2023). Altogether, this work suggests that we implicitly associate and explicitly infer the injunctivity of a behavior from how common it is.

While there is some initial evidence that people infer injunctive normativity from descriptive information—suggesting that descriptive norms might partially contribute to injunctive norm formation—it remains unclear exactly how they might do so and to what extent. For example, although past work suggests that people make basic, explicit binary inferences between descriptive and injunctive norms (e.g., whether a behavior is injunctive or not given that it is common or uncommon; Eriksson et al., 2015), it is less clear how the frequency of the descriptive norm (e.g., the number or proportion of people actually engaging in the normative behavior) influences the extent to which we think others approve of a behavior. Furthermore, little work in this area has examined how people update their own beliefs before and after receiving normative information. Here we investigate this updating process using a repeated measures design which allows us to better understand how individuals adjust their beliefs in response to novel norm information. Namely, while past work has found that descriptive norm information influences injunctive beliefs via average group comparisons, we are among the first, to our knowledge, to measure the extent to which individuals update their own beliefs after receiving norm information as well as to assess whether this relationship between descriptive and injunctive norms varies depending on the normative behavior in question. Additionally, by measuring participants’ normative beliefs before and after the introduction of a descriptive norm, we can examine how individual differences in a priori descriptive norm beliefs moderate and constrain belief updating for new descriptive information, a potentially important factor in the effectiveness of normative interventions and nudges. A more fine-grained approach to exploring the relationship between descriptive and injunctive norms will allow us to better understand how descriptive norm information specifically changes the frequency of injunctive beliefs and, in the process, will help reveal how closely associated these concepts are and how they interact to influence behavior, thereby informing our understanding of the cognitive processes underlying norm cognition. Lastly, it is also important to understand how descriptive and injunctive norms relate to behavior. A large body of work has found that descriptive and injunctive norms influence behavior in important ways (Cialdini et al., 1990, 2006;

writing—original draft, and writing—review and editing. Gordon Kraft-Todd played a supporting role in conceptualization, methodology, supervision, and writing—review and editing. Liane Young played a lead role in funding acquisition and a supporting role in conceptualization, methodology, supervision, and writing—review and editing. Katherine McAuliffe played a lead role in funding acquisition and supervision, a supporting role in

methodology and writing—review and editing, and an equal role in conceptualization.

Correspondence concerning this article should be addressed to Paul Deutchman, Center for Social Norms and Behavioral Dynamics, University of Pennsylvania, McNeil Building, Room 438, 3718 Locust Walk, Philadelphia, PA 19104, United States. Email: pauld924@gmail.com

Raihani & McAuliffe, 2014; Reno et al., 1993). However, it is less clear whether people are more likely to engage in a behavior after receiving new information about the descriptive norm. Thus, while not the main focus of this article, we also explored whether people update their behavioral intentions in response to descriptive norms. Therefore, our first research question was whether participants would update their injunctive norm beliefs and behavioral intentions at the individual level in response to receiving novel descriptive norm information.

Do People Differentially Update Their Personal- and Other-Moral Judgments?

To help generate a clearer picture of the relationship between descriptive and injunctive norms and how they influence beliefs, it is also important to develop a better understanding of how they relate to another important feature of social norm cognition: moral judgments. Previous work suggests that moral judgments are closely tied to descriptive and injunctive norms. For example, in addition to finding a relationship between descriptive and injunctive norms, Eriksson and colleagues found that descriptive norm information influenced participants' moral judgments of characters engaging in those normative behaviors (2015). Other work has more directly explored how descriptive norms intersect with moral judgments of behavior. The "common is moral" heuristic is the hypothesis that the frequency or commonness of a behavior influences its perceived moral status (Lindström et al., 2018). In support of this theory, Lindström and colleagues found that both prosocial and selfish behaviors are evaluated as more moral when common than when rare, suggesting that we infer morality from the frequency of behaviors (2018). Taken together, this work indicates that there is a strong association between descriptive norms, injunctive norms, and moral judgments.

However, while past work suggests that morality is closely linked to normative beliefs¹ and injunctive norms are generally conceived of as moral (Lu et al., 2020; Russell et al., 2022), it is possible that morality and injunctive norms are dissociable, much like the relationship between descriptive and injunctive norms. Thus, while injunctive norms are often thought of as moral (Eriksson et al., 2015), it is possible for injunctive norms to not be perceived as morally good. For example, most people think you should eat dessert after dinner, yet many would also agree it is not necessarily morally wrong to not do so. Furthermore, moral judgments can be divided into two categories: first order (i.e., what I believe to be moral) and second order (i.e., what I think other people believe to be moral). It is possible that descriptive norm information may differentially influence these first- and second-order moral judgments. To date, relatively little work has examined the relationship between descriptive norms beliefs, injunctive norm beliefs, and first- and second-order moral judgments within one experimental design (but see White et al., 2009). Doing so will help us to better understand how descriptive norms shape injunctive norms and moral judgments and will shine light more generally on how these constructs are associated with one another and the social cognition underlying social norms. Thus, our second primary research question asked whether people would update their first-order moral judgments to a different extent than their second-order moral judgments in response to descriptive norm information.

Do People Update to Different Extents Depending on the Type of Behavior?

Another important yet unexplored question is whether the relationship between descriptive norms, injunctive norms, and moral judgments is consistent across different types of norms or whether it varies such that the relationship is stronger for certain kinds of behaviors. A growing body of work suggests there are important differences in cognition between different behavioral domains. For example, there are a number of differences in moral cognition between harm and purity domains—researchers have found an attribution asymmetry between these domains, such that people endorse more person-based attributions for impure acts compared to harmful ones (Chakroff & Young, 2015), while other work finds that people rely more on judgments of intent for harm violations than purity violations (Chakroff et al., 2016). Moreover, some norm theorists make a distinction between social norms and moral norms (Bicchieri, 2006, 2016; Elster, 1989). Namely, in Bicchieri's account of norms, social norms are categorically distinct from moral norms in that social norms are socially conditional behaviors that depend on others' empirical expectations (about what others commonly do, akin to descriptive norms) and normative expectations about how one should behave (akin to injunctive norms). In contrast, moral norms are internalized beliefs that we comply with regardless of what others do or expect of us. Bicchieri also distinguishes personal preferences from social norms as patterns of behavior that, like harm norms, are not socially conditional beliefs (however, unlike harm norms, they are not internalized). Together, this view of norms suggests that people should be less sensitive to information about others' expectations for harm-related norms and personal preferences than other kinds of norms such as fairness norms. This account of norms sits in contrast to other psychological research which does not make a meaningful distinction between moral and social norms and, in fact, often treats them as interchangeable concepts (Lindström et al., 2018; Nielsen & McGregor, 2013; Russell et al., 2022).

Additionally, research on social domain theory finds that people behave differently depending on whether a norm is conventional (e.g., wearing a school uniform) or moral (e.g., bullying and stealing someone's lunch money), such that people view conventional norm violations as less serious and deserving of punishment and view moral norms as more inalterable and objective (Smetana, 2006; Smetana et al., 2014). Because this body of work provides substantial evidence that we treat moral and conventional norms as psychologically distinct, we also wanted to compare belief updating between them to determine whether people are more sensitive to descriptive norm information for conventional than moral norms. Furthermore, some types of behavior fall somewhere between conventional and moral norms, serving as an interesting test of the social conditionality theory of social norms. One such behavior is fairness. Fairness is an interesting behavior to study in this context because while it is often viewed as a moral norm (Baumard et al., 2013; Curry et al., 2019; Folger, 1998), recent work suggests that children perceive fairness norms differently than harm-based norms, viewing harm-based norm violations as more serious (Yucel et al., 2022). Additionally, a large body of work has found substantial

¹ We use this term broadly to denote related beliefs relevant to social norms, including descriptive and injunctive norm beliefs, moral judgments, and behavioral intentions.

cross-cultural differences in fairness norms (Blake et al., 2014; Henrich et al., 2001; House et al., 2020), suggesting they are highly flexible and sensitive to local normative information. Put differently, including fairness behaviors in addition to other-moral norms like harm behaviors serves as an important boundary condition of the social conditionality account. Namely, if moral norms like harm norms are categorically distinct from social norms in that they are not dependent on others' expectations, then people should be significantly less sensitive to normative information for harm norms than fairness norms.

One way of determining whether people are less sensitive to others' expectations for certain kinds of normative content is to examine how sensitive their normative beliefs are to novel descriptive norms that provide information about others' expectations (e.g., if many people are complying with a norm, it suggests that they likely expect others to comply with it as well—i.e., there are both empirical and normative expectations). Therefore, our third primary research question asked whether people would be more sensitive to descriptive norm information for fairness and conventional behaviors than for harm behaviors and preferences. Specifically, we predicted that participants would be less likely to update their normative beliefs for harm norms and personal preferences, which are not conditioned on others' expectations, than their beliefs for fairness norms, which may be more sensitive to others' empirical and normative expectations. We were also interested in exploring whether people view fairness norms more similarly to conventional norms than harm norms or preferences. Consequently, we examined whether participants would update their beliefs for conventional norms to a similar degree as fairness norms and to a greater extent than harm norms or preferences, after receiving novel descriptive norm information.

The Present Studies

In the present set of studies, we explored belief updating from descriptive to injunctive norms at the individual level. First, we examined how specific information about behavioral frequency influences updating at the individual level by measuring pre- to postdescriptive norm manipulation changes. In line with this aim, we explored belief updating for injunctive norms as well as behavioral intentions. Second, to better understand the relationship between moral judgments and normative beliefs, we investigated whether people would update their moral judgments to different extents between first- and second-order moral judgments. Third, we examined whether individual updating was sensitive to the type of behavior in question in order to (a) test whether there is variability in the relationship between descriptive and injunctive norms and (b) test the social conditionality account of social norms which predicts that moral norms (like harm) and personal preferences should be less sensitive to others' normative and empirical expectations than fairness and conventional behaviors.

To investigate these questions, we presented participants with a series of vignettes detailing different normative behaviors. Specifically, we manipulated whether there was an uncommon descriptive norm, in which 20% of people in the vignette were doing the behavior, or a common descriptive norm, in which 80% of people were doing the behavior. To measure belief updating, we assessed beliefs before and after receiving information on the descriptive norm. We used four types of behaviors across the six studies, including conventional norms, that is, arbitrary, contextually dependent rules that coordinate individuals'

behavior (Smetana, 2013; e.g., talking in a library); two types of "moral" norms (Smetana & Ball, 2018; Turiel et al., 1987) related to fairness (e.g., cheating on a test) and harm (e.g., stealing a coworker's lunch); and personal preferences (e.g., wearing socks with sandals), "socially independent" behaviors that should be less sensitive to descriptive norms (Bicchieri, 2006). In analyses across studies, we examined whether the extent of belief updating after receiving descriptive norm information differed between these types of behavior.

If we find that people readily update their injunctive norm beliefs and moral judgments to similar extents after receiving descriptive norm information, this would provide some of the strongest evidence to date that descriptive norms, injunctive norms, and moral judgments are highly related concepts—possibly attesting to the extent to which they align in everyday life or tap into a latent, underlying psychological norm construct. Furthermore, if we find that people update their beliefs after receiving descriptive information, this would suggest that there is a common directional effect of descriptive norms on injunctive norm beliefs and moral judgments such that descriptive norm information plays an important role in the formation of injunctive norms and moral judgments. In other words, simply seeing that many people are engaging in a certain behavior might lead us to infer that most people approve of this behavior and think it is morally good, which may, in turn, influence our own beliefs and decision to engage in the behavior. Alternatively, if people do not update their beliefs, this would suggest that descriptive and injunctive norms and moral judgments are more distinct constructs than previously thought and indicate that descriptive norm information plays little to no role in shaping injunctive norm beliefs and moral judgments.

Lastly, if we find that participants update their beliefs less for harm behaviors and personal preferences than conventional and fairness behaviors, this would provide evidence in favor of Bicchieri's socially conditional account of norms (2006, 2016). Furthermore, if participants update their beliefs to different extents between fairness and harm behaviors, this would provide further evidence that people perceive them as meaningfully distinct and speak against the value of categorizing them as similarly moral norms. On the other hand, if participants update their beliefs to a similar extent between types of behaviors, particularly between fairness and harm behaviors, that would speak against the socially conditional account of norms and indicate that people are as sensitive to others' expectations for social norms as for moral norms.

Method

Overview of Studies

In a series of studies, we assessed whether descriptive norms influence injunctive and moral norm beliefs (pilot study) and whether people update their prior injunctive norm beliefs, moral judgments, and behavioral intentions after receiving novel information about a descriptive norm (Studies 1, 2, 3, 4, and 5). To do this, we measured participants' beliefs before (henceforth *priors*) and after presenting them with novel descriptive norm information, subtracting their postnorm ratings (henceforth, *posteriors*) from their priors to determine their extent of belief updating. In each study, we manipulated the descriptive norm such that it was either common (common descriptive norm) or uncommon (uncommon descriptive norm) to explore how the relative frequency of descriptive norm information influences beliefs. To explore whether updating varies depending on the type of behavior,

we manipulated the norm type across studies, such that participants received some combination of conventional, fairness, harm, and preference behaviors (Studies 1–5). We describe the method for each study in turn; however, because the experimental design was similar across studies, we only detail the new or changed features of subsequent studies. We report how we determined our sample size, all data exclusions, all manipulations, and all measures in the study.

Participants

Study 1

We tested 414 participants (52.42% female), aged 18–73 ($M = 41.65$) from Amazon Mechanical Turk in a preregistered study: <https://osf.io/e3c6n>. Participants earned \$1.80 for participating in a 15-min study. We initially recruited 480 participants, but 66 were excluded for failing one of our prespecified exclusion criteria (e.g., failing an attention check; see preregistration for details).

Study 2

We tested $N = 402$ participants (50.99% female), aged 18–85 ($M = 41.73$) from Amazon Mechanical Turk in a preregistered study: <https://osf.io/gua3q>. The sample size was determined using an observed power simulation on the Study 1 data, which found that we would have 100% power to detect the same size interaction effect between condition and vignette type as observed in Study 1. Participants earned \$3 for participating in a 25-min study. We initially recruited 438 participants, but 36 were excluded for failing our inclusion criteria.

Study 3

We tested $N = 643$ participants (53.81% female) on Mechanical Turk aged 18–80 ($M = 41.08$) in a preregistered study: <https://osf.io/c6r8p>. The sample size was determined using a power simulation that found we would have 100% power to detect an effect of the interaction between condition and vignette type given the effect size observed in Study 2. Participants were compensated \$0.60 for participating in a 5-min survey. We initially recruited 702 participants, but 59 were excluded for failing our inclusion criteria.

Study 4

We tested $N = 416$ participants (60.8% female), aged 20–83 ($M = 43.4$) from Amazon Mechanical Turk in a preregistered study: <https://osf.io/5pz8m>. The sample size was determined using a power simulation using the data structure and effect sizes from Study 3. The power simulation found that we would have 98.64% power to detect an effect of the interaction between norm condition (uncommon, common) and vignette type (fair, harm). Participants earned \$1.70 for participating in a 14-min study. We initially recruited 437 participants, but 25 were excluded for failing our inclusion criteria.

Study 5

We tested $N = 400$ participants (54.3% female), aged 18–78 ($M = 41.33$) from Amazon Mechanical Turk in a preregistered study: <https://osf.io/pdnyx>. The sample size was determined using a power simulation with the data structure and effect sizes from Study 3. The

power simulation found that we would have 98.64% power to detect an effect of the interaction between norm condition (uncommon, common) and vignette type (fairness, harm). Participants earned \$3.15 for participating in a 26-min study. We initially recruited 449 participants, but 49 were excluded for failing our inclusion criteria.

Design and Procedure

Study 1

Participants provided informed consent and were given instructions for the task. Participants saw a series of vignettes that detailed short scenarios involving different normative behaviors (see *Vignettes* in the Supplemental Material), receiving a random subset of six out of 12 total vignettes in a random order. Of the 12 vignettes, six vignettes were negatively emotionally valenced, while six vignettes were positively emotionally valenced. All of the vignettes in the studies reported here were normed on Mechanical Turk prior to their inclusion to ensure that they were perceived as being similar across a range of dimensions including frequency, cost, benefit, etc. (see Supplemental Material for our norming procedure and criteria). In this and all subsequent studies, we varied the type of norm within subject (with the exception of Study 3, see below) to understand whether updating differs depending on the type of behavior as well as to test whether descriptive norms influence updating across a range of different types of social contexts and behaviors.

Participants were randomly assigned to either the uncommon or the common descriptive norm condition in which they received information about how common the behavior detailed in the vignette was. In the uncommon descriptive norm condition, participants were told that 20% of people in the scenario (either 1/5, 2/10, or 4/20) were engaging in the relevant behavior. In the common norm condition, participants were told that 80% (4/5, 8/10, or 16/20) of people were engaging in the relevant behavior. While the proportion of people engaging in the behavior remained constant for each descriptive norm condition, the actual number of people varied in order to make the vignettes sound more realistic. We selected these proportions as a test of whether a relatively common descriptive norm (where the behavior is commonly done) influences beliefs more than a relatively uncommon descriptive norm (where the behavior is uncommonly done).

Participants rated the vignettes on all measures before receiving the descriptive norm information and again after receiving it. Participants saw six vignettes and rated each on several dependent variables (DVs). This study included 12 vignettes in total: the six vignettes from the pilot study (henceforth, *conventional* vignettes) and six new vignettes featuring behaviors related to fairness (*fairness* vignettes). Unlike the pilot study, participants first saw the vignettes without any descriptive norm information and answered the dependent variables in order to assess their prior beliefs. After assessing their priors, participants completed a short filler task consisting of 10 trivia questions in order to mask task demands. Participants then received the same six vignettes they saw previously in a randomized order, this time including either the uncommon or common descriptive norm information, depending on the condition they were assigned to, and answered the dependent measures for a second time. Participants rated the behavior in each vignette on its injunctive normativity, descriptive normativity, and their intentions of engaging in the behavior on a 100-point sliding scale and answered a binary behavioral decision and an open response

question about their behavioral choice. Additionally, this study assessed morality with two questions: the first assessed first-order moral judgments—how moral they personally think the behavior is (*personal*-moral beliefs)—and the second assessed second-order moral beliefs about others' moral judgments, how moral they think other people think the behavior is (*other*-moral beliefs). We included these two different morality questions in order to explore whether first- and second-order moral beliefs dissociate—that is, whether people update their moral judgments more for one than the other in response to descriptive norms—as well as whether one is more strongly associated with normative beliefs.

Study 2

In Study 2, we included a greater range of behaviors to more rigorously test whether people update their beliefs after receiving descriptive norm information as well as to determine whether updating varies depending on the type of normative behavior in question.

Participants were randomly assigned to one of two descriptive norm conditions (uncommon, common) between subjects and saw three vignette types within subjects (fairness, harm, preference). For each vignette type, participants saw four vignettes—two positively valenced and two negatively valenced. In this study, we replaced our conventional norm vignettes with vignettes pertaining to harm and preferences, two behaviors that are less socially conditional (Bicchieri, 2006, 2016) and thus should be less sensitive to novel descriptive norm information. The procedure for Study 2 was otherwise identical to that of Study 1.

Study 3

Our goal in Study 3 was to replicate the findings from Study 2 using a between-subjects design to rule out the possibility of demand characteristics from treating norm type as a within-subject variable. Because responses were near ceiling for the positively valenced vignettes in Studies 1 and 2, we only included the negatively valenced behaviors this study onward (see *Individual Study Analyses* in the Supplemental Material for the findings from the positively valenced vignettes).

Participants were randomly assigned to one of six conditions: one of two descriptive norm conditions (uncommon, common) and—unlike the previous studies—one of three vignette types (fairness, harm, preferences) between subjects. Participants saw two vignettes in total. The procedure of Study 3 was identical to Study 2, and participants completed the same survey with the same measures as in previous studies.

Study 4

In Study 4, we wanted to test whether our injunctive norm updating results were robust to differences in the specific wording of the measure of this construct. Because previous work has operationalized injunctive norms in several different ways (Bicchieri, 2016; Cialdini et al., 1990; Lu et al., 2020; Russell et al., 2022), we replaced the previous injunctive measure with a new question assessing the extent to which participants believe that other people think you *should* engage in the behavior. Additionally, we returned to

treating vignette type as a within-subjects variable in this and all remaining studies.

Participants were assigned to one of two descriptive norm conditions (uncommon, common) and saw all three vignette types (fairness, harm, preference) within subjects. The procedure was identical to that of Study 3, and participants completed the same measures with the exception of a new injunctive measure and the inclusion of a punishment measure (see *Individual Study Analyses* in the Supplemental Material for more information and results). This injunctive dependent variable was framed in terms of what one *should* do, rather than what is *approved of*—another common definition of injunctive norm (“to what extent do other people think you should *X*”).

Study 5

We had three goals for Study 5: We aimed to (a) replicate our findings using a new set of harm behaviors and both injunctive norm measures in one study, (b) validate that participants perceived our norm types as actually falling into the hypothesized types, and (c) explore whether the vignette rating measures predicted updating (or the lack thereof) in the norm updating task (see Study 5 in *Individual Study Analyses* in the Supplemental Material for more information and results). To that end, we included all four sets of vignettes used in earlier studies plus a new set of harm vignettes (“fairness,” “convention,” “harm,” “psychological harm,” “preference”; see Table 1), both injunctive measures as dependent variables (injunctive-approve, injunctive-should; see Table 2 for measure text), and a new vignette rating task.

The design was nearly identical to that of the previous studies, but in this study, participants saw all five vignette types (fairness, harm, psychological harm, preference, conventional) within subjects. The procedure was similar to that of Study 4. Participants completed the same measures, excluding the punishment measure and including the new injunctive-should measure, in addition to the old injunctive measure, and a new set of two harm vignettes. Because the previously used harm vignettes could be perceived as secondarily tapping fairness concerns (e.g., it is both harmful and unfair to steal from someone), we designed a new set of harm behaviors to more exclusively tap harm. While the previous harm vignettes concerned stealing, these new harm vignettes focused on psychological and emotional harm—mocking someone in front of their face and starting a nasty rumor about someone. These new harm vignettes (henceforth *psychological harm* vignettes) went through the same norming criteria as the other vignettes.

Transparency and Openness

We follow the Journal Article Reporting Standards (Appelbaum et al., 2018) and report how we determined our sample size, all data exclusions, and all measures in the study. All of our studies' designs and analysis plans were preregistered. All of our analysis code, data, preregistrations, and Supplemental Material, including our study materials and all preregistered analyses, are available on the Open Science Framework (<https://osf.io/sc842/>; Deutchman et al., 2024). Data was analyzed using R, Version 4.3.0 (R Core Team, 2023), and the packages lme4 (Bates et al., 2015), meta (Balduzzi et al., 2019), and dplyr (Wickham et al., 2023).

Table 1*Table With Design Information on All Studies*

Study	Design	Vignette type	Participant
Pilot study	2 (descriptive norm: uncommon, common) between subject	Conventional	$N = 401$ participants ($M_{\text{age}} = 40.1$, $F = 46.6\%$)
Study 1	2 (descriptive norm: uncommon, common) between subject \times 2 (vignette type) within subject	Fairness, conventional	$N = 414$ participants ($M_{\text{age}} = 41.7$, $F = 52.4\%$)
Study 2	2 (descriptive norm: uncommon, common) between subject \times 3 (vignette type) within subject	Fairness, harm, preference	$N = 402$ participants ($M_{\text{age}} = 41.7$, $F = 50.9\%$)
Study 3	2 (descriptive norm: uncommon, common) between subject \times 3 (vignette type) between subject	Fairness, harm, preference	$N = 643$ participants ($M_{\text{age}} = 41.1$, $F = 53.8\%$)
Study 4	2 (descriptive norm: uncommon, common) between subject \times 3 (vignette type) within subject	Fairness, harm, preference	$N = 411$ participants ($M_{\text{age}} = 43.0$, $F = 60.8\%$)
Study 5	2 (descriptive norm: uncommon, common) between subject \times 5 (vignette type) within subject	Conventional, fairness, harm, psychological harm, preference	$N = 400$ participants ($M_{\text{age}} = 41.3$, $F = 54.3\%$)

Results

Analytic Approach

In each individual study, we conducted a series of preregistered linear mixed effects regression models predicting the dependent measures by descriptive norm condition using the lme4 package in R (Bates et al., 2015; R Core Team, 2023). In all studies (with the exception of the pilot study), the models included the dependent measure difference score, which was found by subtracting the prior ratings from the posterior ratings after receiving the descriptive norm information (the pilot study did not collect prior ratings, and so we could not compute difference scores). For example, if a participant rated their likelihood of talking in the library as 30 prior to receiving the descriptive norm condition and then rated it as 50 after receiving information that the behavior is common, their behavioral intention updating score would be +20. Due to the design differences between the pilot study and the other studies, we only report the belief updating results of Studies 1–5 here and report the results of the pilot in the Supplemental Material (but note here the results from this study are consistent with the other studies). We present all results for each individual study in detail in the Supplemental Material (see *Individual Study Analyses*) as well as a

summary of model results collated across studies (see *Collated Individual Study Results*). See Table 3 for a correlation matrix between the dependent measures collapsed across study. We summarize the key results—the main effect of descriptive norm condition on updating—in Table 4. Lastly, we note that we report all studies collected for this project in the main text and Supplemental Material—the reported studies are the only ones we conducted on belief updating and thus there are no file-drawer concerns.

After conducting analyses across all of our individual studies, we found that the effect size for updating varied considerably across studies, perhaps reflecting variability in the study design and samples between studies. In order to evaluate the totality of evidence, we decided to pool our data and run a cross-study analysis. While not initially preregistered, taking a meta-analytic approach offered a stronger test of whether people update their beliefs after receiving descriptive norms by increasing the number of observations and statistical power. To that end, we conducted a series of meta-analyses across our five studies that assessed belief updating (Studies 1–5) following the method described by Harrer et al. (2021) and using the meta package in R (Balduzzi et al., 2019). We conducted a meta-analytic test for each of our main dependent measures—injunctive normativity, personal-morality, other-morality, and behavioral

Table 2*Table With the Measures From All Studies*

Measure	Measure text	Study
Injunctive normativity:		
1. Approval	1. In general, would most people approve of X? (0–definitely not approve, 100–definitely approve)	1. Studies 1–3, 5
2. Should do	2. To what extent do other people think you should X? (0–not at all, 100–a great deal)	2. Studies 4, 5
Injunctive certainty	Previously you were asked the question. You answered [injunctive norm response]. How certain are you in that response? (0–extremely uncertain, 100–extremely certain)	Studies 2–5
Descriptive normativity	In general, how many people would X? (0–no one, 100–everyone)	Studies 1–5
Morality	1. How moral do you personally think it is for someone to X? 2. How moral do you think other people think it is for someone to X? (0–extremely immoral to 100–extremely moral)	Studies 1–5
Behavioral intentions	How likely would you be to X? (0–extremely unlikely, 100–extremely likely)	Studies 1–5

Note. All measures from Studies 1–5 were on a 100-point sliding scale.

Table 3

Correlation Matrix With Pearson Correlation Coefficients for the Dependent Measure Ratings Prior to Receiving the Descriptive Norm Manipulation Collapsed Across Studies

Variable	1	2	3	4	5	6	7
1. Injunctive-approval	—						
2. Injunctive-should	.79***	—					
3. Injunctive certainty	-.20***	-.18***	—				
4. Morality-other	.67***	.66***	-.18***	—			
5. Morality-self	.63***	.61***	-.16***	.90***	—		
6. Behavioral intention	.60***	.71***	-.13***	.45***	.48***	—	
7. Descriptive	.64***	.73***	-.17***	.40***	.35***	.67***	—

*** $p < .001$.

intentions. Specifically, we looked at the effect of updating by comparing ratings for the dependent measures prior to receiving the descriptive norm to after receiving either the uncommon or common descriptive norm. Additionally, we also ran a set of meta-analyses for the four key measures comparing the degree of updating between the uncommon and the common descriptive norm conditions. Thus, we conducted three models for each dependent measure: one looking at the effect of updating after receiving the common descriptive norm, another looking at updating after receiving the uncommon descriptive norm, and a third comparing the extent of updating between the uncommon and common descriptive norm conditions. Note that this last set of meta-analyses used the absolute value of the updating difference scores in order to solely assess the overall magnitude of updating differences between descriptive norm conditions (i.e., regardless of the direction of the updating). Thus, these comparison analyses use a different, transformed data set to answer a distinct research question from the other meta-analyses that assess both the magnitude and directionality of updating.² For these and all cross-study results present below, we only included the negatively valenced behaviors because we observed ceiling effects for the positively valenced behaviors in Studies 1 and 2 (see *Valence Models* in the Supplemental Material for an analysis of behavior valence on updating).

While our studies had similar demographic information and presumably sampled from the same population (see Table 1 for age and gender breakdowns between studies), we observed statistically significant between-study heterogeneity, and so we used random-effect models in our meta-analyses to account for between-study variability in effect sizes. All models used the restricted maximum-likelihood estimator to calculate τ^2 and Knapp–Hartung adjustments to control for between-study heterogeneity. For determining effect sizes, we used Hedge’s g to correct for small sample bias in calculating the standardized mean difference. We replicate the results of the meta-analyses using a series of linear mixed effect models treating study identity—along with participant and vignette identity—as random effects. We report the results of these models in the Supplemental Material but note that their results are consistent with the results of the meta-analyses reported here (see *Multilevel Regression Models* in the Supplemental Material).

To investigate whether participants updated their beliefs to a lesser extent for harm and preference behaviors than fairness and conventional behaviors, we ran a series of linear mixed effect regression models predicting our four key dependent measures (injunctive, personal-morality, other-morality, behavioral intention)

and that included participant ID, vignette ID, and study number as random effects. We ran four models predicting belief updating by the interaction between descriptive norm condition (uncommon, common) and vignette type (fairness, conventional, harm, preference) to examine whether the direction and size difference in the effect of updating between uncommon and common descriptive norm conditions varied across behaviors. We also report a series of exploratory models subsetting by descriptive norm condition (common, uncommon) to examine the effect of updating between behavior type after receiving either a common or uncommon descriptive norm in the Supplemental Material (see *Multilevel Regression Models*). For all models, we set the fairness behaviors as the reference level for the vignette type variable and make between vignette pairwise comparisons using estimated marginal means adjusted using the multivariate t method to correct for multiple comparisons.

To examine whether participants’ prior descriptive norm beliefs influenced the extent of their updating and to rule out potential alternative explanations such a number matching effect, we conducted several moderation analyses. Specifically, we ran four models, one for each main dependent variable (injunctive, behavioral intentions, self- and other-morality). For all models, we included prior descriptive beliefs as a moderator of belief updating by creating linear mixed effect models predicting the dependent measures by the interaction between the prior descriptive beliefs (moderator) and descriptive norm condition. For these models, we included data from both the negative and positive vignettes.

Lastly, because our dependent measures were all moderately correlated with one another, we wanted to investigate whether our updating results were robust to treating those measures as interrelated, dependent measures as opposed to assuming independence, as in our univariate regression models. Consequently, we ran a series of exploratory multivariate regression models including all four of our main dependent measures together (injunctive, personal-morality, other-morality, behavioral intention). The first model examined updating after receiving the common descriptive norm and predicted the dependent measures by the descriptive manipulation (prior, post). The second model was identical to the first but focused on updating after receiving the uncommon descriptive norm. Lastly, the third model predicted the dependent measure difference score by descriptive norm condition (uncommon, common). We report the results of these

² Consequently, estimates from these two different sets of meta-analyses cannot be directly compared to one another.

Table 4*Model Estimates of the Effect of Descriptive Norm Condition on Belief Updating*

Dependent variable	Study 1	Study 2	Study 3	Study 4	Study 5
Injunctive-approval	$B = 25.70^{***}$	$B = 26.92^{***}$	$B = 23.79^{***}$		$B = 30.27^{***}$
Injunctive-should				$B = 29.81^{***}$	$B = 27.81^{***}$
Behavioral intention	$B = 9.34^{***}$	$B = 4.76^{***}$	$B = 5.61^{***}$	$B = 3.93^{***}$	$B = 5.92^{***}$
Morality-self	$B = 6.22^{***}$	$B = 2.74^{***}$	$B = 2.42^{**}$	$B = 1.41^*$	$B = 2.06^{**}$
Morality-other	$B = 18.66^{***}$	$B = 15.49^{***}$	$B = 15.96^{***}$	$B = 16.26^{***}$	$B = 21.04^{***}$

Note. Table showing unstandardized model estimates and significance levels of the descriptive norm condition term (uncommon, common) for our models predicting each of our five key dependent measures by descriptive norm condition across Studies 1–5. All models included the difference score for each of the dependent variables.

* $p < .05$. ** $p < .01$. *** $p < .001$.

models in the Supplemental Material but note here that they largely replicate the updating results from our univariate models.

Do People Update Their Beliefs After Receiving Descriptive Norm Information?

When looking at the effect of either uncommon or common descriptive norm information on belief updating across studies, we find support for our prediction that people update their injunctive norm beliefs after receiving information about the descriptive norm (Figure 1). There was a significant effect of descriptive norm condition on updating ($d = 0.81$, 95% CI [0.63, 1.00], $p < .001$), such that people positively updated their injunctive beliefs about how approved the norm is after receiving information that there is a common descriptive norm (see Figure 2 for a forest plot of the injunctive meta-analysis). Similarly, we find a significant meta-analytic effect for the uncommon descriptive norm condition but in the opposite direction: People negatively updated their injunctive beliefs after receiving information that the behavior was uncommon, finding the behaviors less approved of ($d = -0.16$, 95% CI [-0.27, -0.06], $p = .01$).

Next, we examined the meta-analytic effect of belief updating for behavioral intentions and found support for our prediction that people update their behavioral intentions after receiving a descriptive norm. Namely, there was a significant effect of the common descriptive norm condition on behavioral intentions ($d = 0.08$, 95% CI [0.02, 0.14], $p = .021$), such that participants reported they would be more likely to engage in the behavior after finding out it was common. We also found a significant effect of the uncommon descriptive norm condition on behavioral intentions ($d = -0.12$, 95% CI [-0.15, -0.09], $p < .001$), such that people were significantly less likely to say they would engage in the behavior after finding out it was uncommon.

When comparing the magnitude of updating between the uncommon and common descriptive norm conditions across studies, we found evidence that people updated their injunctive norm beliefs to different extents depending on the frequency of descriptive norm information. Namely, there was a significant effect of descriptive norm condition on injunctive norm updating ($d = 0.69$, 95% CI [0.47, 0.93], $p = .001$), such that people updated their beliefs more for the common descriptive norm than the uncommon descriptive norm. Turning to the meta-analytic effect of belief updating for behavioral intentions, we found that participants updated their behavioral intentions to a larger extent across studies after receiving a common descriptive norm compared

to an uncommon descriptive norm ($d = 0.13$, 95% CI [0.08, 0.17], $p = .002$).

Do People Differentially Update Their Personal- and Other-Moral Judgments?

Next we examined the meta-analytic effect of belief updating on other-morality—second-order beliefs about how moral other people think the behavior is. We find some support for our prediction that people update their second-order moral beliefs after receiving descriptive norms. Namely, there was a significant effect of the common descriptive norm on other-morality beliefs ($d = 0.47$, 95% CI [0.39, 0.53], $p < .001$), such that people thought that others would think that the behavior is more moral after finding out it was commonly done. However, we did not find an effect of the uncommon descriptive norm condition on updating of other-morality beliefs ($d = -0.11$, 95% CI [-0.23, 0.02], $p = .075$).

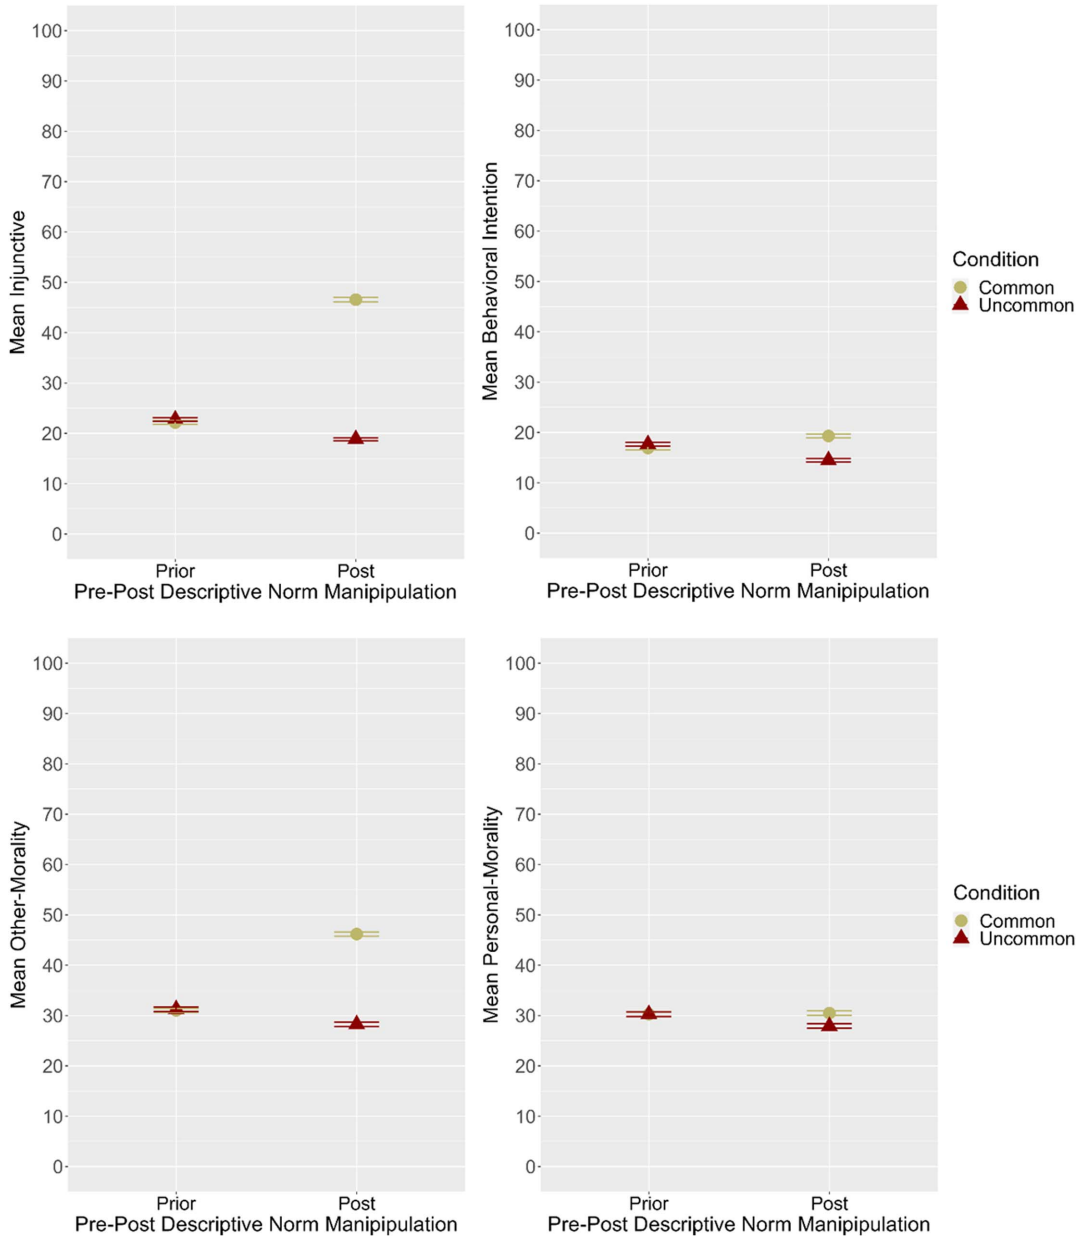
When looking at the meta-analytic effect of belief updating on personal-morality—personal beliefs about how moral the behavior is—we find mixed evidence in support of our belief updating prediction. There was no significant effect of the common descriptive norm condition on personal-morality beliefs ($d = 0.007$, 95% CI [-0.03, 0.04], $p = .638$). However, we find a significant, if small, effect of the uncommon descriptive norm condition on personal-morality beliefs ($d = -0.08$, 95% CI [-0.13, -0.02], $p = .02$), such that participants personally thought the behaviors were less moral after finding out that they were uncommon.

Lastly, when comparing the magnitude of updating between the descriptive norm conditions, we found a significant updating effect across studies for other-moral judgments: Participants were significantly more likely to update their second-order moral beliefs after receiving a common descriptive norm than an uncommon descriptive norm ($d = 0.54$, 95% CI [0.42, 0.65], $p < .001$). When examining the meta-analytic effect of belief updating of personal-moral beliefs about the behavior, we find that participants were significantly more likely to update their personal-moral beliefs after receiving the uncommon than the common descriptive norm ($d = 0.05$, 95% CI [0.001, 0.09], $p = .047$).

Do Participants Update to Different Extents Depending on the Type of Behavior?

We first compared the size and direction of updating between the common and uncommon descriptive norm manipulation to test

Figure 1
Belief Updating Ratings for the Injunctive Norm, Behavioral Intention, Other-Morality, and Personal-Morality Measures




Note. Plot comparing belief ratings for the uncommon and common descriptive norm conditions in comparison to prior beliefs for the key dependent measures, injunctive norm beliefs (approve), behavioral intention, personal-moral beliefs, other-moral beliefs. Error bars show standard error. This figure collapses across data from Studies 1–5 and only includes the negatively valenced vignette results. See the online article for the color version of this figure.

whether participants were more sensitive to descriptive norms for conventional and fairness behaviors than harm and preferences. We found that the difference in injunctive updating between the common and uncommon descriptive norm manipulations was larger for fairness than harm behaviors ($B = 2.97, SE = 1.28, p = .021$), smaller for fairness than conventional behaviors ($B = 3.34, SE = 1.61, p = .037$), and was similar for fairness and preference behaviors ($B = 0.46,$

$SE = 1.28, p = .73$). The difference in updating between descriptive norm conditions was significantly larger for conventional than harm behaviors ($B = 6.31, SE = 1.66, p < .001$), but was not significantly different between conventional and preference behaviors ($B = 3.80, SE = 1.66, p = .09$) nor between harm behaviors and preferences ($B = -2.51, SE = 1.33, p = .23$). Participants positively updated their injunctive beliefs in response to the common descriptive norm and

Figure 2

A Forest Plot of the Meta-Analytic Effects of the Common Descriptive Norm Condition on Injunctive Norm Ratings



Note. CI = confidence interval; SMD = standardized mean difference.

negatively updated in response to the uncommon descriptive norm for all behaviors except the harm behaviors, in which they updated for the common but not uncommon descriptive norm. The difference in behavioral intention updating between the common and uncommon descriptive norm was significantly larger for fairness behaviors than harm behaviors ($B = 5.65$, $SE = 0.92$, $p < .001$), smaller for fairness than conventional behaviors ($B = -5.76$, $SE = 1.15$, $p < .001$), and similar for fairness and preference behaviors ($B = -0.51$, $SE = 0.92$, $p = .58$). The difference in updating between descriptive norm manipulations was greater for conventional than harm behaviors ($B = 11.40$, $SE = 1.19$, $p < .001$) and preferences ($B = 5.24$, $SE = 1.19$, $p < .001$) and for preferences than harm behaviors ($B = 6.16$, $SE = 0.96$, $p < .001$). Unlike injunctive beliefs, participants positively updated their behavioral intentions for the common descriptive norm and negatively updated for the uncommon descriptive norm across all behaviors.

For personal-moral beliefs, participants updated to a larger extent between the common and uncommon descriptive conditions for the fairness behaviors than harm ($B = 3.27$, $SE = 0.77$, $p < .001$) and preference behaviors ($B = 1.86$, $SE = 0.78$, $p = .02$) but less for fairness than for the conventional behaviors ($B = -2.99$, $SE = 0.97$, $p = .002$). The difference between descriptive norm conditions in personal-morality updating was also larger for conventional than harm ($B = 6.26$, $SE = 0.99$, $p < .001$) and preference behaviors ($B = 4.85$, $SE = 1.00$, $p < .001$), while updating differences were similar for harm behaviors and preferences ($B = -1.41$, $SE = 0.81$, $p = .30$). Participants positively updated their personal-moral beliefs in response to the common descriptive norm and negatively updated in response to the uncommon descriptive norm for the fairness and conventional behaviors. In contrast, they negatively updated their beliefs in response to both norm conditions for the harm and preference behaviors, updating to a larger extent for the uncommon than common descriptive norm. Lastly, participants updated their second-order moral beliefs to a larger extent between the common and uncommon descriptive norm conditions for fairness behaviors than preference behaviors ($B = 11.22$, $SE = 1.14$, $p < .001$) but updated to a similar extent between fairness and conventional ($B = -1.19$, $SE = 1.43$, $p = .40$) and harm behaviors ($B = 0.25$, $SE = 1.14$, $p = .82$). The difference between descriptive norm conditions was also larger for conventional behaviors than for preferences ($B = 12.42$, $SE = 1.47$, $p < .001$) and for harm than for preference behaviors ($B = 10.97$, $SE = 1.19$, $p < .001$) but was not

significantly different between conventional and harm behaviors ($B = 1.45$, $SE = 1.47$, $p = .759$). Participants positively updated their other-morality beliefs in response to a common descriptive norm and negatively updated in response to an uncommon descriptive norm for all behaviors.

Do Prior Descriptive Beliefs Moderate the Effect of Descriptive Norm Information on Belief Updating?

Participants' beliefs about the descriptive normativity of the behaviors before receiving the descriptive norm manipulation moderated the effect of descriptive norm information on injunctive belief updating ($B = -2.77$, $SE = 0.39$, $p < .001$), such that the stronger their premanipulation descriptive beliefs, the less likely they were to update their injunctive beliefs in response to a common descriptive norm and the more likely they were to negatively update in response to an uncommon descriptive norm. We found that prior descriptive beliefs moderated the effect of descriptive norm information on behavioral intention updating ($B = 1.84$, $SE = 0.29$, $p < .001$), such that the higher the prior descriptive beliefs, the less likely participants were to update their beliefs after receiving a common descriptive norm and the more likely they were to negatively update after receiving an uncommon descriptive norm. Prior descriptive beliefs also moderated the effect of descriptive norms on personal-morality ($B = -0.79$, $SE = 0.24$, $p = .001$): The higher participants' prior descriptive beliefs, the more likely they were to negatively update in response to an uncommon descriptive norm—finding the behavior less personally moral—while there was no change in updating in response to the common descriptive norm. Lastly, prior descriptive beliefs also moderated the effect of the descriptive norm manipulation on other-morality in a similar manner to injunctive beliefs ($B = 1.78$, $SE = 0.36$, $p < .001$). Namely, participants with lower prior descriptive beliefs updated their second-order moral judgments after receiving a common but not uncommon descriptive norm, while participants with higher prior descriptive beliefs negatively updated their moral judgments after receiving an uncommon but not common descriptive norm.

General Discussion

Across six studies, we explored the relationship between descriptive norms, injunctive norms, moral judgments, and behavioral intentions.

Specifically, we assessed participants' beliefs about normative behaviors both before and after receiving information that there was either a common or uncommon descriptive norm. Across all of these updating studies, we found that people updated their beliefs about the injunctive normativity of a behavior after receiving descriptive norm information. When there was a descriptive norm that the behavior was common, participants thought the behavior was more approved of than before receiving the norm information. When receiving a descriptive norm that the behavior was uncommon, participants thought the behavior was less approved of than they initially believed. Similarly, we also found that participants updated their other-moral beliefs after receiving descriptive norm information, such that they thought that others would think the behavior was more moral after finding out it was common, and less moral after finding out it was uncommon. Yet, we found somewhat ambiguous evidence that participants updated their personal-moral beliefs—they did not significantly update their beliefs in response to the common descriptive norm but did do so in response to the uncommon descriptive norm, although this effect was small. Thus, participants' other-moral beliefs were more sensitive to descriptive norms—and thus more similar to injunctive norm beliefs—than their personal-moral beliefs which were somewhat insensitive to descriptive norm information.

We additionally found a significant but relatively weak pattern of updating results for behavioral intentions. Participants updated their behavioral intentions, such that they reported they were more likely to engage in the behavior after receiving a common descriptive norm (three out of five studies) and less likely to after receiving an uncommon descriptive norm (five out of five studies). In our meta-analyses, we found a small but statistically significant effect of the common descriptive norm in the predicted direction ($d = 0.08$), as well as a statistically significant effect of the uncommon descriptive norm ($d = -0.12$), supporting our prediction that participants would negatively update their behavioral intentions after receiving information that the behavior was uncommon.

Across all studies and dependent measures, we found that, when looking solely at the magnitude of updating, participants updated their beliefs to a greater extent after receiving a common descriptive norm than an uncommon descriptive norm, suggesting that the effect of common descriptive norm information was stronger than that of uncommon descriptive norm information. All together, our results support previous work documenting a common association between descriptive norms, injunctive norms, and moral judgments (Eriksson et al., 2015; Lindström et al., 2018).

Do People Update Their Beliefs After Receiving Descriptive Norm Information?

One of the central goals of this research was to investigate how, and to what extent, people update their injunctive norm beliefs and moral judgments after receiving descriptive norm information. While previous work has found that people make simple, bidirectional inferences between descriptive and injunctive norms (e.g., is a behavior injunctive or not given that it is common or uncommon; Eriksson et al., 2015), no work to our knowledge has examined how people change their beliefs in response to novel normative information. Our study builds on and extends past work by assessing beliefs before and after receiving descriptive norm information to explore how people update their own beliefs about the injunctive normativity and morality of a behavior in the face of novel descriptive

norm information. We show here that not only do people positively update their beliefs after receiving a common descriptive norm that a behavior is common but that people also *negatively* update their beliefs after receiving an uncommon descriptive norm that a behavior is relatively rare. That people update their beliefs in either direction, finding the norm to be more or less injunctive after receiving descriptive norm information, suggests that we are highly sensitive to various frequencies of descriptive norm information—we make rich inferences not only from what we see others commonly do, but also from what most do *not* do. In other words, this novel finding indicates that we rely on descriptive norms not only to infer what others approve of, but that we also rely on that information—or lack thereof—to infer what others do not approve of. These findings extend our understanding of how descriptive norms influence our injunctive norm beliefs and moral judgments, demonstrating the rich social inferences we make from the presence and absence of descriptive norms. Additionally, we believe our pre-post norm design will prove useful for other researchers interested in understanding the relationship between different kinds of normative information. Future work should investigate this question further by titrating the level of descriptive norm information required for individuals to update their injunctive and moral beliefs and examine how this might vary across different behaviors.

While past research has explored how descriptive norms influence injunctive beliefs, little work has explored how descriptive norms influence other relevant normative beliefs. By examining descriptive beliefs, injunctive beliefs, moral judgments, and behavioral intentions within one design, our results offer the clearest evidence to date of how descriptive norms shape normative beliefs and inform our understanding of the relationship between descriptive and injunctive norms, moral judgments, and behavioral intentions. Namely, our results suggest that injunctive norm beliefs, personal and other-moral judgments, and behavioral intentions are moderately to strongly correlated with one another and suggest that people make rich social inferences from descriptive norms.

However, while these measures are interrelated, that people updated their injunctive beliefs and other-moral judgments more readily and to a larger extent than personal-moral judgments and behavioral intentions suggests that these items might be measuring at least somewhat partially independent constructs. In other words, while injunctive beliefs influence personal-moral beliefs, they are not necessarily measuring the same underlying construct.

Additionally, it is important to consider how our findings pertain to the sizeable body of work in social psychology on conformity. How we conform to the behavior of others is one of the foundational areas of research in social psychology, demonstrating that our behavior is highly sensitive to what the majority of people around us are doing (Asch, 1956; te Velde & Louis, 2022). Conformity, and the related but conceptually distinct conformist bias, has been proposed to be two of the more important pathways of social transmission in our species, playing a vital role in cultural evolution (Mesoudi, 2009; Whiten, 2019; but see Eriksson et al., 2007). Yet, there remains disagreement in the field regarding the exact nature in which descriptive norms exert social influence. Our finding that descriptive norm information influences people's intentions to engage in the relevant behavior is aligned with findings from previous work demonstrating that descriptive norms and conformity play an important role in influencing behavior. Namely, that the effect of descriptive norms influenced injunctive and moral beliefs in addition to behavioral intentions—in conjunction with our finding

that injunctive ratings partially mediated the effect of descriptive norms on behavioral intentions (see *Mediation Models* in Supplemental Material)—suggests that descriptive norms influence behavior at least in part because of the additional normative and moral information they convey. This finding is aligned with past work demonstrating that people tend to infer injunctive norms and moral judgments from descriptive norms (Eriksson & Strimling, 2015; Eriksson et al., 2015; Lindström et al., 2018). Together, these findings suggest that, while conformity is often thought of as solely a result of copying the majority behavior, an additional component of its influence on behavior might be due to the rich social inferences we implicitly make from majority behavior—in general, if most people are doing something, then they likely approve of doing it and think others should too. Additionally, the predictiveness of descriptive norms on behavior varies depending on a number of factors including reputational concerns, uncertainty, and power asymmetries (Gelfand & Harrington, 2015). Our findings contribute to the notion that the effect of descriptive norms on behavior is sensitive to social and contextual factors. Namely, our finding that participants updated their behavioral intentions to different extents after receiving descriptive norm information depending on the behavior indicates that the effectiveness of descriptive norms also varies depending on the specific type of behavior in question.

Do People Differentially Update Their Personal- and Other-Moral Judgments?

While often thought of as distinct constructs, recent research has documented a common relationship between descriptive and injunctive norms and moral judgments, with some work even defining injunctive norms in terms of morality (Eriksson et al., 2015; Lindström et al., 2018; Nielsen & McGregor, 2013; Russell et al., 2022). Thus, a secondary goal of this project was to better understand the relationship between descriptive norms, injunctive norms, and other-moral judgments—that is, how exactly do these concepts relate to one another? Are injunctive norms always moral or are they partially independent, measuring related but distinct constructs, such that behaviors can be injunctive but not moral or moral but not injunctive? We were also interested in deconstructing the relationship between injunctive norm beliefs and morality by examining how injunctive beliefs relate to both personal and other-moral judgments. Because injunctive norm beliefs are second-order—beliefs about what others approve of—we wanted to examine whether they are more closely related to second-order, other-moral beliefs—beliefs about what others think is moral—than first-order, personal-moral beliefs—individuals' own moral beliefs. Despite being strongly correlated with one another ($r = 0.90$), our results revealed a key dissociation between personal-moral beliefs and other-moral beliefs. While we found a moderately common association between both of our injunctive norm measures and other-moral judgments (injunctive-approve: $r = 0.67$; injunctive-should: $r = 0.66$) and personal-moral beliefs (injunctive-approve: $r = 0.63$; injunctive-should: $r = 0.61$), our findings indicate that injunctive beliefs were more closely related to other-moral judgments than personal-moral judgments ($Z = 13.6, p < .001$; Hittner et al., 2003). Furthermore, that participants were more likely to update their injunctive and other-moral beliefs across studies, and did so to a larger extent than their personal-moral beliefs—which were somewhat insensitive to descriptive norm information—suggests that participants viewed these measures differently. Together, this suggests that there is

a close, but imperfect relationship between injunctive beliefs and morality. Namely, while injunctive beliefs and moral judgments are interconnected, they are perhaps still partially independent constructs.

Why were participants' beliefs about others' moral judgments more sensitive to descriptive norms than participants' own moral judgments? There is a rich literature in psychology documenting a self-other asymmetry in which people's judgments and beliefs about themselves diverge from how they perceive and judge others (Grossmann & Kross, 2014; Molouki & Pronin, 2015; Pronin & Hazel, 2023; Pronin et al., 2002). For example, research on the bias blind spot finds that, while people are able to see bias in others, they are much less likely to see bias in their own behavior and that this asymmetry is driven by different cognitive mechanisms for self- and social perception (Pronin & Hazel, 2023). That we rely on different cognitive systems for self- and social perception may explain why participants in our study thought others' moral judgments would be more influenced by descriptive norms while their own moral judgments were largely insensitive to descriptive norms. This finding extends previous work on self-other asymmetries into the moral domain, suggesting that even our moral judgments are susceptible to this bias. This has potentially important implications for studying social norms—in models reported in our Supplemental Material (see *Self- and Other-Morality Models*), we find that participants' personal-moral judgments negatively predicted injunctive belief updating (i.e., the higher their personal-moral judgments, the less they updated) while other-moral judgments positively predicted updating (i.e., the higher their other-morality judgments, the more they updated).

More generally, this finding highlights an interesting paradox—people think that other people's moral judgments are easily swayed by descriptive norm information when in reality, most people's personal-moral judgments were largely insensitive to descriptive norms, only updating to a small extent (if at all). In other words, this finding suggests that people think that other people's beliefs are more malleable than they really are. Future work should further explore the relationship between injunctive beliefs and morality in more depth, and consider measuring both personal- and other-moral beliefs when studying moral judgment and social norms.

Do People Update to Different Extents Depending on the Type of Behavior?

While past work has examined the relationship between descriptive and injunctive norms, no work, to our knowledge, has investigated whether the relationship between normative information varies depending on the type of normative behavior in question. Examining the difference in updating allowed us to shed light on two important theoretical questions in the social norm literature. First, it allowed us to test Bicchieri's socially conditional account of norms, which holds that social norms are categorically distinct from moral norms and personal preferences in that social norms are influenced by others' expectations while moral norms and preferences are not (2006, 2016). If moral norms and personal preferences are socially independent behaviors that are not conditioned on others' expectations, then injunctive and moral beliefs about them should be less sensitive to descriptive norm information than for fairness and conventional norms. Second, it allowed us to examine whether people are sensitive to social expectations for fairness behaviors. Fairness behaviors are an interesting test of the social conditionality account because while they are often considered moral norms, recent work suggests that they are

perceived differently from harm norms (Yucel et al., 2022) and thus may fall somewhere between conventional and moral norms. Thus, by studying whether people are more or less sensitive to descriptive norms for fairness behaviors as for conventional and harm norms, we can better understand the role of social conditionality on normative cognition, as well as how people perceive fairness norms, which are an important part of human cooperation.

We found that, on average, participants updated their injunctive and moral norm beliefs to a greater extent for conventional and fairness behaviors compared to harm-related behaviors and, in some cases, preferences. Participants also consistently updated their behavioral intentions and beliefs about others' moral judgments more for fairness and conventional behaviors, while there was mixed evidence for differences in the effect of descriptive norms on personal-moral judgments, which were much less sensitive to descriptive norms. Importantly, participants tended to both positively update their beliefs (i.e., finding the behavior more injunctive) after receiving a common descriptive norm and negatively update their beliefs (i.e., finding the behavior less injunctive) after receiving an uncommon descriptive norm for fairness and conventional behaviors and, to a lesser degree, preferences. In contrast, for harm behaviors, participants were more likely to positively update their beliefs in response to common descriptive norms than they were to negatively update in response to uncommon descriptive norms.

Together, our findings suggest that people are moderately more sensitive to descriptive norms for fairness and conventional behaviors—which they tended to perceive quite similarly—than for harm behaviors and, to a lesser extent, personal preferences. This finding provides some support for Bicchieri's theory of social norms which holds that norms are socially conditional, hinging on our beliefs about how others expect us to behave (Bicchieri, 2006). Specifically, this theory suggests that unlike social norms, harm norms are internalized beliefs that are insensitive to social expectations. That participants in our study updated their beliefs less for harm behaviors than fairness and conventional behaviors suggests that people are indeed less influenced by others' expectations for moral norms (and harm norms specifically). Furthermore, because personal preferences are also independent beliefs that are not socially conditional, they should be less sensitive to descriptive norms and thus others' expectations. We find some evidence for this idea here—participants updated their beliefs less for preferences than conventional and fairness behaviors, although it varied by dependent measure (i.e., there was no difference in updating for the injunctive norm measure), and, in general, these effects were smaller than those observed for harm behaviors. However, it is important to note that participants still, albeit inconsistently, updated their beliefs and behavioral intentions for harm behaviors and preferences after receiving information that the behavior was common (or uncommon), just to a lesser extent than the fairness and conventional behaviors. That participants updated their beliefs and behavioral intentions even for harm behaviors such as stealing or mocking someone highlights how influential descriptive norms are in shaping our beliefs and behavior and suggest that, in contrast to Bicchieri's socially conditional account of norms, even moral norms and personal preferences may be somewhat sensitive to information about others' normative expectations.

Furthermore, our finding that participants' beliefs were more sensitive to descriptive norms for fairness behaviors than harm

behaviors indicates that people were more sensitive to others' expectations—and thus their beliefs were more socially conditional—for fairness than harm behaviors. In general, participants were nearly as sensitive to descriptive norms for fairness behaviors as conventional behaviors, suggesting that people may perceive fairness norms more similarly to conventional norms than harm norms. While belief updating (or the lack thereof) is likely influenced by a number of different factors, our findings suggest that harm may be unique in its insensitivity to normative information, and thus perceived as conceptually distinct from other kinds of normative behaviors. In contrast, our results suggest that people may perceive fairness behaviors more like conventional norms in that, like conventional norms, they are heavily dependent on what others are doing and expect of us, although to a somewhat lesser extent. These findings support recent work which finds that there are important differences in moral norms of fairness and harm—namely, that children do not view distributional fairness as a moral norm, an important finding given that fairness is often considered a moral norm (Yucel et al., 2022). That participants in our study were more sensitive to others' behavior and expectations for fairness behaviors than harm behaviors suggests that there are important differences between these types of behaviors and, more generally, calls into question the validity of grouping fairness and harm norms as equally moral norms, a common practice in the psychological literature on norms. Moving forward, researchers should avoid simply lumping fairness and harm norms in the category of moral norms and should further explore the boundaries of social conditionality by investigating when behaviors will be more or less sensitive to others' expectations.

Overall, our findings suggest that the effectiveness of descriptive norm information depends in part on the type of behavior in question. Why are we more sensitive to descriptive norms for some behaviors than others? Our results are consistent with the possibility that descriptive norms are more influential in contexts where there is more uncertainty about what is considered an appropriate behavior (Gelfand & Harrington, 2015). Namely, because beliefs about harm likely represent internalized norms—from an early age we are taught that hurting others is bad and come to believe it is wrong to do so even absent social pressure or sanctions—it makes sense that descriptive norm information would be less effective at influencing those beliefs as compared to other norms that are not as internalized and are more variable across contexts and groups (such as conventional norms). That participants' beliefs about fairness behaviors were influenced by descriptive norms nearly as much as their beliefs about conventional behaviors suggests that fairness norms are also highly variable and are perhaps more flexible than harm norms. This is supported by work finding that beliefs about harming others are common across cultures (besides in normatively appropriate circumstances like warfare), while fairness norms vary dramatically across cultures (Blake et al., 2015; Henrich et al., 2001; O'Neill & Machery, 2018; Sachdeva et al., 2011). Furthermore, uncertainty may explain why we found that participants were sensitive to descriptive norms for preferences, behaviors that, under the social conditionality account of norms, should be socially independent choices (Bicchieri, 2006, 2016). Namely, while our preference behaviors were designed to reflect personal proclivities, uncertainty about what is acceptable may have still influenced sensitivity to descriptive information (e.g., even if I personally like to wear sandals with socks, I may be less inclined to do so if very few others are and I'm uncertain about what is appropriate in that situation).

Another important dimension between the behaviors we studied was the emotional valence of the different behaviors. We initially categorized behaviors by emotional valence ratings during the behavior norming because behaviors that differed in their valence also differed substantially on the other norming dimensions of interest (e.g., injunctive normativity, morality, etc.). Consequently, we divided the behaviors by valence for norming and ran our models including all behaviors as well as subsetting by negative and positive valence. We note here that while the effects were smaller, we did find evidence of updating when examining solely the positively valenced behaviors. Namely, we found that the direction of updating depended on the specific positively valenced behavior—for most behaviors, participants positively updated their beliefs after receiving a common descriptive norm and negatively updated after receiving an uncommon descriptive, albeit to a lesser extent than for negatively valenced vignettes. However, for other behaviors, and for fairness behaviors in particular, participants negatively updated their beliefs after receiving either descriptive norm information, suggesting that the direction and extent of updating was influenced by both behavior type and valence. We also examined how valence predicts belief updating in Studies 1 and 2—the only studies that included both positive and negative behaviors. We found that participants updated to a larger extent for negative than positive behaviors, perhaps in part because of ceiling effects (see *Valence models* in the Supplemental Material for analyses). Because of these ceiling effects, we chose to exclusively include the negatively valenced behaviors from Study 3 onward and in our combined study analyses. While we did not include valence as a proxy for whether a behavior is prescriptively normative (e.g., you should donate money to charity) or proscriptively normative (e.g., you should not cheat on a test), we acknowledge that there are similarities between valence and whether a norm is prescriptive or proscriptive. That is, in general, behaviors that are rated as positively emotionally valenced are more likely to be prescriptively normative, while negatively emotionally valenced behaviors are more likely to be viewed as proscriptively normative. However, since we did not include valence as a measure of prescriptive–proscriptive normativity, our studies cannot directly speak to updating differences between prescriptive and proscriptive norms. However, given their similarities, we might expect that people would update their beliefs to a greater extent for proscriptive behaviors, in part because, like our negatively valenced behaviors, they have more room for variability since our priors are low to begin with (as compared to prescriptive behaviors). Furthermore, it could be the case that the presented descriptive norm information was discordant with the behavior depending on its valence—in other words, providing participants with a common descriptive norm about a negative behavior that they initially viewed as uncommon might be viewed as counter-normative. Thus, we report a series of models in the Supplemental Material (see *Counternormativity Models*) examining how the degree of counternormativity influences updating and, because it can be difficult to conceptualize the effect of updating for negative behaviors in relation to the positive behaviors, we include figures in the Supplemental Material that reverse the proportions of the descriptive norm condition for the negatively valenced behaviors so that they are on the same scale (e.g., positive behavior; see Supplemental Figures S5–S8) as well as figures showing the effect of updating by each individual vignette (see Supplemental Figures S13–S16).

Do Prior Descriptive Beliefs Moderate the Effect of Descriptive Norm Information on Belief Updating?

Because we studied real, preexisting behaviors, it is possible, if not likely, that participants' prior beliefs influenced their updating in our task. In nearly any situation, our prior beliefs about the commonality of a behavior will influence how we interpret and incorporate novel descriptive norm information. Indeed, one benefit of our pre–post descriptive norm study design is that it allowed us to examine how peoples' individual, preexisting descriptive norm beliefs influenced the efficacy of introducing novel normative information. To investigate this possibility, we examined whether participants' prior descriptive norm beliefs moderated the effect of the descriptive norm manipulation on updating. If prior descriptive beliefs do not moderate the effect of our manipulation, that could be interpreted as evidence for a number matching effect in which participants were simply matching their responses for the dependent measures to the size of the descriptive norm information (e.g., rating a behavior as “80” on an injunctive scale after receiving the common descriptive norm). However, speaking against that possibility, we found that individuals' prior descriptive beliefs moderated the effect of our descriptive norm manipulation on injunctive, behavioral intentions, other- and personal-morality updating. This suggests that the extent to which participants updated their beliefs in response to descriptive norm information was influenced by individual differences in participants' own beliefs about how descriptively normative the behaviors were.

Specifically, when unpacking the moderation analyses, we found that participants were more likely to update their injunctive norm beliefs and moral judgments when the descriptive norm information conflicted with their prior beliefs—the larger the difference between their prior descriptive expectations and the descriptive manipulation, the greater the degree of updating. In other words, if people already believe that most other people would cheat on a test, presenting a descriptive norm that cheating is a common practice will do little to change their injunctive and moral beliefs, whereas presenting a contrasting descriptive norm that cheating is not common will have a much larger influence on their normative beliefs and expectations. This suggests that people are more sensitive to descriptive norm information when it conflicts with their prior beliefs and expectations. Moreover, this finding has potentially important implications for descriptive norm messaging, suggesting that the influence of a normative intervention or nudge on behavior change hinges on individuals' prior descriptive beliefs about the commonality of the behavior. Consequently, it will be important for practitioners to assess individual beliefs about a target behavior before initiating a norm-based intervention.

Limitations

The present work was not without its limitations (see Table 5 for table of limitations). Namely, participants' prior beliefs likely played a role in the effect of descriptive norm information on beliefs; however, as explained above, this finding cannot explain between-condition differences in updating and, if anything, constitutes a strength of our pre–post study design. Additionally, while we normed all behaviors (see *Vignette Norming* in Supplemental Material), it is possible that the specific vignettes used influenced our results and that we might find different patterns of updating with different behaviors. Lastly, because behavioral intentions were self-reported, participants' actual behavior

Table 5*Table of Limitations of the Present Research*

Limitation	Description
Preexisting prior beliefs	Because we studied existing behaviors rather than entirely novel ones, we could not prevent participants from bringing in their own priors about the morality and normativity of the behaviors we studied. In other words, people might have had existing beliefs about how common or approved of a given behavior is based on their personal experiences which could have influenced their decisions in our task. In order to avoid people's prior experiences influencing their beliefs and behavior, future work should explore injunctive belief updating in the context of totally novel behaviors that people do not have priors for. If people update their injunctive beliefs and behavior after receiving descriptive norm information for norms they have no prior experience with or beliefs about, that would provide stronger evidence that descriptive norms influence injunctive norm beliefs and moral judgments.
Generalizability of using vignette and online convenience sample	Another limitation of the present work is that all of our data were collected via online convenience samples. While past work finds little differences in studies conducted in-lab and online (Amir et al., 2012; Horton et al., 2011), it is possible—while unlikely—that our results might not generalize to a broader population. Additionally, despite stringent norming across a number of dimensions, it is possible that our results hinge on the specific vignettes used and that our results might not generalize to different behaviors. However, given that we found effects of updating across all behaviors (fairness, conventional, harm, preference), we think this possibility unlikely; that said, future work should study belief updating using a range of different behaviors than those used here.
Injunctive to descriptive updating	Because we solely focused on the effect of descriptive norms on injunctive norm beliefs, we cannot necessarily make claims about effects in the opposite direction from injunctive to descriptive (i.e., that injunctive norm information influences people's descriptive norm beliefs). However, given the strong relationship between descriptive and injunctive norms, and previous work showing people make bidirectional inferences between them (Eriksson et al., 2015), we expect that people would be as likely to update their descriptive beliefs from injunctive norm information. Future work should explore whether and to what extent people update their descriptive norm beliefs, moral judgments, and behavior after receiving injunctive norm information.

might deviate from what they reported they would do. With that said, a meta-analysis by Webb and Sheeran examining the relationship between intentions and behavior found a significant relationship between them such that a medium-to-large change in intention led to a small-to-medium change in behavior (2006). This suggests that, while people's actual behavior is likely influenced by descriptive norms, it may be influenced to a lesser extent than their behavioral intentions.

Conclusion

Across a series of experiments, we find evidence in support of our three main research questions.

First, we provide strong evidence for belief updating between descriptive and injunctive norms. Namely, we find that, after receiving descriptive norm information that a behavior was either common or uncommon, participants updated their injunctive beliefs, behavioral intentions, and both personal and other-moral beliefs. Second, we find differences between first- and second-order moral judgments, such that participants' beliefs about others' moral judgments were more sensitive to descriptive norms than their own moral judgments. Third, we find that the effect of descriptive norms on updating varies depending on the type of behavior, such that, in general, participants were more sensitive to descriptive norm information for conventional and fairness behaviors than harms and

preferences. Our findings contribute to the literature on social norms and moral psychology, demonstrating the important role descriptive norms play in shaping injunctive norms and moral judgments and, more generally, how our beliefs about what others believe and do play a critical role in social norm cognition. Furthermore, our study advances past work by showing that people update a suite of interrelated normative beliefs in response to descriptive norm information and, in doing so, contributes a useful method for assessing the effect of normative information on beliefs. Specifically, by assessing beliefs before and after providing normative information, we were able to assess the extent to which participants updated their beliefs in response to said information, as well as how individual differences in preexisting normative beliefs moderated the influence of this novel descriptive information. Lastly, our findings inform current theories of social norm cognition, revealing that, as predicted by Bicchieri (2006, 2016), social norms are more strongly influenced by others' normative expectations than moral norms (like harm) and personal preferences. However, in contrast to that theory, we found that participants' beliefs about harm behaviors and preferences were still moderately influenced by descriptive norms, suggesting that even those types of behaviors are likely sensitive to empirical and normative expectations. All told, our findings highlight how our beliefs are flexible and fundamentally influenced by what we see around us in our social world, providing important insight into the social cognition underlying norm cognition and where our sense of what is "appropriate" or "right" comes from.

References

- Amir, O., Rand, D. G., & Gal, Y. K. (2012). Economic games on the internet: The effect of \$1 stakes. *PLOS ONE*, 7(2), Article e31461. <https://doi.org/10.1371/journal.pone.0031461>
- Appelbaum, M., Cooper, H., Kline, R. B., Mayo-Wilson, E., Nezu, A. M., & Rao, S. M. (2018). Journal article reporting standards for quantitative research in psychology: The APA Publications and Communications Board task force report. *American Psychologist*, 73(1), 3–25. <https://doi.org/10.1037/amp0000191>
- Asch, S. E. (1956). Studies of independence and conformity: I. A minority of one against a unanimous majority. *Psychological Monographs*, 70(9), 1–70. <https://doi.org/10.1037/h0093718>
- Axelrod, R. (1986). An evolutionary approach to norms. *The American Political Science Review*, 80(4), 1095–1111. <https://doi.org/10.2307/1960858>
- Balduzzi, S., Rücker, G., & Schwarzer, G. (2019). How to perform a meta-analysis with R: A practical tutorial. *Evidence-Based Mental Health*, 22(4), 153–160. <https://doi.org/10.1136/ebmental-2019-300117>
- Bates, D., Kliegl, R., Vasishth, S., & Baayen, H. (2015). *Parsimonious mixed models*. arXiv.
- Baumard, N., André, J. B., & Sperber, D. (2013). A mutualistic approach to morality: The evolution of fairness by partner choice. *Behavioral and Brain Sciences*, 36(1), 59–78. <https://doi.org/10.1017/S0140525X11002202>
- Bear, A., & Knobe, J. (2017). Normality: Part descriptive, part prescriptive. *Cognition*, 167, 25–37. <https://doi.org/10.1016/j.cognition.2016.10.024>
- Bicchieri, C. (2006). *The grammar of society: The nature and dynamics of social norms*. Cambridge University Press.
- Bicchieri, C. (2016). *Norms in the wild: How to diagnose, measure, and change social norms*. Oxford University Press.
- Blake, P. R., McAuliffe, K., Corbit, J., Callaghan, T. C., Barry, O., Bowie, A., Kleutsch, L., Kramer, K. L., Ross, E., Vongsachang, H., Wrangham, R., & Warneken, F. (2015). The ontogeny of fairness in seven societies. *Nature*, 528(7581), 258–261. <https://doi.org/10.1038/nature15703>
- Blake, P. R., McAuliffe, K., & Warneken, F. (2014). The developmental origins of fairness: The knowledge-behavior gap. *Trends in Cognitive Sciences*, 18(11), 559–561. <https://doi.org/10.1016/j.tics.2014.08.003>
- Blanton, H., Köblitz, A., & McCaul, K. D. (2008). Misperceptions about norm misperceptions: Descriptive, injunctive, and affective “social norming” efforts to change health behaviors. *Social and Personality Psychology Compass*, 2(3), 1379–1399. <https://doi.org/10.1111/j.1751-9004.2008.00107.x>
- Buffalo, M., & Rodgers, J. W. (1971). Behavioral norms, moral norms, and attachment: Problems of deviance and conformity. *Social Problems*, 19(1), 101–113. <https://doi.org/10.2307/799943>
- Cestac, J., Paran, F., & Delhomme, P. (2014). Drive as I say, not as I drive: Influence of injunctive and descriptive norms on speeding intentions among young drivers. *Transportation Research Part F: Traffic Psychology and Behaviour*, 23, 44–56. <https://doi.org/10.1016/j.trf.2013.12.006>
- Chakroff, A., Dungan, J., Koster-Hale, J., Brown, A., Saxe, R., & Young, L. (2016). When minds matter for moral judgment: Intent information is neurally encoded for harmful but not impure acts. *Social Cognitive and Affective Neuroscience*, 11(3), 476–484. <https://doi.org/10.1093/scan/nsv131>
- Chakroff, A., & Young, L. (2015). Harmful situations, impure people: An attribution asymmetry across moral domains. *Cognition*, 136, 30–37. <https://doi.org/10.1016/j.cognition.2014.11.034>
- Christensen, P. N., Rothgerber, H., Wood, W., & Matz, D. C. (2004). Social norms and identity relevance: A motivational approach to normative behavior. *Personality and Social Psychology Bulletin*, 30(10), 1295–1309. <https://doi.org/10.1177/0146167204264480>
- Cialdini, R. B., Demaine, L. J., Sagarin, B. J., Barrett, D. W., Rhoads, K., & Winter, P. L. (2006). Managing social norms for persuasive impact. *Social Influence*, 1(1), 3–15. <https://doi.org/10.1080/15534510500181459>
- 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>
- Curry, O. S., Mullins, D. A., & Whitehouse, H. (2019). Is it good to cooperate? Testing the theory of morality-as-cooperation in 60 societies. *Current Anthropology*, 60(1), 47–69. <https://doi.org/10.1086/701478>
- Deutchman, P., Kraft-Todd, G., Young, L., & McAuliffe, K. (2024, June 18). *Descriptive and Injunctive Norm Updating*. <https://osf.io/sc842/>
- Deutchman, P., Marshall, J., Lee, Y., Sansom, E., Warneken, F., & McAuliffe, K. (2023). *Children selectively update their injunctive norm beliefs and moral evaluations after receiving descriptive norm information*. PsyArXiv. <https://doi.org/10.31234/osf.io/2vd5t>
- Elster, J. (1989). Social norms and economic theory. *The Journal of Economic Perspectives*, 3(4), 99–117. <https://doi.org/10.1257/jep.3.4.99>
- Eriksson, K., Enquist, M., & Ghirlanda, S. (2007). Critical points in current theory of conformist social learning. *Journal of Evolutionary Psychology*, 5(1), 67–87. <https://doi.org/10.1556/JEP.2007.1009>
- Eriksson, K., & Strimling, P. (2015). Injunctive versus functional inferences from descriptive norms: Comment on Gelfand and Harrington. *Journal of Cross-Cultural Psychology*, 46(10), 1330–1332. <https://doi.org/10.1177/0022022115605387>
- Eriksson, K., Strimling, P., & Coultas, J. C. (2015). Bidirectional associations between descriptive and injunctive norms. *Organizational Behavior and Human Decision Processes*, 129, 59–69. <https://doi.org/10.1016/j.obhdp.2014.09.011>
- Folger, R. (1998). Fairness as a moral virtue. In M. Schminke (Ed.), *Managerial ethics: Moral management of people and processes* (pp. 13–34). Lawrence Erlbaum.
- Ge, W., Sheng, G., & Zhang, H. (2020). How to solve the social norm conflict dilemma of green consumption: The moderating effect of self-affirmation. *Frontiers in Psychology*, 11, Article 566571. <https://doi.org/10.3389/fpsyg.2020.566571>
- Gelfand, M. J., & Harrington, J. R. (2015). The motivational force of descriptive norms: For whom and when are descriptive norms most predictive of behavior? *Journal of Cross-Cultural Psychology*, 46(10), 1273–1278. <https://doi.org/10.1177/0022022115600796>
- Goldstein, N. J., Cialdini, R. B., & Griskevicius, V. (2008). A room with a viewpoint: Using social norms to motivate environmental conservation in hotels. *The Journal of Consumer Research*, 35(3), 472–482. <https://doi.org/10.1086/586910>
- Grossmann, I., & Kross, E. (2014). Exploring Solomon’s paradox: Self-distancing eliminates the self-other asymmetry in wise reasoning about close relationships in younger and older adults. *Psychological Science*, 25(8), 1571–1580. <https://doi.org/10.1177/0956797614535400>
- Harrer, M., Cuijpers, P., Furukawa, T. A., & Ebert, D. D. (2021). *Doing meta-analysis with R: A hands-on guide*. Chapman and Hall/CRC. <https://doi.org/10.1201/9781003107347>
- Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., Gintis, H., & McElreath, R. (2001). In search of homo economicus: Behavioral experiments in 15 small-scale societies. *The American Economic Review*, 91(2), 73–78. <https://doi.org/10.1257/aer.91.2.73>
- Hittner, J. B., May, K., & Silver, N. C. (2003). A Monte Carlo evaluation of tests for comparing dependent correlations. *The Journal of General Psychology*, 130(2), 149–168. <https://doi.org/10.1080/00221300309601282>
- Horton, J. J., Rand, D. G., & Zeckhauser, R. J. (2011). The online laboratory: Conducting experiments in a real labor market. *Experimental Economics*, 14(3), 399–425. <https://doi.org/10.1007/s10683-011-9273-9>
- House, B. R., Kanngiesser, P., Barrett, H. C., Broesch, T. C., Cebiglu, S., Crittenden, A. N., Erut, A., Lew-Levy, S., Sebastian-Enesco, C., Smith, A. M., Yilmaz, S., & Silk, J. B. (2020). Universal norm psychology leads to societal diversity in prosocial behaviour and development. *Nature Human Behaviour*, 4(1), 36–44. <https://doi.org/10.1038/s41562-019-0734-z>

- Jacobson, R. P., Mortensen, C. R., & Cialdini, R. B. (2011). Bodies obliged and unbound: Differentiated response tendencies for injunctive and descriptive social norms. *Journal of Personality and Social Psychology, 100*(3), 433–448. <https://doi.org/10.1037/a0021470>
- Keizer, K., Lindenberg, S., & Steg, L. (2011). The reversal effect of prohibition signs. *Group Processes & Intergroup Relations, 14*(5), 681–688. <https://doi.org/10.1177/1368430211398505>
- Lindström, B., Jangard, S., Selbing, I., & Olsson, A. (2018). The role of a “common is moral” heuristic in the stability and change of moral norms. *Journal of Experimental Psychology: General, 147*(2), 228–242. <https://doi.org/10.1037/xge0000365>
- Lu, H., Zou, J., Chen, H., & Long, R. (2020). Promotion or inhibition? Moral norms, anticipated emotion and employee’s pro-environmental behavior. *Journal of Cleaner Production, 258*, Article 120858. <https://doi.org/10.1016/j.jclepro.2020.120858>
- Meisel, M. K., & Goodie, A. S. (2014). Descriptive and injunctive social norms’ interactive role in gambling behavior. *Psychology of Addictive Behaviors, 28*(2), 592–598. <https://doi.org/10.1037/a0036444>
- Mesoudi, A. (2009). How cultural evolutionary theory can inform social psychology and vice versa. *Psychological Review, 116*(4), 929–952. <https://doi.org/10.1037/a0017062>
- Molouki, S., & Pronin, E. (2015). Self and other. In M. Mikulincer, P. R. Shaver, E. Borgida, & J. A. Bargh (Eds.), *APA handbook of personality and social psychology, Vol. 1. Attitudes and social cognition* (pp. 387–414). American Psychological Association. <https://doi.org/10.1037/14341-013>
- Neighbors, C., O’Connor, R. M., Lewis, M. A., Chawla, N., Lee, C. M., & Fossos, N. (2008). The relative impact of injunctive norms on college student drinking: The role of reference group. *Psychology of Addictive Behaviors, 22*(4), 576–581. <https://doi.org/10.1037/a0013043>
- Nielsen, L., & McGregor, S. L. (2013). Consumer morality and moral norms. *International Journal of Consumer Studies, 37*(5), 473–480. <https://doi.org/10.1111/ijcs.12032>
- O’Neill, E., & Machery, E. (2018). The normative sense: What is universal? What varies? In A. Zimmerman, K. Jones, & M. Timmons (Eds.), *The Routledge handbook of moral epistemology* (pp. 38–56). Routledge. <https://doi.org/10.4324/9781315719696-3>
- Ostermann, C. M., & de Barcellos, M. D. (2021, September 8–10). *156 Conflicts in meat consumption: Exploring their effects on social norms and consumer behavior* [Conference session]. 20th European Round Table on Sustainable Consumption and Production 2021, Graz, Austria.
- Pronin, E., & Hazel, L. (2023). Humans’ Bias blind spot and its societal significance. *Current Directions in Psychological Science, 32*(5), 402–409. <https://doi.org/10.1177/09637214231178745>
- Pronin, E., Lin, D. Y., & Ross, L. (2002). The bias blind spot: Perceptions of bias in self versus others. *Personality and Social Psychology Bulletin, 28*(3), 369–381. <https://doi.org/10.1177/0146167202286008>
- R Core Team. (2023). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing.
- Raihani, N. J., & McAuliffe, K. (2014). Dictator game giving: The importance of descriptive versus injunctive norms. *PLOS ONE, 9*(12), Article e113826. <https://doi.org/10.1371/journal.pone.0113826>
- Reno, R. R., Cialdini, R. B., & Kallgren, C. A. (1993). The transsituational influence of social norms. *Journal of Personality and Social Psychology, 64*(1), 104–112. <https://doi.org/10.1037/0022-3514.64.1.104>
- Roberts, S. O., Gelman, S. A., & Ho, A. K. (2017). So it is, so it shall be: Group regularities license children’s prescriptive judgments. *Cognitive Science, 41*(Suppl. 3), 576–600. <https://doi.org/10.1111/cogs.12443>
- Roberts, S. O., Guo, C., Ho, A. K., & Gelman, S. A. (2018). Children’s descriptive-to-prescriptive tendency replicates (and varies) cross-culturally: Evidence from China. *Journal of Experimental Child Psychology, 165*, 148–160. <https://doi.org/10.1016/j.jecp.2017.03.018>
- Russell, P. S., Smith, D. M., Birtel, M. D., Hart, K. H., & Golding, S. E. (2022). The role of emotions and injunctive norms in breastfeeding: A systematic review and meta-analysis. *Health Psychology Review, 16*(2), 257–279. <https://doi.org/10.1080/17437199.2021.1893783>
- Sachdeva, S., Singh, P., & Medin, D. (2011). Culture and the quest for universal principles in moral reasoning. *International Journal of Psychology, 46*(3), 161–176. <https://doi.org/10.1080/00207594.2011.568486>
- Smetana, J. G. (2006). Social-cognitive domain theory: Consistencies and variations in children’s moral and social judgments. In M. Killen & J. G. Smetana (Eds.), *Handbook of moral development* (pp. 119–153). Lawrence Erlbaum.
- Smetana, J. G. (2013). Moral development: The social domain theory view. In P. D. Zelazo (Ed.), *The Oxford Handbook of Developmental Psychology (Vol. 1): Body and mind* (pp. 832–863). Oxford University Press.
- Smetana, J. G., & Ball, C. L. (2018). Young children’s moral judgments, justifications, and emotion attributions in peer relationship contexts. *Child Development, 89*(6), 2245–2263. <https://doi.org/10.1111/cdev.12846>
- Smetana, J. G., Wong, M., Ball, C., & Yau, J. (2014). American and Chinese children’s evaluations of personal domain events and resistance to parental authority. *Child Development, 85*(2), 626–642. <https://doi.org/10.1111/cdev.12140>
- Smith, J. R., Louis, W. R., Terry, D. J., Greenaway, K. H., Clarke, M. R., & Cheng, X. (2012). Congruent or conflicted? The impact of injunctive and descriptive norms on environmental intentions. *Journal of Environmental Psychology, 32*(4), 353–361. <https://doi.org/10.1016/j.jenvp.2012.06.001>
- Staunton, M., Louis, W. R., Smith, J. R., Terry, D. J., & McDonald, R. I. (2014). How negative descriptive norms for healthy eating undermine the effects of positive injunctive norms. *Journal of Applied Social Psychology, 44*(4), 319–330. <https://doi.org/10.1111/jasp.12223>
- te Velde, V. L., & Louis, W. (2022). Conformity to descriptive norms. *Journal of Economic Behavior & Organization, 200*, 204–222. <https://doi.org/10.1016/j.jebo.2022.05.017>
- Turiel, E., Killen, M., & Helwig, C. C. (1987). Morality: Its structure, functions, and vagaries. In J. Kagan & S. Lamb (Eds.), *The emergence of morality in young children* (pp. 155–243). The University of Chicago Press.
- Tworek, C. M., & Cimpian, A. (2016). Why do people tend to infer “ought” from “is”? The role of biases in explanation. *Psychological Science, 27*(8), 1109–1122. <https://doi.org/10.1177/0956797616650875>
- Webb, T. L., & Sheeran, P. (2006). Does changing behavioral intentions engender behavior change? A meta-analysis of the experimental evidence. *Psychological Bulletin, 132*(2), 249–268. <https://doi.org/10.1037/0033-2909.132.2.249>
- White, K. M., Smith, J. R., Terry, D. J., Greenslade, J. H., & McKimmie, B. M. (2009). Social influence in the theory of planned behaviour: The role of descriptive, injunctive, and in-group norms. *British Journal of Social Psychology, 48*(1), 135–158. <https://doi.org/10.1348/014466608X295207>
- Whiten, A. (2019). Conformity and over-imitation: An integrative review of variant forms of hyper-reliance on social learning. *Advances in the Study of Behavior, 51*, 31–75. <https://doi.org/10.1016/bs.asb.2018.12.003>
- Wickham, H., François, R., Henry, L., Müller, K., & Vaughan, D. (2023). *dplyr: A grammar of data manipulation* (R package Version 1.1.4). <https://dplyr.tidyverse.org>
- Yucel, M., Drell, M. B., Jaswal, V. K., & Vaish, A. (2022). Young children do not perceive distributional fairness as a moral norm. *Developmental Psychology, 58*(6), 1103–1113. <https://doi.org/10.1037/dev0001349>

Received July 4, 2023

Revision received July 12, 2024

Accepted July 14, 2024 ■