

Bidirectional associations between descriptive and injunctive norms

Kimmo Eriksson, Pontus Strimling, Julie C Coultas

Publication date

01-07-2015

Licence

This work is made available under the **Copyright not evaluated** licence and should only be used in accordance with that licence. For more information on the specific terms, consult the repository record for this item.

Citation for this work (American Psychological Association 7th edition)

Eriksson, K., Strimling, P., & Coultas, J. C. (2015). *Bidirectional associations between descriptive and injunctive norms* (Version 1). University of Sussex. <https://hdl.handle.net/10779/uos.23417966.v1>

Published in

Organizational Behavior and Human Decision Processes

Link to external publisher version

<https://doi.org/10.1016/j.obhdp.2014.09.011>

Copyright and reuse:

This work was downloaded from Sussex Research Open (SRO). This document is made available in line with publisher policy and may differ from the published version. Please cite the published version where possible. Copyright and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners unless otherwise stated. For more information on this work, SRO or to report an issue, you can contact the repository administrators at sro@sussex.ac.uk. Discover more of the University's research at <https://sussex.figshare.com/>

Bidirectional associations between descriptive and injunctive norms

Article (Unspecified)

Citation:

Eriksson, Kimmo, Strimling, Pontus and Coultas, Julie C (2015) Bidirectional associations between descriptive and injunctive norms. *Organizational Behavior and Human Decision Processes*, 129. pp. 59-69. ISSN 0749-5978

This version is available from Sussex Research Online: <http://sro.sussex.ac.uk/55039/>


This document is made available in accordance with publisher policies and may differ from the published version or from the version of record. If you wish to cite this item you are advised to consult the publisher's version. Please see the URL above for details on accessing the published version.

Copyright and reuse:

Sussex Research Online is a digital repository of the research output of the University.


Copyright and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable, the material made available in SRO has been checked for eligibility before being made available.

Copies of full text items generally can be reproduced, displayed or performed and given to third parties in any format or medium for personal research or study, educational, or not-for-profit purposes without prior permission or charge, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.



Contents lists available at ScienceDirect

Organizational Behavior and Human Decision Processes

journal homepage: www.elsevier.com/locate/obhdp

Bidirectional associations between descriptive and injunctive norms

Kimmo Eriksson^{a,b,*}, Pontus Strimling^{a,c}, Julie C. Coultas^{a,d}^aStockholm University, Centre for the Study of Cultural Evolution, Stockholm, Sweden^bMälardalen University, Västerås, Sweden^cLinköping University, Institute for Analytical Sociology, Norrköping, Sweden^dUniversity of Sussex, Brighton, UK

ARTICLE INFO

Article history:

Received 10 December 2012

Accepted 30 September 2014

Available online xxx

Accepted by Michael Morris, Ying-yi Hong and Chi-yue Chiu

Keywords:

Descriptive norms

Injunctive norms

Automatic association

Priming

Implicit association test

Recall

Moral judgments

Frequency dependence

Conformity

Social influence

ABSTRACT

Modern research on social norms makes an important distinction between descriptive norms (how people commonly behave) and injunctive norms (what one is morally obligated to do). Here we propose that this distinction is far from clear in the cognition of social norms. In a first study, using the implicit association test, the concepts of “common” and “moral” were found to be strongly associated. Some implications of this automatic common–moral association were investigated in a subsequent series of experiments: Our participants tended to make explicit inferences from descriptive norms to injunctive norms and vice versa; they tended to mix up descriptive and injunctive concepts in recall tasks; and frequency information influenced participants’ own moral judgments. We conclude by discussing how the common–moral association could play a role in the dynamics of social norms.

© 2014 Elsevier Inc. All rights reserved.

Introduction

Social norms and their power to govern behavior have been studied extensively in the social sciences. There are several different theoretical approaches to social norms, such as the focus theory of Cialdini and colleagues (e.g., Cialdini, Kallgren, & Reno, 1991), the social norms approach of Berkowitz and Perkins (e.g., Berkowitz, 2004), and Bicchieri’s theory of dynamics of norms in social dilemmas (e.g., Bicchieri, 2006). These different approaches share a theoretical division of social norms into two distinct types: *injunctive norms*, referring to people’s beliefs about how one ought to behave, and *descriptive norms*, referring to people’s beliefs about what most people actually do. For a review of various lines of work based on these concepts, see Lapinski and Rimal (2005).

In this paper we are concerned with the relation between injunctive and descriptive norms. The two types of norms are often *congruent*, by which we shall mean that what is common to do is also what you ought to do. For instance, at a formal meeting the descriptive norm is that most individuals will be silent and

attentive. This norm is also injunctive, as noncompliance is likely to incur social sanctions (example from Lapinski & Rimal, 2005, p. 131). Several studies have found that injunctive and descriptive norms tend to be congruent (e.g., Brauer & Chaurand, 2010; Thøgersen, 2008). This statistical correlation is not a logical necessity, though. People are quite capable of endorsing the moral desirability of a certain behavior, yet not practice it (Cialdini, Reno, & Kallgren, 1990). Thus, the concepts of descriptive and injunctive norms are logically distinct, echoing David Hume’s famous law that one cannot derive what *ought* to be from what *is*.

The reason for Hume to state this law was that he saw other writers violating it, that is, they made claims about what ought to be on the basis of what is. An intriguing question is why this fallacy is so common. Can it be that the distinction between the descriptive and the injunctive is not very clear in people’s minds? Specifically, can it be that commonness and morality are automatically associated with each other? This fundamental question about cognition of social norms seems not to have been asked before. However, it is related to a (rather loose) proposal of Kelley (1971) according to which people depend on others’ behavior as the evidence of what is right and wrong and therefore “concepts of what ‘ought to be’ tend to drift toward conceptions of what ‘is’” (Kelley, 1971, p. 298).

* Corresponding author at: Stockholm University, Centre for the Study of Cultural Evolution, Wallenberglab., SE-10691 Stockholm, Sweden.

E-mail address: kimmo.eriksson@mdh.se (K. Eriksson).

Observations suggestive of a common–moral association

The possibility of a cognitive link between the descriptive and the injunctive is supported by a number of observations. First, as Hume noted, his contemporaries were in the habit of blurring the distinction. It is not difficult to find modern examples of such blurring too. For instance, parents raising children may interchangeably use expressions like the injunctive “you shouldn’t do that” and the descriptive “we don’t do that” (Boyer, 2012). Lapinski and Rimal (2005, p. 130) note that even researchers of norms sometimes conflate the distinction between injunctive and descriptive norms.

Moreover, a large body of research has established a general tendency of people to conform not only to injunctive norms but also to descriptive norms (e.g., Bicchieri & Xiao, 2009; Borsari & Carey, 2001; Borsari & Carey, 2003; Cialdini, 2003; Cialdini & Goldstein, 2004; Cialdini et al., 1990, 1991; Claidière & Whiten, 2012; Perkins & Berkowitz, 1986). This similarity in the effects of injunctive and descriptive norms is consistent with a mental association such that both kinds of norms activate the same behavioral schemas.

The strongest evidence suggestive of a common–moral association comes from studies finding that moral judgments of socially undesirable behavior tend to be less harsh when the behavior is perceived to be common (McGraw, 1985; Trafimow, Reeder, & Blising, 2001; Welch et al., 2005), and that people seem to find fault both with singular selfishness and singular generosity in others (Parks & Stone, 2010).

In sum, several empirical observations suggest that people’s ideas about the morality of a behavior are linked to their beliefs about how common it is. We shall now argue that it is plausible that this association is automatic.

The common–moral association hypothesis

There is a large literature on the formation of automatic associations. While much of this literature focuses on the formation of evaluative associations (i.e., attitudes), most theoretical assumptions are equally applicable to non-evaluative domains (Gawronski & Sritharan, 2010). It is well-established that automatic associations can come about through classical conditioning, that is, repeated co-occurrence of stimuli (Gawronski & Bodenhausen, 2011; Olson & Fazio, 2001). Thus, the typical co-occurrence between injunctive norms and descriptive norms – when it is salient what one ought to do, most people will also be seen doing it – should tend to make people form a corresponding mental association. Hence, both observations and theory support the hypothesis that commonness and morality tend to be automatically associated with each other in people’s minds. In Study 1 we test this hypothesis using the most well-established method for measuring automatic associations, the Implicit Association Test (IAT; Greenwald, McGhee, & Schwartz, 1998).

Note that the concept of morality is in itself multi-faceted. For instance, it encompasses both moral obligation and moral character (Kelley, 1971; McGraw, 1985). However, the conceptual difference between these aspects of morality seems very subtle compared to the conceptual difference between commonness and morality. Consequently, the hypothesized automatic association with commonness is expected to apply across the different facets of morality.

Below we describe several implications of the common–moral hypothesis, which we then test in a series of experimental studies. In order to cover the injunctive–moral spectrum from obligation to goodness, these studies used a variety of terms (obligated/up to you, OK/not OK, moral/immoral, etc.). Studies also varied the set of behavioral stimuli that these terms applied to, in order to cover both the *proscriptive* domain of socially desirable behaviors and the

proscriptive domain of socially undesirable behaviors (Janoff-Bulman, Sheikh, & Hepp, 2009). Stimuli also covered a range of base rates and drew upon a range of different moral foundations (Graham, Haidt, & Nosek, 2009).

Implication 1: explicit inferences between commonness and morality

Automatic associations have interesting downstream effects on cognition and behavior (see, for instance, the introductory paragraph of Gawronski & Sritharan, 2010). To begin with – unsurprisingly but importantly – automatic associations influence people’s explicit, verbally reported evaluations. The most comprehensive model of the relationship between explicit evaluations and implicit measures of associations is the APE model of Gawronski and Bodenhausen (2006, 2011). According to the APE model, agreement between explicit and implicit measures is expected as long as the association, when made explicit as a verbal proposition, is not inconsistent with other relevant beliefs (Gawronski & Bodenhausen, 2011, p. 63). For instance, consider a question like: “To do X is common in group G; what do you think is the moral status of doing X?” The automatic association of our hypothesis will suggest that X is a moral thing to do. According to the APE model, the cognitive process then proceeds with this suggestion being weighed against other beliefs the respondent may have about the appropriateness to verbalize that X is a moral thing to do. To the extent that no strong contrary beliefs exist, the respondent should then verbalize the moral status that is congruent with the frequency information.

In other words, the common–moral association hypothesis implies that *explicit inferences between commonness and morality should tend to be congruent*. We test this prediction in Study 2.

Implication 2: distortion of memory

Memory research has established that schema-based expectancies can cause memory distortion (Schacter, 1999, p. 194). For instance, Banaji and Bhaskar (2000) reported a study where participants read a list of African American and European American names and were asked to identify names of criminals that had appeared in the media. In fact, none of the names were names of criminals. African American names were incorrectly recalled as those of criminals almost twice as often as European American names.

Information about social norms tends to be particularly well remembered according to a study by O’Gorman, Wilson, and Miller (2008). However, their study did not distinguish between injunctive and descriptive norms. We propose that an automatic common–moral association may be a source of error in recall of social norms. Reminiscent of the abovementioned effect for memory of names of criminals, we expect information about one type of norm to interfere with the memory of previous information of the other type (e.g., information that behavior X is common might cause information that X is immoral to be incorrectly recalled as X being moral). We test this prediction in Studies 3 and 4.

We also propose an even stronger type of recall error: The cognitive conflation of descriptive and injunctive norms may be so potent that information about one type of norm is misremembered as the congruent norm of the other type (e.g., information that behavior X is moral might be recalled as X being common). As far we know, this kind of recall error across associated categories has not been studied before. We test this prediction in Studies 5 and 6.

Implication 3: influence of frequency information on own moral judgments

Moral judgments largely rely on automatic evaluations (Haidt, 2001). We should therefore expect the automatic common–moral

association to lead to descriptive information actually influencing people's own moral judgments. Indeed, we have already mentioned some studies that have found people tend to judge behaviors in the proscriptive domain as less immoral when common (McGraw, 1985; Trafimow et al., 2001; Welch et al., 2005). In the prescriptive domain, however, the prediction from the common–moral association competes with a previous prediction from trait attribution theory. Specifically, it has been suggested that uncommon behaviors should generally lead to more extreme attributions of traits (Reeder & Brewer, 1979). Applying this suggestion to moral traits, McGraw (1985) hypothesized that someone behaving in an immoral and uncommon way should be judged as particularly immoral, whereas someone behaving in a moral and uncommon way should be judged as particularly moral. Note that the common–moral association hypothesis makes the same prediction in the proscriptive domain but the opposite prediction in the prescriptive domain.

McGraw (1985) ran two experiments, each of which tested the prediction in both domains. In the proscriptive domain results were as predicted. In the prescriptive domain, the first experiment found a tendency for a positive relation between perceived commonness and morality – contrary to McGraw's hypothesis but in line with the common–moral association hypothesis. Whereas the first experiment studied a single behavior, McGraw's second experiment used several behaviors of different degrees of commonness. More common behaviors showed a tendency for higher moral attributions, again contrary to McGraw's hypothesis and in line with the common–moral association. However, the evidence was far from conclusive. For one thing, neither of these tendencies was statistically significant. Moreover, the methodology of correlating perceptions of commonness with moral judgments is subject to the third variable problem (i.e., a causal relation cannot be concluded).¹

In two experiments (Studies 7 and 8) we assessed the influence of information about descriptive norms on own moral judgments in the prescriptive domain. We manipulated the information between conditions to make it possible to draw conclusions about causality.

Sampled populations and modes of data collection

The sampled population and mode of data collection varied across our studies. Studies 1 and 3 were conducted in a laboratory with Swedish undergraduates. The other studies were conducted online with American participants recruited through the Amazon Mechanical Turk (mturk.com). This is a convenient and reliable data source for online behavioral studies (Buhrmester, Kwang, & Gosling, 2011; Paolacci, Chandler, & Ipeirotis, 2010). Participants who had already taken part in a previous study in this project were identified by their mturk user ID number and excluded. Thus, the data presented in this paper are based on unique sets of participants for each online study.

Overview of studies

In sum, we present a series of eight studies. The first study tests the core hypothesis, the presence of an automatic common–moral association. The remaining studies experimentally investigate implications for norm cognition in terms of explicit inferences of norms, false recall of norms, and influences on own moral judgments.

¹ In addition, McGraw also reported a negative correlation between individually perceived commonness and moral attributions in the second experiment. However, as every individual rated several different behaviors it is unclear how this correlation was calculated and if it has any meaningful interpretation.

Study 1

Our core hypothesis is that the common–moral association is automatic. To enable this conclusion an implicit measure must be used. The Implicit Association Test (Greenwald et al., 1998) is a widely used paradigm to measure the strength of the automatic association between different concepts. It rests on the assumption that it should be easier to make the same behavioral response to concepts that are more strongly associated. A few studies have used IAT to measure associations between morality and other concepts, such as the association between moral and self (Perugini & Leone, 2009). Here we are interested in the association between concepts related to the moral status of behaviors and concepts related to the commonness of behaviors.

In the IAT participants are asked to sort stimuli into categories. In our study these are Moral, Immoral, Common, and Uncommon. Sorting is done by pressing one of two keys. In initial practice blocks of trials only two categories are used (e.g., the left key for Moral and the right key for Immoral). During the main tasks, all four categories are shown simultaneously (e.g., Moral and Uncommon on the left key and Immoral and Uncommon on the right key). The pairing of categories switches between blocks. Differences in reaction time are used to estimate the relative strength of associations between categories (Nosek, Greenwald, & Banaji, 2005).

Our hypothesis is that moral status and commonness concepts are associated, such that participants' response latencies should be shorter when the key for Immoral was also used for Uncommon rather than for Common.

Method

Participants

Participants were recruited from a large pool of volunteers at Stockholm University, at a compensation of 150 Swedish kronor. 103 participants (31% female, 3% unknown) with age ranging from 19 to 47 years ($M = 27$ years, $SD = 7$ years) took part in this study.

Materials, design and procedure

The study was conducted in the GameLab at Stockholm University, a laboratory where participants in separate cubicles take part in computer based studies. The IAT was administered using the Inquisit 4 software. During the IAT, participants used the E- and I-keys to sort stimulus words either left or right. Stimuli were presented in the center of the computer screen and sorting categories were shown in the upper left and right of the screen. In translation from Swedish, the following categories and stimuli were used.² The target category was Immoral (represented by the stimuli "wrong", "not okay", "bad"), which was contrasted with Moral ("right", "okay", "good"). The paired categories were Common ("almost everyone", "majority", "typical") and Uncommon ("almost no-one", "minority", "untypical"). Previous meta-analyses have shown as few as two words to be sufficient to represent a category in an IAT (Nosek et al., 2005).

Presentation of stimuli followed a standard sequence of blocks described in Table 1. The order of the sequence was counterbalanced such that for half the participants the pairing Immoral–Uncommon preceded the pairing Immoral–Common.

Scoring

Scoring followed the D_2 algorithm (Greenwald, Nosek, & Banaji, 2003). This means that latencies below 400 ms and above 10,000 ms were removed. Errors were implicitly penalized in terms

² The Swedish terms used were: *Moralisk* (rätt, okej, god), *O-moralisk* (fel, inte okej, dålig), *Vanlig* (nästan alla, majoritet, typisk), *Ovanlig* (nästan ingen, minoritet, otypisk).

of larger latencies due to the software signaling the error and requiring a correct response to be given. For each participant, the difference between the mean latency of the two practice blocks was divided by the standard deviation of latencies in these blocks pooled together. The analogous ratio was also calculated for the test blocks. The *D* score of a participant was then computed as the average of these two ratios. This is the IAT effect size measure, with a theoretical range between -2 and 2 . Differences were computed such that a positive *D* score indicates a positive association between Immoral and Uncommon (i.e., faster responses when Immoral used the same key as Uncommon rather than Common).

Results

D scores were positive for 102 out of 103 participants, indicating a near universal association between the categories Immoral and Uncommon. The mean *D* score was 1.09 ($SD = 0.36$) with a 99.9% confidence interval between 0.97 and 1.21. This is a very large effect, at least as large as that found in studies of well-known associations between gender and science or implicit race attitude (Greenwald et al., 2003).

Discussion

The core assumption of the common–moral association hypothesis was supported: Using the implicit association test we found evidence for a strong automatic association between injunctive and descriptive concepts, exhibited by almost every single participant.

Study 2

The first study assessed implicit associations between descriptive and injunctive norms. Our next aim was to investigate whether the association is still present when people are explicitly asked to make inferences between descriptive and injunctive norms. We studied both the prescriptive and proscriptive domains.

Method

Participants

Two hundred participants (35% female) with age ranging from 18 to 74 years ($M = 30$ years, $SD = 11$ years) were recruited online among American users of Mturk at a compensation of half a US dollar.

Materials

Participants were asked to imagine that they traveled to a new place where customs may be different from what they are used to. The questionnaire introduced two behaviors called “phooshing” and “quining”. Phooshing was described as a generic socially undesirable behavior (“other people might appreciate if you don’t do it”),

whereas quining was described as a generic socially desirable behavior (“other people might appreciate if you did it”). Participants were first asked what would be most important for them to know about the new place: the actual frequencies of behaviors or the general opinions about what you ought or ought not to do. This question served the purpose of emphasizing the conceptual difference between descriptive and injunctive norms.

We then presented participants with information about the descriptive norms, i.e., the frequency of phooshing and quining in this strange place. Given this information, participants were asked for each behavior what they thought the injunctive norm would be. For example, the exact wording (in one condition) of the question about the undesirable behavior was: “Suppose you learn about some new place that it is not common to phoosh. What is your intuition on the general opinion about phooshing?” Participants were then given a binary choice between “one ought not to phoosh” and “it is up to you whether you want to phoosh”. The question about the desirable behavior was similar, with a binary choice between “one ought to quine” and “it is up to you whether you want to quine”.

Then followed similar questions where participants were instead presented with injunctive norms for each behavior and asked what they thought the descriptive norms would be. For example, the exact wording (in one condition) of the question about the undesirable behavior was: “Suppose you learn about some new place that the general opinion about phooshing is that you ought not to do it. What is your intuition about whether phooshing is common in this new place?” Participants were then given a binary choice between “it is common to phoosh” and “it is not common to phoosh”.

Across participants information was counterbalanced such that different combinations of pieces of information occurred equally often.³

Results

Results were overwhelmingly in favor of a common–moral association. The inference that you ought not to phoosh was much more frequent when phooshing was presented as not common (80%) rather than common (15%), $\chi^2(1, N = 199) = 80.82$, $p < .001$, odds ratio = 22.38. Similarly, the inference that you ought to quine was much more frequent when quining was presented as common (80%) rather than not common (33%), $\chi^2(1, N = 198) = 41.42$, $p < .001$, odds ratio = 7.90.

Similar results held for inferences from injunctive to descriptive norms. Thus, it was much more frequent to infer that it is uncommon to phoosh when phooshing was presented as something you ought not to do (88%) rather than up to you (45%), $\chi^2(1, N = 199) = 41.45$, $p < .001$, odds ratio = 8.96. Similarly, it was much more frequent to infer that it is common to quine when quining was presented as something you ought to do (93%) rather than up to you (50%), $\chi^2(1, N = 199) = 45.37$, $p < .001$, odds ratio = 13.29.

Discussion

This simple study asked people to make inferences from descriptive norms to injunctive norms and vice versa. A strong tendency for a common–moral association was in evidence in these explicit inferences. As expected, this tendency was found in both inferential directions and across the prescriptive and proscriptive domains.

³ For instance, both “phooshing common” and “phooshing not common” occurred equally often together with “quining common” and “quining not common”, and equally often together with “you ought not to phoosh” and “it is up to you to phoosh”.

Table 1
Sequence of trial blocks in IAT.

Block	Trials	Function	Left key response	Right key response
1	20	Practice	Immoral	Moral
2	20	Practice	Common	Uncommon
3	20	Practice	Immoral + Common	Moral + Uncommon
4	40	Test	Immoral + Common	Moral + Uncommon
5	20	Practice	Moral	Immoral
6	20	Practice	Moral + Common	Immoral + Uncommon
7	40	Test	Moral + Common	Immoral + Uncommon

Note: For half the subjects, the positions of Blocks 1, 3, and 4 were switched with those of Blocks 5, 6, and 7, respectively. The procedure in Blocks 3, 4, 6, and 7 was to alternate trials that presented a stimulus from the Moral/Immoral categories with trials that presented a stimulus from the Common/Uncommon categories.

Study 3

In the same lab sessions in which Study 1 was run we also tested the hypothesis of interference between descriptive and injunctive information in recall. We presented injunctive information for a set of behaviors, followed by descriptive information for the same set of behaviors. The information was manipulated such that injunctive and descriptive norms were randomized to be either congruent or incongruent. Participants were then asked what injunctive norm had been presented for each behavior. Our hypothesis is that participants automatically associate the descriptive information to injunctive norms and that this association can interfere with the memory of the earlier injunctive information. An incorrect response to an incongruent item (e.g., a behavior presented as “not OK” and “common” recalled as “OK”) would be a *congruent recall error*, because it is consistent with recall of the congruent associated norm instead of the presented norm. An incorrect response to a congruent item (e.g., a behavior presented as “OK” and “common” recalled as “not OK”) would be an *incongruent recall error*. Our hypothesis predicts congruent recall errors to dominate over incongruent recall errors.

Method

Participants

This study used the same 103 participants as Study 1.

Materials, design and procedure

We used a set of twelve behaviors adapted from [Graham et al. \(2009\)](#) to draw on all moral foundations (harm, fairness, ingroup, authority, and purity).⁴ The first part of the survey presented manipulated information about injunctive norms found in a survey in a foreign country. Each behavior was presented as being either “OK” or “not OK” (randomized per item) in the foreign group. For instance, one item read “To slap your child was typically judged as [OK/not OK] in the group. Do you agree with this judgment?” The main purpose of the question about agreement was to force participants to pay attention to the manipulated information without disclosing that a recall task would be given later in the study. However, responses to different items were also used as an indication of the moral status of each behavior in our sample.

The second part of the survey presented similarly manipulated information about typical judgments of descriptive norms in the foreign group, either “common” or “uncommon” (randomized per item). Again, participants were asked whether they agreed with this judgment.

The third part declared that we were interested in how well participants could recall the moral judgments of the foreign group. We asked for each behavior in turn if the typical judgment in the foreign group had been “OK” or “not OK”.

At subsequent debriefing, no participant claimed to have suspected that the information on typical judgments in a foreign group was fabricated.

⁴ In English translation, the items were: slap your child; not give money to beggars in need; laugh at a politician having personal problems; cut in with a friend on a long line for a concert; take more than half of some sweets you have purchased to share with a friend; report to the boss that your closest colleague does a poor job; help a friend cheat on a test; question the boss every time he has a bad idea; request of children to follow family traditions; not follow the law when it clashes with your moral persuasion; shake hands when you have a bad cold; have many sex partners in the same year.

Results

Judgments of stimuli

Across stimuli, the proportion of participants who judged a behavior as OK ranged from 6% to 77% and the proportion who judged a behavior as common ranged from 16% to 83%. Thus, stimuli covered behaviors of varying moral status and varying commonness.

Error in recall

For each participant we calculated the total rate of recall errors ($M = 0.21$, $SD = 0.18$) as well as separate rates for congruent recall errors ($M = 0.29$, $SD = 0.28$) and incongruent recall errors ($M = 0.12$, $SD = 0.19$). The difference between these error rates was statistically significant, $t(102) = 5.32$, $p < .001$, $d = 0.52$, a medium sized effect.

Individual differences

Recall that the IAT (Study 1) was conducted in the same lab sessions. We realized after the studies were conducted that it would be interesting to analyze whether individual differences in response latencies to congruent and incongruent pairings in the IAT would predict corresponding recall errors. Unfortunately, participants were not identifiable across the two studies because of the different software used. However, through an analysis of time stamps it was possible for the lab assistant to infer identities across studies with reasonable (but not absolute) certainty for a subset of 62 participants. For this subsample we conducted a tentative analysis.

As expected, there was a positive correlation ($r = .31$, $p = .013$) between the mean response latency for incongruent pairings and the rate of recall error for incongruent norm pairs. Similarly, there was a positive correlation between the mean response latency for congruent pairings and the recall error rate for congruent norm pairs ($r = .35$, $p = .005$). To rule out that these results simply reflected an overall correlation between response latencies and recall errors, we also computed the partial correlation between congruent response latencies and congruent recall error rates while controlling for the incongruent measures, and vice versa. Both correlations remained positive (partial r s $> .26$, p s $< .05$).

Discussion

We found information about a behavior’s descriptive norm to interfere with recall of information about the injunctive norm. The rate of congruent recall errors was more than twice as high as the rate of incongruent recall errors. This suggests that automatically associated injunctive norms can supersede earlier injunctive information in recall.

Further, our analysis of individual differences in recall errors for congruent and incongruent items indicated that they were predicted by response latencies for corresponding pairings in the implicit association task. Conclusions are limited by the uncertainty in the identification of participants across the two tasks, but the result is consistent with our assumption that the same cognitive association underlies both the observed pattern in response latencies and the observed pattern in recall errors.

Study 4

The previous study tested whether descriptive information interferes with recall of injunctive information. In this study we extended this test to include also the other direction, that is, whether injunctive information interferes with recall of descriptive

information. We also used a different set of stimuli to cover both the prescriptive and proscriptive domain.

Method

Participants

Participants were recruited online among US users of the Amazon Mechanical Turk at a compensation of one US dollar. There were 200 participants (55% female) with age ranging from 18 to 69 years ($M = 33$ years, $SD = 12$ years).

Materials, design and procedure

As our set of stimuli, we used a set of sixteen behaviors taken from the Moralism scale (Janoff-Bulman et al., 2009). Of the 16 items, half were taken from the prescriptive domain and half from the proscriptive domain.⁵ Participants were split into two conditions, *Injunctive Recall* ($N = 100$) and *Descriptive Recall* ($N = 100$).

The Injunctive Recall condition replicated Study 3: For each behavior participants were told that a certain moral judgment had been typical in a foreign group (either *OK* or *wrong* for proscriptive items; either *obligated* or *up to you* for prescriptive items), and asked whether they agreed with this judgment. The injunctive information was manipulated such that half of the participants were told one judgment was typical (say, 'OK'), and the other half was told the opposite judgment was typical (say, 'wrong').⁶ Following Study 3, participants were subsequently told that a certain frequency estimate had been typical in the foreign group (either *common* or *uncommon*),⁷ and asked whether they agreed with this estimation. The descriptive information was manipulated such that half of the participants were told the typical estimate was 'common' and the other half was told the typical estimate was 'uncommon'. Further, the descriptive and injunctive information was congruent – e.g., 'wrong' and 'uncommon' – for half the items (eight items out of sixteen, four in each domain), and incongruent for the other half of the items. Finally, participants were asked to recall for each item the typical moral judgment of the foreign group. Depending on the domain of the item, the response options were either *OK/wrong* or *obligated/up to you*.

So far we have described the Injunctive Recall condition. In the Descriptive Recall condition, the first and second parts above were interchanged so that descriptive information was presented first and the final part then asked for recall of this descriptive information, with response options common/uncommon.

At the end of the survey, participants were informed of the deception and the purpose of the study. This practice was followed in all on-line studies involving deception.

Results

Judgments of stimuli

For each behavior we calculated the proportion of participants who judged it as common. Proportions ranged from 27% to 74%, indicating that the set of stimuli included both common and uncommon behaviors.

⁵ An example of a prescriptive item: "Cory is in the supermarket, where he sees an elderly woman having trouble carrying her groceries. He is in a hurry and knows he could ignore her, but considers instead helping the elderly woman carry her groceries." An example of a proscriptive item: "Melanie and Scott have just bought a house in a quiet, middle-class neighborhood. The homes are not fancy, but are modest and well-kept. Melanie and Scott are considering ignoring the community and painting their house bright orange with green trim."

⁶ For instance, one item read: "To paint one's house bright orange with green trim in this situation was generally considered [OK/wrong]."

⁷ For instance, one item read: "To help the elderly woman with her groceries in this situation was generally considered to be a [common/uncommon] behavior."

Error in recall

The total rate of error in recall was computed for each participant as the proportion of items recalled incorrectly. There was no significant difference in error rates between injunctive recall ($M = 0.27$, $SD = 0.16$) and descriptive recall ($M = 0.30$, $SD = 0.17$), $t(198) = 1.30$, $p = .20$. Pooling injunctive recall and descriptive recall yielded a mean error rate of 0.29 ($SD = 0.17$). The rate of congruent recall errors ($M = 0.36$, $SD = 0.21$) was higher than the rate of incongruent recall errors ($M = 0.22$, $SD = 0.19$), $t(199) = 8.79$, $p < .001$, $d = 0.62$, a medium effect size. Separating conditions showed that the error rate difference was similar for injunctive recall ($M = 0.15$, $SD = 0.23$) and descriptive recall ($M = 0.13$, $SD = 0.22$), $t(198) = 0.58$, $p = .56$.

Discussion

In Study 3, using Swedish students in a lab, we found that descriptive information interfered with the recall of injunctive information. Here we replicated this finding in an online study with American users of the Amazon Mechanical Turk, using another set of stimulus behaviors. Thus, the finding seems robust with respect to various methodological variations.

Further, this study extended the investigation to include also interference of injunctive information on recall of descriptive information. The effect was the same in both directions, indicating that the association between injunctive and descriptive can interfere with memory of both types of norms.

Study 5

In Studies 3 and 4 we found that information about one type of norm, injunctive or descriptive, could interfere with recall of earlier information about the other norm type. A related prediction is that when only one type of norm information is given for a behavior, it may be misremembered as being about the other type. For instance, a behavior presented as "uncommon" may be recalled, incorrectly but congruently, as "wrong" and vice versa.

Socially undesirable behaviors were presented together with manipulated information about norms in a strange land. Each behavior was presented as either *common*, *uncommon*, *OK* or *wrong*. Participants were later asked to recall the items on a multiple-choice format. We predicted that participants would tend to mix up 'common' with 'OK', and 'uncommon' with 'wrong' (to a much larger extent than they would make other errors in recall).

Method

Participants

Participants were recruited online among US users of the Amazon Mechanical Turk at a compensation of one US dollar. There were 40 participants (30% female) with age ranging from 18 to 58 years ($M = 32$ years, $SD = 12$ years).

Materials, design and procedure

A stimulus set from the proscriptive domain was obtained by adapting 16 socially undesirable behaviors from a previous survey of injunctive and descriptive norms (Brauer & Chaurand, 2010).

The survey consisted of three parts. The first part ostensibly presented an excerpt from a guidebook to some unknown country. Participants were told that the task was to guess which country the excerpted guidebook was describing and therefore encouraged to read the excerpt carefully. The excerpt consisted of the list of 16 behaviors, four of which were described as 'uncommon', and equally many described as 'common', 'OK', or 'wrong' (counterbalanced across different versions). For example, one item read: *It is*

[common/uncommon/OK/wrong] to leave a store and let the door fall on the person behind.

After guessing on a country, participants were asked to write a short essay about their general experiences with guidebooks. This served as a distracter task before the subsequent recognition task, in which participants were presented with a list of 27 behaviors. This list included all of the 16 original behaviors, reordered and interlaced with 11 new behaviors from the same domain. For each behavior, participants were asked a multiple-choice question about how the behavior had been presented in the first part: as common, uncommon, OK, wrong, or not mentioned at all.

Coding

In the recognition task, all incorrect responses were coded as belonging to one of four categories of error: *opposite* (mixing up 'common' with 'uncommon' or 'OK with 'wrong'), *congruent* (mixing up 'common' with 'OK' or 'uncommon' with 'wrong'), *incongruent* (mixing up 'common' with 'wrong' or 'uncommon' with 'OK'), and *no recall* (if the response was 'not mentioned' or if no response option was chosen at all).

Results

Results from the recognition task are presented in Table 2. On average, participants correctly recalled slightly less than half of all presented norms: 44% of presented descriptive norms and 47% of presented injunctive norms. Our interest lies in when recall was incorrect. Ignoring the cases when there was no recall at all, we consider three categories of error (opposite, congruent, incongruent). If error in recall was the result of chance responses these categories of error should be equally common. Instead, as predicted, congruent errors dominated, see Table 2.

We conducted an ANOVA of the number of norm recall errors with type of presented norm (2 levels: descriptive or injunctive) and error category (3 levels) as within-subject factors. This analysis confirmed that the observed difference between error categories was highly unlikely to occur by chance, $F(2,78) = 32.4$, $p < .001$. The type of presented norm had no statistically significant main effect, nor did norm type interact with error category, all $ps > .5$.

Discussion

In this study we found a strong pattern in errors made in recall of information about norms: Errors were dominated by incorrectly recalling a norm as the congruent norm of the different type (such as recalling 'common' as 'OK' or vice versa). Thus, participants tended to mix up descriptive and injunctive concepts as predicted by the earlier established association between commonness and morality. The independence of norm type implies a bi-directional effect on memory of the common–moral association, consistent with the findings of Study 4.

Study 6

The previous study showed a strong bias for congruence in errors in recall of norm information in the prescriptive domain.

The next study investigated the same phenomenon in the prescriptive domain. It also added one methodological feature, based on the notion that congruent recall errors are not complete memory failures but rather successful recall of an associated norm. Thus, the dominance of congruent errors should be pronounced among responses where complete memory failure can be ruled out. In order to test this prediction, a task asking for free-recall of the list of behaviors was included before the recognition task. The purpose of the free-recall task was to establish a set of behaviors that were definitely recalled, for which complete memory failure could be ruled out. We predicted the bias for congruence in errors in the subsequent recognition task to be particularly strong in responses where complete memory failure could be ruled out.

Method

Participants

Participants were recruited online among US users of the Amazon Mechanical Turk at a compensation of one US dollar. There were 40 participants (38% female) with age ranging from 18 to 59 years ($M = 30$ years, $SD = 11$ years).

Materials, design and procedure

See the online Supplementary material for complete instructions to Study 6. The material in Study 5 was adapted to the prescriptive domain, with 16 socially desirable behaviors adapted from the Moralism scale (Janoff-Bulman et al., 2009). For example, one item in the "excerpted guidebook" read [It is common/It is uncommon/You are obligated/You are not obligated] to keep your garden looking proper and well-kept. The norm labels (common, uncommon, obligated, not obligated) were distributed and counterbalanced in the same way as in Study 5.

The same recognition task as in Study 5 was here preceded by a free-recall task. As mentioned above, its purpose was to identify responses for which complete memory failure could be ruled out. Participants were cued to recall as many as possible of the behaviors presented in the guidebook excerpt. In order not to interfere with the subsequent recognition task, participants were here told not to bother with recalling the norms but only the behaviors.

Coding

In the free-recall task, each item was binary coded as either successfully recalled (1) or not (0). The coder was unaware of the hypothesis. As a reliability test, a second coder recoded 25% of the data; codings were 98% identical.

Responses to the recognition task were coded in analogy with Study 5. For instance, *congruent* errors were those mixing up 'common' with 'obligated' or 'uncommon' with 'not obligated'.

Results

Results in the recognition task are presented in Table 3. First note that, on average, participants correctly recalled slightly more than half of all presented norms: 52% of presented descriptive norms and 58% of presented injunctive norms. Consistent with our aim to identify responses for which complete memory failure

Table 2

Study 5 average numbers of descriptive and injunctive items in each of several recall categories.

Presented norm	Correct recall	Opposite recall	Congruent recall	Incongruent recall	No recall	Total
Descriptive	3.48 ± 0.34	0.70 ± 0.16	1.77 ± 0.24	0.93 ± 0.19	1.13 ± 0.16	8
Injunctive	3.75 ± 0.35	0.50 ± 0.16	1.98 ± 0.25	0.70 ± 0.15	1.08 ± 0.18	8

Note: Every participant gave a total of 16 responses about their recall of 8 items presenting descriptive norms and 8 items presenting injunctive norms. Incorrect recall was categorized as either opposite (e.g., 'common' recalled as 'uncommon'), congruent (e.g., 'common' recalled as 'OK'), or incongruent (e.g., 'common' recalled as 'wrong'). Entries give mean ± se for the number of participants' responses that belonged to each category.

Table 3
Study 6 average numbers of descriptive and injunctive items in each of several recall categories.

Presented norm	Preceding free-recall	Correct recall	Opposite recall	Congruent recall	Incongruent recall	No recall	Total
Descriptive	Not successful	3.05 ± 0.29	0.50 ± 0.17	1.68 ± 0.23	0.28 ± 0.09	0.75 ± 0.13	6.25 ± 0.26
	Successful	1.15 ± 0.24	0	0.50 ± 0.11	0.05 ± 0.05	0.05 ± 0.03	1.75 ± 0.26
	Total	4.20 ± 0.32	0.50 ± 0.17	2.17 ± 0.27	0.32 ± 0.10	0.80 ± 0.13	8
Injunctive	Not successful	3.63 ± 0.35	0.30 ± 0.11	1.25 ± 0.23	0.48 ± 0.12	0.80 ± 0.17	6.45 ± 0.20
	Successful	1.05 ± 0.16	0.08 ± 0.04	0.35 ± 0.14	0.05 ± 0.03	0.03 ± 0.02	1.55 ± 0.20
	Total	4.68 ± 0.37	0.38 ± 0.13	1.60 ± 0.29	0.52 ± 0.13	0.83 ± 0.18	8

Note: Every participant gave a total of 16 responses about their recall of 8 items presenting descriptive norms and 8 items presenting injunctive norms. For each participant these items were further split into two classes based on whether the participant had successfully recalled the presented behavior in a preceding free-recall task. Incorrect recall was categorized as either opposite (e.g., 'common' recalled as 'uncommon'), congruent (e.g., 'common' recalled as 'obligated'), or incongruent (e.g., 'common' recalled as 'not obligated'). Entries give mean ± se for the number of participants' responses that belonged to each category.

can be ruled out, correct recall was especially high (66% for descriptive norms and 69% for injunctive norms) for those items where participants had previously succeeded in recalling the presented behavior in the free-recall task.

When descriptive norms were presented, the congruent error category was more than 4 times as large as any of the other two error categories – and at least 10 times as large when looking only at those items where participants had previously succeeded in recalling the presented behavior. For injunctive norms, the congruent error category was more than 3 times larger than any of the other categories – and more than 4 times larger when looking only at successfully free-recalled items.

We conducted an ANOVA of the number of norm recall errors with type of presented norm (2 levels: descriptive or injunctive), error category (3 levels) and behavior recall success (2 levels: successful or not) as within-subject factors. This analysis revealed that the observed difference between error categories was highly unlikely to occur by chance, $F(2,78) = 49.1$, $p < .001$, and similarly for the interaction between category and behavior recall success, $F(2,78) = 13.0$, $p < .001$. The type of the presented norm had no statistically significant main effect or interaction with any other factor, all $ps > .20$.

Discussion

This study replicated Study 5 using stimuli from the prescriptive domain. As in the previous study, the dominating recall error was to recall a norm as the congruent norm of the different type (such as recalling 'common' as 'obligated' or vice versa). As in Studies 4 and 5, the effect was bi-directional.

The present study also included a free-recall task of the behaviors used as stimuli. Consistent with the bias for congruent errors reflecting successful recall of associated norms, the bias was particularly strong among responses to items that were successfully free-recalled.

Study 7

The final two studies focused on whether perceptions of commonness influence own moral judgments of specific behaviors. In the proscriptive domain this is already supported by several studies: Own moral judgments of an undesirable behavior tend to be less harsh when the behavior was thought to be common (McGraw, 1985; Trafimow et al., 2001; Welch et al., 2005). For this reason, our studies focused on the prescriptive domain instead. As discussed in the introduction, it has previously been predicted that higher moral status would be attributed to uncommon behaviors than to common behaviors (McGraw, 1985; Reeder & Brewer, 1979). The common-moral association hypothesis makes the opposite prediction. We tested these competing predictions in a

study where participants were primed with information about descriptive norms in a foreign group before making own moral judgments of these behaviors. The study was conducted in two versions, differing in whether judgments were of the *moral obligation* to behave in a certain way or of the *moral character* of a person who behaves in that way. The purpose of this was to establish whether these different aspects of morality are influenced by commonness in the same way.

Method

Participants

The moral obligation version of the study was completed by 116 unique participants (37% female) with age ranging from 19 to 82 years ($M = 34$ years, $SD = 12$ years) recruited online among American users of Mturk at a compensation of 0.75 US dollar. The moral character version of the study run at a later date and was completed by 113 unique participants (36% female) with age ranging from 19 to 66 years ($M = 31$ years, $SD = 10$ years), recruited in the same way.

Materials

Participants were told they were taking part in a study of the moral status of various behaviors in different countries. We used nine behaviors adapted from the set of stimuli used in Study 3. All these behaviors could be judged as morally prescribed depending on what moral foundations one relies on (Graham et al., 2009).⁸ For each behavior, participants were presented with the ostensible descriptive norm in a foreign group. Descriptive norms were manipulated across behaviors and across different versions of the survey, such that each behavior was presented equally often as "uncommon", "neither common nor uncommon", and "common". In the moral obligation version, participants were then asked for their own judgment of the moral obligation of the behavior, on a scale from 0 = absolutely not morally obligated to 100 = absolutely morally obligated. In the moral character version, participants were then asked for their own judgment of the moral character of someone who behaves in the specified way, on a scale from 0 = absolutely not moral to 100 = absolutely moral.

Results

Moral obligations

Mean ratings of moral obligations varied across behaviors from 17.5 to 62.5. For each behavior we normalized ratings by subtracting the mean and dividing by the standard deviation. For each participant we calculated three averages of normalized ratings. As

⁸ Some examples of items: *to give money to beggars in need; never slapping one's child; to take only half of some goodies you have purchased to share with a friend.*

predicted, normalized moral obligation ratings were lower for behaviors presented as “uncommon” ($M = -0.13$, $SD = 0.62$) than for behaviors presented as “common” ($M = 0.11$, $SD = 0.70$), $t(115) = 3.28$, $p = .001$, $d = 0.30$. Consistent with a monotonous effect, ratings of behaviors presented as “neither common nor uncommon” had an intermediate level ($M = 0.01$, $SD = 0.58$). A repeated measures ANOVA with three levels of commonness labels indicated a significant effect on ratings of moral obligations, $F(2, 114) = 5.56$, $p = .005$, $\eta_p^2 = 0.089$.

Moral character

Mean ratings of moral character varied across behaviors from 26.4 to 87.8. Ratings were normalized and averaged as described above. Results were similar to the moral obligation version. Normalized moral character ratings were lower for behaviors presented as “uncommon” ($M = -0.06$, $SD = 0.55$) than for behaviors presented as “common” ($M = 0.13$, $SD = 0.73$), $t(112) = 2.66$, $p = .009$, $d = 0.25$. Ratings of behaviors presented as “neither common nor uncommon” had an intermediate level ($M = -0.04$, $SD = 0.64$). A repeated measures ANOVA with three levels of commonness labels indicated a significant effect on ratings of moral character, $F(2, 111) = 3.86$, $p = .024$, $\eta_p^2 = 0.065$.

Discussion

Recall that Study 2 presented a society’s descriptive norms for some generic behaviors and found these lead to explicit expectations about the injunctive norms in that society (and vice versa). Study 7 differed in two important respects, showing that information about another society’s descriptive norms for various *specific* behaviors influence participants *own* judgments of the moral obligation of these behaviors and of the moral character of someone who behaves in these ways.

Our findings are as predicted from the common–moral association hypothesis. It is noteworthy that they are in conflict with a theory from the literature on moral trait attribution, according to which moral attributions should be more extreme for uncommon behaviors (McGraw, 1985; Reeder & Brewer, 1979). In the prescriptive domain, higher moral status should be attributed to uncommon behaviors according to this theory. On the contrary, we found a positive effect of commonness on moral attributions (as well as on judgments of moral obligations).

Study 8

Study 7 investigated how moral judgments of well-known behaviors were influenced by priming participants with information about descriptive norms in a foreign group. To complement that study, we primed participants with information about their own group’s typical behavior in a novel situation, the economic game known as the *dictator game*. Previous research on the dictator game has shown that participants tend to adapt their behavior to descriptive norms presented to them (Bicchieri & Xiao, 2009). The common–moral association hypothesis predicts that also *moral judgments* of dictator game behavior will be sensitive to descriptive norm information.

Method

Participants

Two hundred participants (33% female) with age ranging from 18 to 75 years ($M = 31$ years, $SD = 12$ years) were recruited online among American users of Mturk at a compensation of 0.30 US dollar.

Materials

Participants were presented with two scenarios describing economic game experiments that we had ostensibly run in our lab. An initial scenario, identical for all participants, presented an investment game (details available on request). The real experiment was the second scenario, in which information about typical behavior was manipulated between subjects. The scenario read as follows:

“Two participants, anonymous to each other, are in different rooms in the lab. The experimenter gives 20 dollars to one of the participants with the following instructions: ‘You have two options. Either (a) you keep this money (20 dollars), or (b) you give it back to me, in which case I’ll give you 5 dollars instead and I’ll give 15 dollars to the other participant (who will otherwise get nothing).’ When we ran this experiment, about 10 [90] percent chose to give the money back (i.e., take 5 dollars instead of 20 dollars, so that the other participant would get 15 dollars). Thus, about 90 [10] percent chose to keep the 20 dollars (so that the other participant got nothing). In the present survey we are interested in your thoughts about the moral norms in this situation. How strong is the moral obligation to give the money back, as about 10 [90] percent did (and 90 [10] percent did not do)?”

Half of the participants received the information that giving back the money was uncommon (namely, that about 10% did it), the other half received the information that it was common (namely, that about 90% did it). The moral obligation to give back the money was estimated on a four-step Likert type scale anchored in 0 = *no obligation at all* and 3 = *very strong obligation*. Note that by asking about the obligation to give back the money we emphasized the prescriptive aspect of morality.

Results

As predicted, moral judgments of the obligation to give money in the dictator game tended to be stricter in the condition where giving was said to be common (rating 0: 34%, 1: 27%, 2: 23%, 3: 16%) than in the condition where giving was said to be uncommon (rating 41%, 1: 37%, 2: 17%, 3: 5%). The effect was small but statistically significant according to the Mann-Whitney U test, mean ranks 109.03 vs. 91.97, $p = .028$, $r = 0.15$.

Discussion

In this study we found a tendency for judgments of the moral obligation to give in the dictator game to be sensitive to information about how most people behave in the game. This indicates that commonness influences own moral judgments in a novel situation.

General discussion

A key feature of all human societies is that behavior is governed by social norms to some extent. In this paper we have addressed a fundamental question about the cognition of social norms: Is information about descriptive and injunctive norms processed separately? Based on the typical co-occurrence of descriptive and injunctive norms in real life we proposed the common–moral association hypothesis, according to which commonness and morality should tend to be automatically associated. This hypothesis, as well as various implications of it, was supported in a number of experiments. A strong association was found using implicit as well as explicit measures (Studies 1–2). In memory experiments we found a systematic tendency for injunctive and descriptive information to be mixed up in recall (Studies 3–6). Finally, priming participants with descriptive information about specific behaviors

influenced their moral judgments of these behaviors (Studies 7–8). Taken together, these studies establish the existence of an interesting psychological phenomenon. They also open up several lines of inquiry for the future.

Our findings are consistent with the social intuitionist thesis that moral judgments largely rely on automatic evaluations, in addition to rational thought and principles learnt from moral authorities, etc. (Haidt, 2001, 2007). The existence of several sources of moral judgments suggests future research on moderators of the influence of the common–moral association. Following the APE model (Gawronski & Bodenhausen, 2011) we should expect this influence to be moderated by the extent to which a person's morality has other important sources (such as rational thought or religious authorities). We should also expect the influence to vary across different behaviors (Trafimow et al., 2001). Following the APE model we should expect the influence of the common–moral association to be greatest for behaviors to which no other beliefs about morality strongly apply.

The latter point suggests the theoretical possibility that the common–moral association can contribute to the emergence of new injunctive norms. Assuming that automatic common–moral associations are common in a population (as Study 1 indicated), the same injunctive inferences from commonness of behaviors may be made by many individuals in the group. These mass inferences could potentially aggregate into a new injunctive norm shared by the group. In other words, for a behavior that was previously not morally loaded, a group might start thinking of it as morally obligated if it is common, or as morally suspect if it is uncommon. A related argument was made already by Kelley (1971). Nonetheless, this prediction remains to be tested.

Obviously, all common behaviors do not become morally obligated. One reason for this, we speculate, is that a behavior may have to reach a high degree of commonness for the common–moral association to be activated. This is consistent with research on conformity where perceived consensus has been found to play an important role (Cialdini & Goldstein, 2004). Thus, it is typically not sufficient that a behavior is in moderate majority for spontaneous conformity to be elicited (Claidière & Whiten, 2012; Eriksson & Coultas, 2009). The effects of different degrees of commonness on the common–moral association should be similarly investigated.

Once norms are established, for whatever reason, the common–moral association would also potentially play a role in sustaining norms. For example, consider societies where corruption is endemic. It is notoriously difficult to eradicate corruption in such societies (Persson, Rothstein, & Teorell, 2013). One reason might be that when most people are corrupt, anti-corruption campaigns that highlight this fact will be in vain because the automatic association will be that corrupt behavior is not so wrong after all.

In sum, all studies in this paper were focused on establishing the common–moral association at the individual level. Future research on the common–moral association should investigate (a) moderators of its effects on individual moral judgments and (b) its potential to cause emergence of shared injunctive norms in groups and maintenance of norms when the group faces norm-conflicting information.

Finally, note that our memory experiments used a novel paradigm for recall errors related to automatic associations. This paradigm could be applied to the study of other associations than the common–moral association.

Acknowledgments

This work was supported by the Swedish Research Council grants 2009–2390 and 2009–2678. Thanks to Per Andersson for col-

lecting the IAT data. We are grateful to Jonathan Haidt, Robert Cialdini, Andrew Whiten, Ying-yi Hong, and several anonymous reviewers for encouragement and many helpful comments.

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.obhdp.2014.09.011>.

References

- Banaji, M. R., & Bhaskar, R. (2000). Implicit stereotypes and memory: The bounded rationality of social beliefs. In D. L. Schacter & E. Scarry (Eds.), *Memory, brain, and belief* (pp. 139–175). Cambridge, MA: Harvard University Press.
- Berkowitz, A. D. (2004). *The social norms approach: Theory, research and annotated bibliography*. Higher Education Center for Alcohol and Other Drug Abuse and Violence Prevention. US Department of Education.
- Bicchieri, C. (2006). *The grammar of society. The nature and dynamics of social norms*. Cambridge: Cambridge University Press.
- Bicchieri, C., & Xiao, E. (2009). Do the right thing: But only if others do so. *Journal of Behavioral Decision Making*, 22, 191–208.
- Borsari, B., & Carey, K. B. (2001). Peer influences on college drinking: A review of the research. *Journal of Substance Abuse*, 13, 391–424.
- Borsari, B., & Carey, K. B. (2003). Descriptive and injunctive norms in college drinking: A meta-analytic integration. *Journal of Studies on Alcohol*, 64, 331–341.
- Boyer, W. (2012). Cultural factors influencing preschoolers' acquisition of self-regulation and emotion regulation. *Journal of Research in Childhood Education*, 26, 169–186.
- Brauer, M., & Chaurand, N. (2010). Descriptive norms, prescriptive norms, and social control: An intercultural comparison of people's reactions to uncivil behaviors. *European Journal of Social Psychology*, 40, 490–499.
- Buhrmester, M. D., Kwang, T., & Gosling, S. D. (2011). Amazon's Mechanical Turk: A new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science*, 6, 3–5.
- Cialdini, R. B. (2003). Crafting normative messages to protect the environment. *Current Directions in Psychological Science*, 12, 105–109.
- Cialdini, R. B., & Goldstein, N. J. (2004). Social influence: Compliance and conformity. *Annual Review of Psychology*, 55, 591–621.
- Cialdini, R. B., Kallgren, C. A., & Reno, R. R. (1991). A focus theory of normative conduct: A theoretical refinement and reevaluation of the role of norms in human behavior. *Advances in Experimental Social Psychology*, 24, 1–243.
- 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, 1015–1026.
- Claidière, N., & Whiten, A. (2012). Integrating the study of conformity and culture in humans and non-human animals. *Psychological Bulletin*, 138, 126–145.
- Eriksson, K., & Coultas, J. C. (2009). Are people really conformist-biased? An empirical test and a new mathematical model. *Journal of Evolutionary Psychology*, 7, 5–21.
- Gawronski, B., & Bodenhausen, G. V. (2006). Associative and propositional processes in evaluation: An integrative review of implicit and explicit attitude change. *Psychological Bulletin*, 132, 692–731.
- Gawronski, B., & Bodenhausen, G. V. (2011). The associative-propositional evaluation model: Theory, evidence, and open questions. *Advances in Experimental Social Psychology*, 44, 59–127.
- Gawronski, B., & Sritharan, R. (2010). Formation, change, and contextualization of mental associations: Determinants and principles of variations in implicit measures. In B. Gawronski & B. K. Payne (Eds.), *Handbook of implicit social cognition: Measurement, theory, and applications* (pp. 216–240). New York, NY: Guilford Press.
- Graham, J., Haidt, J., & Nosek, B. A. (2009). Liberals and conservatives rely on different sets of moral foundations. *Journal of Personality and Social Psychology*, 96, 1029–1046.
- Greenwald, A. G., McGhee, D. E., & Schwartz, J. K. L. (1998). Measuring individual differences in implicit cognition: The implicit association test. *Journal of Personality and Social Psychology*, 74, 1464–1480.
- Greenwald, A. G., Nosek, B. A., & Banaji, M. R. (2003). Understanding and using the implicit association test: I. An improved scoring algorithm. *Journal of Personality and Social Psychology*, 85, 197–216.
- Haidt, J. (2001). The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review*, 108, 814–834.
- Haidt, J. (2007). The new synthesis in moral psychology. *Science*, 316, 998–1002.
- Janoff-Bulman, R., Sheikh, S., & Hepp, S. (2009). Proscriptive versus prescriptive morality: Two faces of moral regulation. *Journal of Personality and Social Psychology*, 96, 521–537.
- Kelley, H. H. (1971). Moral evaluation. *American Psychologist*, 26, 293–300.
- Lapinski, M. K., & Rimal, R. N. (2005). An explication of social norms. *Communication Theory*, 15, 127–147.
- McGraw, K. (1985). Subjective probabilities and moral judgments. *Journal of Experimental Social Psychology*, 21, 501–518.

- Nosek, B. A., Greenwald, A. G., & Banaji, M. R. (2005). Understanding and using the implicit association test: II. Method variables and construct validity. *Personality and Social Psychology Bulletin*, *31*, 166–180.
- O’Gorman, R., Wilson, D. S., & Miller, R. R. (2008). An evolved cognitive bias for social norms. *Evolution and Human Behavior*, *29*, 71–78.
- Olson, M. A., & Fazio, R. H. (2001). Implicit attitude formation through classical conditioning. *Psychological Science*, *12*, 413–417.
- Paolacci, G., Chandler, J., & Ipeirotis, P. G. (2010). Running experiments on Amazon Mechanical Turk. *Judgment and Decision Making*, *5*, 411–419.
- Parks, C. D., & Stone, A. B. (2010). The desire to expel unselfish members from the group. *Journal of Personality and Social Psychology*, *99*, 303–310.
- Perkins, H. W., & Berkowitz, A. D. (1986). Perceiving the community norms of alcohol use among students: Some research implications for campus alcohol education programming. *Substance Use & Misuse*, *21*, 961–976.
- Persson, A., Rothstein, B., & Teorell, J. (2013). Why anticorruption reforms fail—Systemic corruption as a collective action problem. *Governance*, *26*, 449–471.
- Perugini, M., & Leone, L. (2009). Implicit self-concept and moral action. *Journal of Research in Personality*, *43*, 747–754.
- Reeder, G. D., & Brewer, M. B. (1979). A schematic model of dispositional attribution in interpersonal perception. *Psychological Review*, *86*, 61–79.
- Schacter, D. L. (1999). The seven sins of memory: Insights from psychology and cognitive neuroscience. *American Psychologist*, *54*, 182–203.
- Thøgersen, J. (2008). Social norms and cooperation in real-life social dilemmas. *Journal of Economic Psychology*, *29*, 458–472.
- Trafimow, D., Reeder, G. D., & Blissing, L. M. (2001). Everybody is doing it: The effects of base rate information on correspondent inferences from violations of perfect and imperfect duties. *Social Science Journal*, *38*, 421–433.
- Welch, M. R., Xu, Y., Bjarnason, T., Petee, T., O’Donnell, P., & Magro, P. (2005). “But everybody does it . . .”: The effects of perceptions, moral pressures, and informal sanctions on tax cheating. *Sociological Spectrum*, *25*, 21–52.