

Cognitive Properties of Norm Representations

Bertram F. Malle

(bfmalle@brown.edu)

Department of Cognitive, Linguistic,
& Psychological Sciences, Brown
University, 190 Thayer Street,
Providence, RI 02906

Joseph L. Austerweil

(austerweil@wisc.edu)

Department of Psychology, University
of Wisconsin, 1202 West Johnson
Street, Madison, WI 53706

Vivienne B. Chi

(vivienne_chi@brown.edu)

Department of Cognitive, Linguistic,
and Psychological Sciences, Brown
University, 190 Thayer Street,
Providence, RI 02906

Yoed N. Kenett

(yoedk@technion.ac.il)

Faculty of Industrial
Engineering &
Management, Technion
3200003, Haifa, Israel

Emorie D. Beck

(emorie.beck@gmail.com)

Medical Social Sciences,
Feinberg School of Medicine,
Northwestern University, 420
E. Superior Street, 60611
Chicago IL

Stuti Thapa

(sthapama@purdue.edu)

Department of Psycho-
logical Sciences, Purdue
University, West Lafayette,
703 3rd Street, IN 47907

Mowafak Allaham

(mmallaham1@gmail.com)

McCormick School of
Engineering, Northwestern
University, 70 Arts Circle
Drive, Evanston IL 60208

Abstract

Norms are central to social life. They help people select actions that benefit the community and facilitate behavior prediction and coordination. However, little is known about the cognitive properties of norms. Here we focus on norm activation, context specificity, and how those properties differ for the two major types of norms: prescriptions and prohibitions. In two studies, participants are exposed to a variety of contexts by way of scene images and either (a) freely generate norms that apply to the context or (b) decide whether each of a series of candidate norms applies to a given context. Across both studies, people showed high levels of context specificity and fast norm activation, and these properties were substantially stronger for prescriptions than for prohibitions.

Keywords: Social norms, moral norms, negation, cognitive structure, network, deontics

Introduction

Norms are a crucial part of social life. They guide action that benefits the community (Goldstein & Cialdini, 2007), they offer strong priors for predicting others' behavior (Cialdini, Reno, & Kallgren, 1990), and they coordinate and facilitate social interaction (Ullmann-Margalit, 1977). Research on social norms has primarily focused on identifying different types of norms and their influence on behavior (Cialdini et al., 1991; Horne & Mollborn, 2020). Few studies have explored the *cognitive structure* of norms—how they are represented and organized in the mind and how they are activated in specific contexts. For example, when we stand behind a bookshelf in a library we keep the level of noise down even if no other person is around, simply because the norm is to be quiet in libraries (Aarts & Dijksterhuis, 2003). Can a physical stimulus such as bookshelves directly activate such a context-specific norm? And what does such activation tell us about the cognitive structure of norms?

The literature on scene recognition provides initial insights into the power of physical contexts to activate norm representations (Rifkin, 1985; Tversky & Hemenway, 1983). For example, Tversky and Hemenway (1983) asked

participants to list activities that were “appropriate” to perform in various scenes. Participants generated systematic and organized activities suitable for those scenes, and while the authors interpreted these results as revealing properties of *scene categorization*, we believe that they also reveal properties of *norm representations*—the mental structures that guide socially appropriate actions in specific contexts.

Our goal is to make norm representations measurable and study some of their cognitive properties. We ground our studies in a working definition of norms that integrates previous analyses (Andrighetto et al., 2013; Bicchieri, 2006; Brennan et al., 2013; Cialdini et al., 1991; Malle et al., 2017):

A norm is an instruction, in a given community, to (not) perform an action in a specific context, provided that a sufficient number of community members demand of each other to follow the instruction and do in fact follow it.

This working definition highlights several features. First, norms come in two major types: as prescriptions (instructions to act) and as prohibitions (instructions to not act). Second, they govern and direct actions in specific communities and conditional on specific contexts (Horne & Mollborn, 2020). Third, norms exert a demand on people, and people will comply with the norm to varying degrees (Cialdini et al., 1990; Lindenberg, 2013). Thus, *norm representations* must code for these numerous features: the relevant community, norm type, action content, context conditions, and more. Here we focus on the properties of context specificity, norm activation, and how those properties differ for the two major types of norms, prescriptions and prohibitions.

Humans interact in a vast number of settings that vary in location, time, goals, roles, and social relations—constituting different “contexts.” How does the mind store norms for these contexts? The solution may lie in a tight link between specific contexts and specific bundles of norms. Humans are excellent at context-dependent learning and memory (Gould & Bevins, 2012; Rovee-Collier & Haye, 1987); and even infants demonstrate context-specific rule learning (Werchan et al., 2015; Wyman et al., 2009). We therefore examine to

what degree, and how quickly, different contexts can uniquely activate their associated norms.

These context-specific activation patterns may differ for the two major types of norms: prescriptions and prohibitions. The two are typically framed as tagging actions as, respectively, desirable or aversive (e.g., Cushman, 2013). In this framing, we might expect a negativity effect (Baumeister et al., 2001) in that prohibitions, compared to prescriptions, are more tightly linked to their contexts and thus cognitively accessible. Alternatively, the distinct functional roles of the two norm types (Janoff-Bulman et al., 2009) may have an opposite impact on their cognitive accessibility. Because prescriptions guide appropriate action, the context a person enters may directly trigger a set of prescription norms and thus enable appropriate courses of action. Prohibitions, by contrast, block undesirable courses of action. It would be cognitively inefficient if contexts activated a large number of prohibitions to stop actions that the person would not have thought of in the first place; instead, prohibitions may be activated when the person considers an action that, in the current context, must be suppressed. For example, because some beaches permit nudity, and a given beach may enable this possible action, explicit prohibitions against nudity will be needed for certain beaches. However, no beaches allow murder, and so no explicit prohibition is necessary. On this account, when a person enters a context, prohibitions would on average be less accessible than prescriptions.

To examine these questions of context-specific activation, and the relative accessibility of prohibitions and prescriptions, we developed two new experimental tasks, corresponding roughly to the recall-recognition distinction in the memory literature. The first is the *norm generation task*, which borrows from research on scene categorization (Tversky & Hemenway, 1983) and on semantic memory representations (Kumar, 2021). In this procedure, we exposed people to a set of “contexts” (pictures of everyday scenes) and asked them to generate norms that are relevant to that context (e.g., “What are you supposed to do here?”). The second is the *norm recognition task*. It also relies on the power of pictured scenes to invoke contexts but borrows from the recognition memory and signal detection theory traditions to assess how quickly and accurately people recognize experimenter-presented norms as relevant to a particular context, in contrast to those that are not relevant. This task affords analyses of recognition sensitivity independent of any speed-accuracy tradeoffs.

Study 1

Participants

121 participants were recruited via Amazon’s Mechanical Turk and randomly assigned to a norm type (prescription, permission, prohibition) and one of two context sets. Two individuals had duplicate IDs, leaving 119 for analysis. Mean age was 34.7 (SD = 10.5), 42.4% were female, 82.2% White, 9.3% Asian American, 5.9% African American, 2.5% Latinx.

Methods

Materials To select a representative sample of everyday contexts we considered a corpus of 40 physical contexts that are typical of contemporary culture (e.g., classroom, elevator, restaurant, library) and classified them on the dimensions of (a) public vs. private; (b) work vs. recreational; and (c) human-made vs. natural. We decided to examine only public contexts because they have the greatest likelihood of eliciting socially shared norms, and we selected scenes in equal numbers from combinations of the human-made/natural work/recreational categories. We chose ones that were likely to be familiar to a broad range of people, were distinct from each other, and could be easily captured by a picture. To ensure generalizability we created two sets of four scenes (Vegetable Harvest, Jogging Path, Restaurant, Public Restroom; Cave, Beach, Boardroom, Library).

Procedure After providing consent, participants received instructions, a practice trial, and four experimental trials. Each trial featured a different scene. Participants were assigned either the four scenes from set 1 or the four scenes from set 2, and the four scenes were presented in randomized order. In all four trials, participants were asked to generate norms of one and the same type, elicited by a prescription probe (*What are you supposed to do here?*), or a permission probe (*What are you allowed to do here?*), or a prohibition probe (*What are you not allowed to do here?*).

The norm-governed actions for a given scene can be very different depending on one’s role (e.g., waiter vs. customer in a restaurant). Each trial therefore began with a preparatory sentence (presented for 2 sec) that specified the role participants were to adopt in the scene (e.g., for the harvesting scene: “You are a farm worker doing this:…”). Next, the scene image appeared below the preparatory statement, and the norm probe question (e.g., *What are you supposed to do here?*) appeared below the picture. Participants were asked to type answers to the probe question into a text box below. They were instructed to type as many answers as possible for up to 60 seconds and to hit *Enter* after each answer.

Because people express the same norm in linguistic variants (e.g., *be quiet, be silent, don’t talk*), two coders inspected the typed norms for each of the eight contexts and assigned a common norm category to responses with highly similar meaning (interrater agreement > .80%). However, responses that could refer to different norms (e.g., *serve customers vs. deliver food*) were kept separate.

Results

All reported results were consistent across the two stimulus sets, with some expected variation from scene to scene. We report aggregated analyses across all scenes or mixed-effects analyses with scene as a random effect, allowing generalization to similar other scenes. Controlling for age and gender did not qualify any of the results.

Norms Generated People generated on average 6.11 norms, with substantial individual differences and some scene

differences. Two thirds of responses occurred in the first 30 seconds, and the vast majority of people did not use the full 60 seconds to write norm entries. The number of entries people did produce showed a marked effect of norm type. A random-effects analysis, predicting number of generated norms from norm type as a fixed effect and intercepts for scene and subject as random effects, confirmed that people generated fewer prohibition norms ($M = 4.95, SD = 2.02$) than prescription norms ($M = 6.37, SD = 2.98$), $t(114) = 2.78, p = .006$, while permission norms ($M = 7.03, SD = 3.25$) did not differ from prescription norms, $p = .21$. Aggregating across four scenes per person, the effect size of the prescription-prohibition difference was $d = 0.66$.

We also examined whether the norm types differed in the length of their verbal entries (measured in number of words per entry). Though several variables significantly predicted entry length (e.g., older participants had longer entries; earlier entries were shorter), norm type did not predict length. Further, when we controlled for length of entry, participants with shorter entries generated more norms, but people still generated significantly fewer norms of prohibition (estimated marginal $M = 5.38$) than norms of prescription ($M = 6.87$) or permission ($M = 7.45$), $F(2, 114) = 8.4, p < .001$.

Norm agreement Next, we computed the “community agreement” for each norm, defined as the percentage of participants (in a given norm type \times scene combination) who generated that norm for the invoked context. We identified the top-10 norms with highest agreement for each scene and each norm type. Figure 1 shows the resulting agreement gradients, averaged across scenes, from the most highly agreed-on norms (rank = 1), declining in a mixed linear and quadratic pattern. This pattern was consistent across scenes but, as Figure 1 shows, displayed substantially higher agreement for prescriptions than prohibitions (permissions followed the prescription pattern).




Figure 1. Between-participant agreement in the norms they spontaneously generate, ordered from 1st to 10th-most agreed-on norms, averaged across eight contexts.

Predicting agreement from norm type and rank as fixed effects and varying intercepts across scene and varying slopes for norm types across scenes as random effects, we found that, aside from the obvious rank effect, prescription norms had higher average agreement ($M = 0.37$) than prohibition norms ($M = 0.27$), $t(7) = 3.4, p = .011$, while permission norms ($M = 0.38$) did not differ from prescription norms, $p = .59$. In addition, agreement differences by norm type were primarily located in the most frequently mentioned norms, reflected in an interaction term of norm type and rank, such that prescriptions (and permissions) showed a steeper decline, with significant linear and quadratic contrasts, $ts > 2.9, ps < .004$.

Context Specificity We operationalized context specificity with respect to the top-10 norms of each context (scene). A norm is maximally context-specific if it occurs only in the context in response to which it was generated and does not recur among the top 10 norms of any of the other seven contexts. The more often a norm recurs in other contexts, the more it loses specificity. A conservative index of specificity is to count, across all 8 contexts, how many top-10 norms appeared uniquely in a single context and, if some recurred, how often they did. Among 80 prescription norms, 87.5% were unique; five norms each recurred in one other context. Among prohibitions norms, 48.8% were unique, a rate that was substantially lower than that for prescriptions, $\chi^2 = 24.7, p < .001$. Among permissions, 82.5% were unique, with norms recurring one to three times.

Structural Indications We explored whether the order in which people generated norms may reflect the importance or centrality of a norm, indexed by its level of community agreement. The results showed that more consensual norms were reported earlier across the top-10 norms of all three norm types ($r = -.26$).

Contents of Prescriptions vs. Permissions Prescriptions and permissions overlap conceptually because all prescribed actions are also permitted (though some permitted actions are not prescribed). This overlap was reflected in permissions and prescriptions showing high similarity in numbers of norms generated, community agreement, and levels of context specificity. However, the particular permissions and prescriptions people generated were not identical. Of the 80 top-10 prescriptions, 37.5% did not recur in permissions for the same contexts, and of the 40 top-5 prescriptions, 50% of prescriptions did not recur in permissions.

Discussion

Study 1 found that the norm generation task can activate norm representations that (a) show considerable agreement; (b) demonstrate substantial context specificity; and (c) hint at structural organization such that more agreed-upon norms are accessed earlier. Importantly, the number of norms generated, level of community agreement, and level of context specificity were consistently lower for prohibitions than for prescriptions and permissions. Permissions and

prescriptions were similar in activation patterns but distinguishable by their specific norm contents.

Replication. We sought to replicate these results in an additional study, but because of space constraints we only summarize its results. The study contained a number of improvements. First, we collected a larger sample of participants (120 per norm type) and allotted 30 seconds per context. We found the same decrement among prohibition norms for the number of generated norms, community agreement, and context specificity. Second, we used different verbal probes to elicit norm types. In particular, we removed a concern with Study 1's prohibition probe (*not allowed*), which, as a negation, may have been processed more slowly and thereby reduced the accessibility of prohibition norms. Using the probe *forbidden*, the new study replicated all prescription-prohibition differences. We also found that the patterns of results for prescription norms were consistent across three eliciting probes (*should*, *obligated*, *required*). Finally, we explored two non-normative comparisons: one to elicit what people would *want to do* in the given contexts and the other to elicit what *comes to mind* in the given contexts. The desire probe elicited overlapping but clearly distinguishable sets of actions compared with prescriptions, suggesting that the generation task does not simply elicit desirable actions but rather norm-conforming actions. The association probe elicited very different responses from those of the norm probes. For example, for the restaurant scene, only 1 out of the top 10 prescription norms shared content with a top-10 response from the association probe. Thus, the generation task does not elicit semantic or image associations but genuine norm representations.

Study 2

One hypothesis to explain the unique features of prohibitions, documented in the norm generation task, is that contexts less readily activate prohibitions compared with prescriptions. Instead, people must often consider an action first before a prohibition may stop them from acting. As a result, prohibitions are not linked as tightly with specific contexts (i.e., show less context specificity), they are more difficult to generate, and, due to these variabilities, people also agree less in the prohibition norms they generate.

To test this hypothesis we developed the *norm recognition* task that does not require people to generate norms on their own but rather presents them with candidate actions one by one and asks them to judge whether the particular action falls under a selected norm type (e.g., “Are you forbidden from doing this here?”). If a given scene activates context-specific norm representations, then people should be fast at recognizing offered candidates that match the pre-activated representations.

Importantly, we elicited these recognition judgments both for actions that were “local” to the given context (governed by a norm frequently generated for that context in Study 1) and nonlocal actions (governed by frequent norms generated in *other* contexts). If a context (scene image) immediately

activates its associated norms (as it may be for prescriptions), then local candidate actions governed by these norms should be primed and therefore quickly and accurately detected as falling under the context-specific norms. Nonlocal candidate actions, by contrast, would have to be retrieved, considered, and often rejected—and even if accepted, the processing will take longer than for local actions. For prohibitions, the process may be different. If a context does not immediately activate “its” prohibition norms but instead requires people to consider candidate actions one by one (e.g., when presented in the experiment), people would have to actively retrieve from memory whether the action is forbidden in the context and then make the appropriate judgment, leading to longer response times. Assuming such contrasting processes, people would more quickly detect local than nonlocal prescription norms, because the local ones are immediately activated by the scene image. By contrast, if both local and nonlocal prohibitions are retrieved after being “offered” (in a given experimental trial), no such speed advantage of local over nonlocal norms should exist.

Methods

Participants A total of 360 participants were recruited through Amazon Mechanical Turk and compensated \$0.75. Of those, 350 completed the study, and additional ones were excluded because of invalid data (see below).

Design Participants were randomly assigned to one of three norm types, implemented by one of two probe variants: “Are you *supposed to do this here?*” and “*Should you do this here?*” for prescriptions; “Are you *allowed to do this here?*” and “Are you *permitted to do this here?*” for permissions; and “Are you *not allowed to do this here?*” and “Are you *forbidden to do this here?*” for prohibitions. Half of participants were randomly assigned to see scene pictures from Set 1, the other half from Set 2, as in Study 1.

Procedure After general instructions and a practice round, participants worked through four blocks of norm recognition trials, each block consisting of a scene picture shown on the screen, the assigned norm probe below it, and a random ordering of 14 actions one might take in the scene. For each action, one after another, people answered the norm probe with “Yes” or “No” (e.g., “Are you supposed to do this here?”—*Swimming*). There were no time restrictions, and people received no feedback. Key assignment (F or J for “Yes” or “No”) was counterbalanced. A response triggered display of the next action phrase. Participants worked with the same norm type probe for the whole four blocks.

Materials We worked with the same two sets of scene images (four each) as in Study 1. Among the 14 actions accompanying each scene image, 7 were “local” to this scene (for a given norm type)—that is, most frequently generated in Study 1 as normative actions for the particular scene and norm type. The other 7 were “nonlocal”—selected from among the most frequently generated normative actions for the *other* seven scenes, again for the given norm type. In a

few cases, a selected local norm had to come from just outside the top 7 in Study 1 (in order to avoid repetitions), but the nonlocal norms were never outside the top 7. For each set of four scenes, the 28 local norms (7 per four contexts) were unique, but the nonlocal norms had to be reused across contexts. That was because some candidate nonlocal norms were confusable with a local norm, and several candidate nonlocal norms were so unique that they would have been implausible to consider in any but their local context (e.g., checking out books in a public bathroom). These necessary repetitions were equally distributed over norm types and stimulus sets.

Data preparation We used 400 ms as the reaction time cut-off below which it would be impossible to read a stimulus and select a response (Malle & Holbrook, 2012), and for the upper limit, we selected 5000 ms. (All reported results held robustly both with higher and lower cut-offs, as well as under various transformations, such as square root or inverse of reaction times). Almost half of the resulting out-of-range values came from 11 individuals who each had more than half such values and were eliminated from analysis. Among the remaining 339 participants, only 2.5% of all reaction times were out of range. We then examined “Yes” response rates for the 28 local items in each set to detect participants who were noncooperative or misunderstood instructions. Local items were, by design, expected to elicit high rates of “Yes” responses, and 95% of participants indeed answered Yes to more than half of the local items. The remaining 17 participants responded Yes to 5 or fewer items (out of 28), and they were eliminated from analysis.

Results

We first examined whether norm endorsements (Yes responses to the norm probes) were higher for local than for nonlocal actions and whether this locality effect varied by norm type. A generalized mixed-effects logistic model, predicting endorsements (1, 0) from norm type and locality as fixed effects and varying intercepts across participants and scenes as random effects, confirmed an overall locality effect: local norms were endorsed far more often (estimated marginal $M = 0.94$) than nonlocal norms ($M = 0.62$), $z = 42.9$, $p < .001$. However, the locality effect was much stronger for prescriptions ($M_s = 0.95$ vs. 0.51) than prohibitions ($M_s = 0.83$ vs. 0.68), $z = 18.5$, $p < .001$. Permissions ($M_s = 0.98$ vs. 0.65) were also stronger than prohibitions, $z = 17.0$, $p < .001$, but did not significantly differ from prescriptions.

We next examined response times. Of particular interest was the interaction between norm type and the speed of norm endorsements (“Yes” responses) for local vs. nonlocal norms. By hypothesis, a locality speed advantage suggests that the scene context directly activated relevant norms, whereas a lack of such an advantage suggests a retrieval-driven response. A mixed-effects model, predicting “Yes” reaction times from norm type and locality as fixed effects and intercepts across participants and across scenes as random effects, confirmed that people were faster at accepting local

(estimated marginal $M = 1235$ ms) than nonlocal norms ($M = 1341$ ms), $t(13081) = 8.9$, $p < .001$. However, whereas prescriptions showed a strong locality speed advantage (1087 ms vs. 1330 ms), as did permissions (1095 ms vs. 1318 ms), $t_s > 11.5$, $p_s < .001$, prohibitions showed a surprising reversal (1522 ms vs. 1374 ms), $t = 6.9$, $p < .001$ (see Figure 2). These patterns were highly consistent across the two stimulus sets.




Figure 2. Average speed of endorsing local actions and nonlocal actions as falling under a given norm type (“Yes” reaction times with 95% CIs)

Turning to the speed of rejecting the context-appropriateness of norms (“No” responses), we found that people were overall faster at rejecting nonlocal norms ($M = 1552$ ms) than local norms ($M = 1738$ ms); $t(3188) = 4.7$, $p < .001$. But, as Figure 3 shows, this locality effect in rejections occurred in a very different speed range for prescriptions ($M = 1443$ ms) compared with prohibitions ($M = 1779$ ms), $t(436) = 5.5$, $p < .001$, and also compared with permissions ($M = 1715$ ms), $t(808) = 4.7$, $p < .001$.

All results were consistent across the different probes for each norm type. Importantly, although the prohibition probe *not allowed* yielded about 100 ms longer RTs than the probe *forbidden*, both probes were substantially slower than all prescription and permission probes.




Figure 3. Average speed of rejecting local and nonlocal actions as falling under a given norm type (“No” reaction times with 95% CIs)

Discussion

In a norm recognition task, people clearly discriminated between local norms (i.e., norms frequently generated for a displayed context) and nonlocal norms (i.e., norms frequently generated for *other* contexts), thus providing further evidence for the context specificity of norms. But this evidence is much stronger for prescriptions (and permissions) than for prohibitions, whose small locality effect is consistent with their lower context specificity in Studies 1 and 2. Most importantly, people were faster at recognizing (i.e., “Yes” reaction times) local norms than at recognizing nonlocal norms, again only for prescriptions and permissions. These faster endorsements of local norms may result from contexts (scene images) activating relevant prescription (and permission) norms right when the image appears—perhaps due to a norm analog of affordances (Gibson, 1966), where norms reveal themselves in some subset of the perceptible features of the physical environment. Probing local norms therefore has the speed advantage of matching experimenter-presented norms to already pre-activated norms.

For prohibitions, by contrast, people appear to retrieve each norm they are asked about (whether local or nonlocal), and this process slows them down for all decisions—that a prohibition is applicable or that it is not applicable. We do not have a ready explanation for the *slower* reaction times in recognizing local, rather than nonlocal, prohibition norms. An ongoing replication with entirely new scene contexts shows equally slow RTs for endorsing local and nonlocal prohibitions, so the present result for prohibitions may be a slight aberration; but this in no way qualifies the strong difference between prohibitions and other norm types.

Rejections (“No” reaction times) of nonlocal norms were faster than (the relatively rare) rejections of local norms, for all three norm types (see Figure 4). However, prohibitions again showed this effect at a much slower speed than prescriptions. Interestingly, people were also slower at rejecting permissions (even though permission endorsements were as fast as prescription endorsements). One possible explanation for this pattern is that denying that an action is permissible is tantamount to saying that it is *impermissible*, which equates to the action being *prohibited*; and because processing prohibitions is cognitively slower, we see slower reaction times for permission rejections.

General Discussion

As a first step in identifying the cognitive properties of norms, we examined the process of norm activation by contexts and the context specificity of norms—with a particular emphasis on how the processes might differ for prescriptions and prohibitions. To investigate these questions we developed a norm generation task, which asks people to actively report the norms they consider applicable to a given context, and a norm recognition task, which asks people to endorse or reject whether candidate norms, offered by the experimenter, are applicable to a given context.

Our main results illustrate the considerable context specificity of norms. In Study 1, the mentioned norms in one

context were often unique to that context (72.9% overall), but this specificity was higher for prescriptions (and permissions) than for prohibitions. In Study 2, people distinguished clearly between norms that were local (previously generated in the given context) vs. nonlocal (previously generated in other contexts), but they discriminated contexts for prescriptions and permissions more clearly than contexts for prohibitions. Of note, the community agreement in generating the most central norms for a given context (see Figure 1) can be compared to the endorsement (“Yes” rates) for these central (local) norms in Study 2. The high endorsements of 98% for permissions and 95% for prescriptions suggest that people agree more strongly on the applicability of local norms than the norm generation task in Study 1 had been able to detect. There, agreement even for the top five permissions and prescriptions ranged between 30% and 75%—though this is perhaps still impressive, given that the task is open-ended rather than a two-alternative forced choice. At the same time, participants in Study 2 were willing to say Yes (the norm applies) to about half of nonlocal norms, which suggests that true context specificity may be lower than the norm generation task had assessed. Even if norms generated in context A are not actively generated in context B, some of those norms local to A may still apply (though more weakly) in context B.

Given that the top-10 prohibitions are not as context-specific as the top-10 prescription norms (according to Study 1), one might suspect that the smaller locality effect in Study 2 may be due to this greater cross-context applicability of prohibition norms identified in Study 1 and used in Study 2. However, in a signal detection analysis, prohibitions showed substantially lower discrimination ($M_d' = 0.50$) than prescriptions did ($M_d' = 1.44$), $t(319) = 14.2$, $p < .001$, and a less “generous” response bias ($M_\beta = 0.71$) than prescriptions ($M_\beta = 0.37$), $t(319) = 11.2$, $p < .001$. Thus, people do not simply say Yes to nonlocal prohibitions because they might be more applicable across contexts but because they have genuine difficulties discriminating between local and nonlocal prohibitions.

The present studies have a number of limitations. For one, even though we selected eight everyday scenes as contexts, their generalizability to a larger number of scenes has yet to be established. The selected scenes were relatively distinct, so overall context specificity may have been lower than, say, across scenes depicting eight different eating establishments. Nonetheless, even formal and informal restaurants or Asian and French restaurants come with different sets of norms.

The specific scenes we selected may also have influenced the overall lower accessibility of prohibitions. Certain contexts—for example, those of high danger or risk—may activate more prohibitions, at a higher level of community agreement, and perhaps at a higher level of context specificity. However, the consistency and strong effect sizes in the present studies do make it plausible that, even across a broader range of everyday contexts, prescriptions (and permissions) will show a greater propensity, agreement, and specificity of context-based activation than prohibitions do.

The variability and vast number of contexts in social life raise the specter of one of the thorniest problems in cognitive science: what a context is, what distinguishes two contexts, or what makes them the same. Somehow, people are able to keep track of a massive number of everyday contexts and usually (though certainly not always) conjure up the appropriate norms for a given context. How they achieve this feat is currently unknown. Perhaps our initial work on norms in contexts will encourage researchers to tackle once more the question of how people deal with the challenge of context identification and context differentiation.

A final limitation of the present studies is that we have no direct evidence yet for the organizational structure of norm representations, aside from a correlation between serial position and degree of agreement of norms. Are norms organized like networks, with properties such as centrality and small-world connectivity (Steyvers & Tenenbaum, 2005)? Are norm representations directly associated with each other in memory or are they activated jointly by co-present features of contexts? And how do people suppress norm contents that have semantic associations with each other but very different context applicability? Tables and chairs are strongly connected in semantic memory, but it would be unfortunate if norm representations blindly followed this association. After all, we are expected to eat at a table and sit on a chair, not the other way round.

Acknowledgments

This project was supported in part by a grant from the Office of Naval Research, No. N00014-14-1-0144 and from and from the Defense Advanced Research Projects Agency (DARPA), SIMPLEX 14-46-FP-097. The opinions expressed here are our own and do not necessarily reflect the views of ONR or DARPA.

References

Andrighetto, G., Castelfranchi, C., Mayor, E., McBreen, J., Lopez-Sanchez, M., & Parsons, S. (2013). (Social) norm dynamics. In G. Andrighetto, G. Governatori, P. Noriega, & L. W. N. van der Torre (Eds.), *Normative multi-agent systems* (Vol. 4, pp. 135–170). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. <https://doi.org/10.4230/DFU.Vol4.12111.135>

Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad is stronger than good. *Review of General Psychology*, 5(4), 323–370. <https://doi.org/10.1037/1089-2680.5.4.323>

Bicchieri, C. (2006). *The grammar of society: The nature and dynamics of social norms*. Cambridge University Press.

Brennan, G., Eriksson, L., Goodin, R. E., & Southwood, N. (2013). *Explaining norms*. Oxford University Press.

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. In M. P. Zanna (Ed.), *Advances in Experimental Social Psychology* (Vol. 24, pp. 201–234). Academic Press.

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>

Cushman, F. (2013). Action, outcome, and value: A dual-system framework for morality. *Personality and Social Psychology Review*, 17(3), 273–292. <https://doi.org/10.1177/1088868313495594>

Gibson, J. J. (1966). *The senses considered as perceptual systems*. Houghton Mifflin.

Goldstein, N. J., & Cialdini, R. B. (2007). Using social norms as a lever of social influence. In A. R. Pratkanis (Ed.), *The science of social influence: Advances and future progress* (pp. 167–191). Psychology Press.

Gould, T. J., & Bevins, R. A. (2012). Context conditioning. In N. M. Seel (Ed.), *Encyclopedia of the Sciences of Learning* (pp. 794–797). Springer US.

Horne, C., & Mollborn, S. (2020). Norms: An integrated framework. *Annual Review of Sociology*, 46(1), 467–487. <https://doi.org/10.1146/annurev-soc-121919-054658>

Janoff-Bulman, R. J., Sheikh, S., & Hepp, S. (2009). Proscriptive versus prescriptive morality: Two faces of moral regulation. *Journal of Personality and Social Psychology*, 96(3), 521–537. <https://doi.org/10.1037/a0013779>

Kumar, A. A. (2021). Semantic memory: A review of methods, models, and current challenges. *Psychonomic Bulletin & Review*, 28(1), 40. <https://doi.org/10.3758/s13423-020-01792-x>

Lindenberg, S. (2013). How cues in the environment affect normative behaviour. In L. Steg, A. E. van den Berg, & J. I. M. de Groot (Eds.), *Environmental psychology: An introduction*. (2012-19546-012; pp. 119–128). BPS Blackwell.

Malle, B. F., & Holbrook, J. (2012). Is there a hierarchy of social inferences? The likelihood and speed of inferring intentionality, mind, and personality. *Journal of Personality and Social Psychology*, 102(4), 661–684. <https://doi.org/10.1037/a0026790>

Malle, B. F., Scheutz, M., & Austerweil, J. L. (2017). Networks of social and moral norms in human and robot agents. In M. I. Aldinhas Ferreira, J. Silva Sequeira, M. O. Tokhi, E. E. Kadar, & G. S. Virk (Eds.), *A World with Robots: International Conference on Robot Ethics: ICRE 2015* (pp. 3–17). Springer.

Rifkin, A. (1985). Evidence for a basic level in event taxonomies. *Memory & Cognition*, 13(6), 538–556. <https://doi.org/10.3758/BF03198325>

Rovee-Collier, C., & Haye, H. (1987). Reactivation of infant memory: Implications for cognitive development. *Advances in Child Development and Behavior*, 20, 185–238. [https://doi.org/10.1016/S0065-2407\(08\)60403-6](https://doi.org/10.1016/S0065-2407(08)60403-6)

Steyvers, M., & Tenenbaum, J. B. (2005). The large-scale structure of semantic networks: Statistical analyses and a model of semantic growth. *La Structure à Grande Échelle Des Réseaux Sémantiques : Des Analyses Statistiques et Un Modèle de Croissance Sémantique*, 29(1), 41–78.

Tversky, B., & Hemenway, K. (1983). Categories of environmental scenes. *Cognitive Psychology*, 15(1), 121–149. [https://doi.org/10.1016/0010-0285\(83\)90006-3](https://doi.org/10.1016/0010-0285(83)90006-3)

Werchan, D. M., Collins, A. G. E., Frank, M. J., & Amso, D. (2015). 8-month-old infants spontaneously learn and generalize hierarchical rules. *Psychological Science*, 26(6), 805–815. <https://doi.org/10.1177/0956797615571442>

Wyman, E., Rakoczy, H., & Tomasello, M. (2009). Normativity and context in young children’s pretend play. *Cognitive Development*, 24(2), 146–155. <https://doi.org/10.1016/j.cogdev.2009.01.003>