Pre-print: not the final version. Draft 2.1 (Sep 2021)

Published in: *Nature Communications* volume 12, Article number: 5776 (2021)

Version of record available at <https://www.nature.com/articles/s41467-021-26067-4>

How social relationships shape moral wrongness judgments

Brian D. Earp, Killian L. McLoughlin, Joshua T. Monrad, Margaret S. Clark*, and Molly J. Crockett*

Yale University, Department of Psychology

* Shared senior authorship. Corresponding author: Molly J. Crockett

Email: molly.crockett@yale.edu

Brian D. Earp: <https://orcid.org/0000-0001-9691-2888>

Killian L. McLoughlin: <https://orcid.org/0000-0002-6591-1986>

Joshua T. Monrad: <https://orcid.org/0000-0002-7377-2074>

Margaret S. Clark: <https://orcid.org/0000-0001-6511-113X>

Molly J. Crockett: <https://orcid.org/0000-0001-8800-410X>

Author Contributions

Designed research: BDE, JTM, MSC, MJC

Performed research: BDE, JTM

Analyzed data: KLM, BDE

Wrote the paper: BDE, JTM, MSC, MJC, KLM

This PDF file includes:

Main Text

Figures 1 to 6 [embedded]

Tables 1 to 2 [embedded]

Supplementary Information

Abstract

Judgments of whether an action is morally wrong depend on who is involved and the nature of their relationship. But how, when, and why do social relationships shape moral judgments? Here we provide evidence to address this question, measuring cooperative expectations and moral judgments in the context of common social relationships such as romantic partners, housemates, and siblings. In a pre-registered study of 423 U.S. participants nationally representative for age, race and gender, we show that people normatively expect different relationships to serve cooperative functions of care, hierarchy, reciprocity, and mating to varying degrees. In a second pre-registered study of 1,320 U.S. participants, we show that these relationship-specific cooperative expectations (i.e., relational norms) make highly precise out-of-sample predictions about the moral wrongness of actions in the context of particular relationships. We further show that our relational norms model better predicts patterns of moral judgments across relationships than alternative models based on genetic relatedness, social closeness, or interdependence. Our findings show how the perceived morality of actions depends not only on the actions themselves, but also on the relational context in which those actions occur.

Main Text

Introduction

Moral psychology has been dominated by studies of judgments and behaviors concerning strangers: individuals who stand in no particular relationship with one another, and who may or may not interact in the future (1). Researchers conducting such studies commonly ask participants to make decisions that impact anonymous others (2–4) or to judge the moral acceptability of hypothetical actions taken by thinly-described agents, as in sacrificial dilemmas where participants must judge the permissibility of killing one person in order to save a greater number (5–8). Of course, people often do encounter strangers as they go about their lives, and the interpersonal standing implied by such encounters can be seen as a bare-bones social relationship involving certain minimal obligations: for example, a “duty of easy rescue” in the case of emergencies (9). The copious research on moral judgments in the context of stranger-stranger relationships thus sheds important light on at least one important aspect of our moral psychology.

However, the vast majority of our moral judgments in everyday life do not concern strangers. Rather, they concern familiar others with whom we stand in particular, often ongoing relationships (10). The stakes of such moral judgments for the maintenance of our personal social networks typically are higher than the stakes of analogous judgments pertaining to strangers. Moreover, moral judgments about interactions between strangers often will differ in systematic ways from judgments about interactions between friends, family members, or other familiar individuals in the same situation (11–13). For example, consider someone who could easily feed a hungry individual but fails to do so. If this person is a mother failing to feed her own child, she likely will be seen as highly blameworthy. But if the person is a local restaurant owner failing to feed a non-paying customer, the same behavior likely will not be seen as blameworthy under ordinary conditions (14).

A number of theorists have highlighted relational context as likely to be important for understanding moral judgment and behavior (10,12,15–18). In line with these developments, there is now a small but growing empirical literature which explores how moral judgments of particular actions vary across different types of social relationships (19–34). How these relationships are theorized depends on the study. For example, one recent study characterized relationships in terms of the *genetic relatedness* of the interaction partners, and showed how varying this factor affects moral judgments about helping behavior (35). Another recent study characterized relationships in terms of the authors’ intuitive sense of the *social closeness* and

relative *interdependence* of the interaction partners – regardless of genetic relatedness – and tested the influence of these factors on judgments about violations of care (36). Researchers have also sought to predict moral wrongness judgments of actions in relational context from a single *cooperative function* thought to characterize a given relationship (e.g., care for a sibling relationship, hierarchy for a teacher-student relationship, and so on) (28).

These studies demonstrate that moral judgments of one and the same action often differ across different types of relationships, depending on how relationship “type” is understood. What is missing, however, is a systematic, data-driven account of the *multiple* cooperative functions that can characterize any given social relationship (14), and an explicit comparison of how well such cooperative functions predict relationally-situated moral judgments relative to alternative models such as genetic relatedness, social closeness, and interdependence. We aim to fill that gap with the present research.

In contrast to genetic relatedness, which can be determined objectively, and the constructs of social closeness and interdependence, both of which have been carefully defined within the relationship science literature, there is no agreed-upon set of cooperative functions prescribed for different social relationships to solve characteristic coordination problems. Recognizing both the theoretical overlap and diversity among the various existing taxonomies of cooperative functions (16,37,38), we build on work by Bugental (39). This work describes a distinctive set of cooperative functions that serve to coordinate behavior in interpersonal relationships. Each function represents an efficient, socially acceptable solution to a particular type of recurrent coordination problem (37), enabling cooperation partners to mutually benefit over repeated interactions (38,40,41). We focus here on four cooperative functions that solve dyadic or two-party coordination problems: care, reciprocity, hierarchy, and mating (**Table 1**).

As has been noted previously, any given relationship may serve multiple cooperative functions, either characteristically or in a specific context [e.g., (39) p. 192; (16) p. 60]. We propose that within a given society, there are prescriptive norms for the set of cooperative functions different relationships should serve (‘relational norms’). In the present work, we sought to (i) describe patterns of relational norms for a large set of common dyadic relationships in a U.S. cultural context; (ii) use these patterns of relational norms to predict out-of-sample judgments of moral wrongness for actions that violate those norms across relationships; and (iii) to compare this ‘relational norms’ model with alternative ways of characterizing dyadic relationships, i.e., in terms of genetic relatedness, social closeness, and interdependence.

Cooperative function	Coordination problem to be solved
<i>Care</i>	Securing basic welfare needs through non-contingent provision (or acceptance) of help or support; maintaining safety; encouraging learning
<i>Reciprocity</i>	Coordinating behavior between individuals with functionally similar (or equal) status, power, authority, or claim on a resource
<i>Hierarchy</i>	Coordinating behavior between individuals with different (unequal) status, power, authority, or claim on a resource
<i>Mating</i>	Finding and maintaining sexual partners; ultimately, producing and ensuring the survival of offspring

Table 1. Cooperative functions of dyadic relationships, adapted from Bugental (39). We note that the care function is based on the work of Clark and colleagues concerning “communal” relationships [e.g. (42)]; it conceptually overlaps with, and replaces, the “attachment” function in Bugental’s model. Because our model is focused on dyadic interactions, we also do not include Bugental’s group-level “coalition” function in this table (see *SI Appendix* for data pertaining to the coalition function).

With respect to aims (i) and (ii), we predicted that relational norms would robustly predict moral judgments about the wrongness of actions in relational context. The basis for this prediction is straightforward: the more a particular set of cooperative functions matters for a given relationship, the morally worse it should be judged to be to neglect or frustrate those same functions within that relationship. Because relationships vary in terms of the set of cooperative functions they are expected to serve, a given action may be judged to be seriously wrong in the context of one relationship but entirely acceptable in the context of another.

We further predicted that our relational norms model would better explain the variance in moral judgments across relationship dyads than genetic relatedness, social closeness, and interdependence, which we believe offer incomplete predictive accounts of moral judgments in relational context. To illustrate, imagine that Person A fails to behave in a deferential manner toward Person B. Insofar as the relationship is normatively expected to be governed by the hierarchy function (see Table 1), with Person A in the subordinate position, such behavior will likely be judged as morally wrong. By contrast, consider how genetic relatedness might explain the wrongness of this action. Some genetically close relationships, such as the parent-child relationship, may, indeed, normatively rely on the hierarchy function to coordinate behavior, and to the extent they do, the action might be judged to be wrong. However, other genetically close relationships, such as siblings of a similar age, are less likely to rely on the hierarchy function, while some genetically distant relationships, such as a typical boss-employee relationship, might be equally or even more likely to rely on the function. Thus, genetic relatedness may ultimately prove to be largely *independent* of the question of what makes certain actions liable to be judged morally wrong.

In summary, unlike most prior work in moral psychology, which has been designed to predict moral judgments from features of actions *regardless* of who performs the action or their relationship to the affected other, here we consider features of common social relationships that we predict will shape moral judgments of actions that occur in the context of *specific* relational dyads. We further predicted that similarity between relationship dyads in terms of their prescribed cooperative functions—or relational norms—would correspond to similarity in moral judgments between relationships. Put another way, dyads with similar relational norms within a given society should be associated with similar patterns of moral judgments across actions, whereas dyads with dissimilar relational norms should be associated with divergent patterns of moral judgments across actions. Finally, we predicted that relational norms would more strongly predict patterns of moral judgments across relationships than alternative predictors, including genetic similarity, social closeness, or interdependence.

Experimental Designs and Results

Relational norms vary across common dyadic relationships

We first measured relationship-specific patterns for prescribed cooperative functions (i.e., relational norms) for a set of common dyadic relationships in the U.S. (study design, sampling plan, and exclusion criteria pre-registered at <https://aspredicted.org/blind.php?x=et9qt8>). Participants (final $n = 423$, U.S. nationally representative for age, race, and gender; “Sample 1”) rated 20 common dyads on the extent to which each is normatively expected to serve the functions of care, reciprocity, hierarchy, mating, and coalition. Results for all 20 relationships across the four functions from **Table 1** are depicted in **Figure 1** (for coalition data, see *SI Appendix*).

As can be seen in **Figure 1**, relational norms varied markedly across dyads in several respects. The reciprocity function was generally prescribed for most dyads (M across dyads = 54.23, $SD = 49.64$; higher than scale midpoint with Bonferroni corrected $\alpha = .0125$, $t(1,422) = 100.47$, $p < .001$, $d = 1.09$; note that all tests reported in the manuscript are two-sided). Meanwhile, the mating function was *negatively* prescribed (i.e., proscribed) for most dyads (M across dyads = -63.02, $SD = 62.01$; lower than the scale midpoint with same correction, $t(1,422) = -93.48$, $p < .001$, $d = -1.02$), with a few obvious exceptions (romantic partners, $M = 95.12$, $SD = 12.94$; friends-with-benefits, $M = 58.43$, $SD = 51.21$). Participants demonstrated higher levels of agreement about whether dyads were expected to serve the mating (SD_{mean} across relationships = 32.26) and care (SD_{mean} across dyads = 37.82)

functions, relative to the reciprocity (SD_{mean} across dyads = 42.25) and hierarchy (SD_{mean} across dyads = 53.72) functions.

Figure 1. Kernel density plots of prescribed cooperative functions for 20 common relationship dyads. Dots represent the population mean prescription for each cooperative function within each relationship, with 95% confidence intervals. The height of the curve represents density: the likely proportions of scores (relative to each function) that fall within the given range along the x-axis.

Figure 2 depicts the four-dimensional relational norm profiles (i.e., sets of prescribed cooperative functions) for a subset of relationships studied, and illustrates several additional features of our data (see *SI Appendix* for functional profile plots for all 20 relationships). First, some relationships are highly functionally “polarized,” showing substantial deviation in mean prescriptions across the four cooperative functions, with at least one function anchored at an extreme end of the scale. An example is the parent and under-18 child relationship (SD across cooperative functions = 85.33, 85.2 for the mother-child and father-child relationship, respectively), which is characterized by a strongly positive expectation for care and a strongly negative expectation for mating. By contrast, other relationships are less functionally polarized, such as the relationship between strangers (SD across functions = 37.93). In these relationships, prescribed cooperative functions are relatively evenly spread across the measured spectrum.

Second, some relationships are functionally “specific,” that is, they are only strongly expected to serve a single cooperative function. For example, the roommate/housemate relationship is strongly expected to serve the reciprocity function ($M = 87.30$, $SD = 21.71$), but less so the care ($M = 24.9$, $SD = 43.64$), hierarchy ($M = -4.48$, $SD = 63.00$), and mating functions ($M = -52.39$, $SD = 49.85$). Similarly, the boss-employee relationship is strongly expected to serve the hierarchy function ($M = 84.75$, $SD = 24.68$), but less so reciprocity ($M = 29.14$, $SD = 58.93$) or care ($M = 7.86$, $SD = 50.21$), and not at all mating ($M = -92.17$, $SD = 23.98$). By contrast, other relationships are functionally “pluralistic,” that is, they are strongly expected to serve multiple cooperative functions. A key example is the romantic partner relationship (M across functions = 64.61), which is strongly expected to serve three of the four cooperative functions: care ($M = 92.43$, $SD = 17.06$), mating ($M = 95.12$, $SD = 12.92$), and reciprocity ($M = 84.95$, $SD = 27.28$) but not, in this sample from the United States, the hierarchy function. See *SI Appendix* for the rankings of all 20 relationship dyads on the dimensions of polarization and specificity.

We also find gender differences in prescribed cooperative functions across relationships. After scaling the raw scores to each participant’s mean rating, we built a mixed linear effects regression model controlling for demographic information (age, race, income, religiosity, and political orientation entered as fixed effects), including participant and relationship dyad type as random effects. With a Bonferroni correction ($\alpha = .0125$ for each of the following effects), the model revealed that women ($M = 0.43$, $SD = 0.75$), compared to men ($M = 0.37$, $SD = 0.79$), reported stronger average expectations that relationships will serve a function of care ($p = 0.012$, 95% CI [0.01, 0.12]), consistent with the existing literature

(43–46). This divergence was most apparent for the roommate/housemate ($M_{female} - M_{male} = 0.17$), customer-seller ($M_f - M_m = 0.15$), teacher-student ($M_f - M_m = 0.14$), neighbor ($M_f - M_m = 0.14$), and colleague/classmate ($M_f - M_m = 0.13$) relationships. Regarding mating, the opposite pattern was found, also consistent with the existing literature (47,48). Men ($M = -1.12$, $SD = 0.97$), compared to women ($M = -1.2$, $SD = 0.89$), reported stronger average expectations that relationships will serve a mating function ($p < .01$, 95% CI [0.02, 0.13]). This divergence was most apparent for the friends-with-benefits ($M_m - M_f = 0.27$), roommate/housemate ($M_m - M_f = 0.25$), acquaintance ($M_{male} - M_{female} = 0.24$), close friend ($M_m - M_f = 0.24$), colleague/classmate ($M_m - M_f = 0.21$), stranger ($M_m - M_f = 0.17$), and neighbor ($M_m - M_f = 0.17$) relationships. Men and women did not significantly differ in their prescriptions for reciprocity and hierarchy functions across relationships. Additional demographic analyses are reported in the [SI Appendix](#).

Figure 2. Relational norm profiles for a subset of 10 relationships. Note: Mother/Father and under-18 child have been combined into a single plot. Plots for all 20 relationships are in the *SI Appendix*.

Common relationships are hierarchically clustered around relational norms

Next, we sought to quantify the distinctiveness of each relationship in four-dimensional relational norm space. Because in many instances patterns of prescribed cooperative functions were not normally distributed in our study population (see **Figure 1**), characterizing relationship differences in terms of their average relational norm scores would sacrifice considerable information. We therefore calculated the Kolmogorov-Smirnov (K-S) distance statistic (a quantification of the difference in overall shape between any two empirical distributions) for each cooperative function for each possible pair of relationships, and averaged across functions to calculate the overall dissimilarity in relational norms for each relationship pair. This approach is conceptually similar to representational similarity analysis (49), but incorporates information about the shapes of the relational norm distributions in addition to distribution means.

We used the relational norm dissimilarity values to conduct a hierarchical clustering analysis using a farthest-point algorithm: $d(u,v)=\max(\text{dist}(u[i],v[j]))$ (50). This revealed four main clusters, depicted in **Figures 3a** and **3b**, which align with intuitive relational categories. The first cluster consists of sexual relationships (romantic partners and friends-with-benefits). The second cluster consists of hierarchical relationships with highly unequal authority between individuals (parents and their minor children, teacher-student, boss-employee). The third cluster includes relationships characterized largely by reciprocal interactions between equals (e.g., customer-seller, roommates/housemates, strangers). And the fourth, final cluster includes familial or other caring relationships (e.g., siblings, extended family members, parents and their adult children).

Based on these analyses, we identified a subset of 10 relationships with relatively distinctive relational norms (see *SI Appendix* for the selection procedure). This subset included long-term romantic partners, friends with benefits, boss and employee, colleagues or classmates, mother/father and under-18 child, siblings, close friends, roommates or housemates, teammates, and strangers. Relational norm profiles for these relationships are depicted in **Figure 2**. We next sought to predict moral judgments of actions performed in the context of these relationships on the basis of their relational norm profiles.

Figure 3a. Circular dendrogram visually representing the mean Kolmogorov-Smirnov (K-S) distance between relationships in four-dimensional relational norm space, clustered hierarchically according to the Voorhees (50) method. Relationships selected for Study 2 are highlighted in a darker shade.

Figure 3b. Radar plots derived from the hierarchical cluster model depicted in Figure 3a. The left panel shows the overlapping clusters; the right panel shows each cluster on its own set of axes.

Relational norm profiles predict relationship-specific moral judgments out of sample

To test the hypothesis that relational norm profiles would predict patterns of moral judgments across common relationships, we first assembled a set of common behaviors that would plausibly weaken or violate one or more of the cooperative functions. Fifteen trained judges rated 86 action statements of the form “Person A does X to Person B” on the extent to which each described action would characteristically weaken (that is, violate or impair) or strengthen each of the cooperative functions, *setting moral questions aside* (that is, they were instructed not to think about whether an action might be right or wrong in any relationship, but only whether it would weaken or strengthen each function). There was very high interrater

agreement in these ratings ($ICC(3, k) = .97$). Using these data, we selected a final set of 12 characteristic function-violating action statements, with 3 statements for each of the 4 dyadic functions (see [SI Appendix](#) for the algorithm used to select the final sub-set). See **Table 2**.

Cooperative function violations		Care	Hierarchy	Mating	Reciprocity
		Mean (SD)	Mean (SD)	Mean (SD)	Mean (SD)
1	Person A sees Person B crying and walks away from them	-87.9 (15.5)	-15.9 (24.3)	-40.1 (35.0)	-19.0 (32.3)
2	Person A keeps checking their cellphone while Person B tells a sad personal story	-75.1 (27.8)	-22.1 (29.6)	-36.5 (39.1)	-35 (35.1)
3	Person A watches passively while Person B carries several heavy boxes up the stairs, even though they could easily help	-73.9 (28.6)	-27.9 (31.3)	-17.4 (29.4)	-29.1 (28.6)
4	Person A refuses to follow a reasonable order from Person B	-22.1 (29.5)	-89.5 (17.7)	-10.7 (13.1)	-16.9 (24.6)
5	Person A repeatedly interrupts Person B while they are speaking	-42.2 (28.5)	-71.6 (27.7)	-15.9 (20.3)	-50.3 (37.0)
6	Person A decides to skip a meeting scheduled with Person B without a good excuse	-42.5 (34.7)	-69.8 (30.8)	-21.3 (31.0)	-38.4 (32.4)
7	Person A refuses to have sex with Person B	-9.3 (17.7)	-10.7 (32.6)	-95.3 (10.7)	0.1 (2.1)
8	Person A repeatedly turns down Person B's offer to go on a romantic date	-10.5 (20.2)	-15.3 (19.4)	-77.7 (23.3)	-9.9 (34.3)
9	Person A invests time and energy in a romantic relationship with someone other than Person B	-26.5 (34.8)	-4.5 (27.6)	-74.0 (28.8)	-31.9 (35.8)
10	Person A decides not to pay Person B back, hoping Person B won't remember	-37.1 (31.3)	-27.0 (33.9)	-12.4 (17.8)	-85.2 (26.7)
11	Person A decides not to return Person B's nice favor	-35.5 (29.9)	-9.2 (24.1)	-22.7 (21.6)	-82.0 (21.8)
12	Person A charges Person B \$50 for an item worth \$25.	-34.9 (27.9)	-28.1 (31.2)	-13.4 (23.6)	-69.4 (30.7)

Table 2. Heatmap showing mean ratings of judges ($n = 15$) of the extent to which each action would characteristically neglect or weaken the care, hierarchy, mating, and reciprocity functions, respectively, where -100 represents the most characteristic function-weakening effect. Note: when rating actions on the “hierarchy” dimension, judges were asked to imagine that Person A was in the *subordinate* role, specifically; when rating actions on the “care” dimension, judges were asked to imagine that Person A was in a *caregiving* (as opposed to care-seeking) role, specifically.

As can be seen in **Table 2**, each action was rated by the judges as having both a main (i.e., “target”) effect on a given function, as well as “side effects” on the other cooperative functions. For example, “Person A sees Person B crying and walks away from them” was rated as most characteristic in weakening the care function ($M = -87.9$, $SD = 15.5$), but also was rated as characteristically weakening the mating function, albeit to a lesser extent ($M = -40.1$, $SD = 35.0$). The fact that one and the same action might simultaneously weaken several interpersonal functions is to be expected, depending on the “logic” of each function and the nature of the action. To account, then, for the specificity of each action as a function-violator, we computed a “target specificity” variable (i.e., main effect minus the mean of side effects) for each action for use in subsequent analyses.

Having identified a set of actions, drawn from everyday life, that were judged to characteristically weaken or violate one or more prescribed cooperative functions, our next step was to assess moral judgments concerning those actions in the context of specific relationships. To do this, we recruited a new group of participants (online US convenience sample, final $n = 1,320$; “Sample 2”), after pre-registering our hypothesis, study design, sampling plan, exclusion criteria, and analysis approach at <https://aspredicted.org/blind.php?x=ta5yj5>. These “naïve” Sample 2 participants were given no information about cooperative functions. Rather, each participant was assigned randomly to consider 1 of the 10 functionally distinctive relationships identified above, and was asked to rate the moral wrongness of all 12 actions listed in **Table 2** in the context of that relationship (e.g., “Imagine that an employee refuses to follow a reasonable order from their boss. How morally wrong would that be, if at all?”). We also asked participants to rate each action on how *likely* it would be to occur in real life, in order to be able to control for violations of nonmoral (i.e., social-conventional) expectations (51) in a pre-registered secondary analysis (see “action likelihood” variable below). For each participant, we computed the mean moral wrongness rating for each of the four cooperative function-violation categories within their assigned dyad. See **Figure 4** for distributions of moral wrongness ratings for each function-violation for each relationship (for demographic analyses, see the *SI Appendix*).

Figure 4. Sample 2 moral wrongness judgments for cooperative function violations in different relationships: kernel density plot of wrongness judgments (0 = not at all morally wrong, 100 = very morally wrong) concerning characteristic function-weakening actions for each of four dyadic cooperative functions across 10 relationships. Dot represents the mean, with 95% confidence intervals. Height of the curve represents density (see Figure 1 for explanation). Note that actions which *weaken* the mating function (e.g., refusing to have sex with someone) were judged closer to “not at all wrong” than “very wrong” for all dyads apart from the romantic partner relationship. Otherwise, the relative lack of visually dramatic differences in the shape of the moral wrongness judgment distributions between relationships can likely be explained by the mild or ‘everyday’ nature of the function-weakening actions employed in this study (see Table 2). Such actions were deliberately chosen to contrast with the more extreme, unusual, or bizarre actions often studied in moral psychology; thus the ability of our model to predict even subtle variance in moral wrongness judgments between relationships for common, non-extreme actions (see analysis below) can be seen as a strength.

We turn now to our main, pre-registered hypothesis. As a first approach, we sought to predict Sample 2 moral wrongness judgments (i.e., for violating each of the four cooperative functions) directly from Sample 1 relational norm profiles in a linear mixed regression model. Sample 2 participants were entered as the highest-level grouping variable, with relationship dyad and function-violation type then entered as crossed random factors. This variance structure accounts for the fact that for each relationship a judgment was made for every function-violation type (i.e., a crossed design). The mean relational norm estimates from Sample 1 were entered alongside both “action likelihood” and “target specificity” as continuous fixed factors for the reasons given above.

The results from this model supported our hypothesis. Relational norms derived from Sample 1 significantly predicted the moral wrongness judgments of Sample 2 participants ($p <$

.001, 95% CI [15.63, 16.88]), accounting for 63% of the variance in mean moral wrongness judgments according to an R^2 analysis (52). Breaking the model down further, we find that target specificity was positively correlated with moral wrongness judgments ($p < .001$, 95% CI [.34, .40]), indicating that the more “on-target” the effect of an action in violating a given function, the more harshly that action was judged. Action likelihood was also negatively correlated with moral wrongness judgments ($p < .001$, 95% CI [-.21, -.18]), indicating that rarer actions were judged more harshly, consistent with past research (53). These results are robust when controlling for demographic factors. For the full regression tables, see *SI Appendix*.

The “action likelihood” variable serves an additional, theoretically important purpose. As we alluded to previously, it can help account for the variance in moral judgments that is due to potentially non-moral violations of social-conventional expectations (i.e., deviations from what is socially expected, whether or not the expectation tracks a perceived moral obligation) (51) as opposed to violations of relational norms specifically. By comparing the R^2 effect size estimates and AIC goodness-of-fit scores (i.e., of relational norm versus action likelihood models) we can judge the relative impact of each metric in explaining moral judgments across relationships. We find that, in a model with no information about relational norms, action likelihood alone does significantly predict moral wrongness judgments in the absence of other predictors ($p < .001$). However, this model explains much less variance, with a poorer goodness-of-fit score (marginal $R^2 = .08$, AIC = 136,496.9) than a model based only on relational norms (marginal $R^2 = .30$, AIC = 130,804). Moreover, the beta value for relational norms (16.26) is more than 80 times larger than that for the action likelihood ratings (-.20) when both are included in the same model (see Table S11f in *SI Appendix*). This shows that relational norms explain moral judgments in this study far better than do merely conventional norms regarding what is socially expected.

Having confirmed that relational norms predict between-relationship variation in moral judgments, over and above mere uncommonness or unexpectedness of behavior, we sought to further explore the nature of this predictive relationship. Specifically, we sought to predict the distance between each pair of relationships in moral judgment space (based on Sample 2 patterns of moral judgment) from their corresponding distances in four-dimensional relational norm space (from Sample 1). To do this, we relied on the same K-S distance approach as described above, comparing the moral judgment distributions for each type of function violation for each possible pair of relationships, and averaging across functions to produce an overall moral judgment dissimilarity score for each relationship pair. We then ran a Spearman’s correlation between these moral judgment dissimilarity values and the previously computed

relational norm dissimilarity values, hypothesizing that the average K-S distance between every pair of relationships in relational norm space would predict the corresponding K-S distance between the same pairs of relationships in moral judgment space. As can be seen in **Figure 5**, this hypothesis was confirmed ($r = .43$, $p = .003$). Looking at the same K-S distances, but on a function-by-function basis (see *SI Appendix* for the corresponding scatterplots), we find that the positive correlation between prescribed cooperative functions and moral judgment K-S distances holds for care ($r = .48$, $p < .001$), mating ($r = .73$, $p < .001$), and hierarchy ($r = .31$, $p = .041$), but not for reciprocity ($r = -.13$, $p = .39$). We will return to this unexpected result for reciprocity in the general discussion.

Figure 5. Scatterplot showing the predicted correlation in K-S distance between each pair of relationship dyads in relational norm space (x-axis) and the K-S distance between those same dyads in moral judgment space (y-axis). Spearman's $r = .43$, $p = .003$. Note that the color of each relationship reflects the cluster in which it is located from Figures 3a and 3b.

Relational norms explain more variance in moral judgments than alternative models

Prior work has sought to predict relational variance in moral judgments from factors such as genetic relatedness (54), social closeness (55), and degree of interdependence (56) of the interaction partners. How do these alternative predictors compare to relational norms in terms of explaining variance in relationally situated moral judgments?

To address this question, we asked a third sample of participants (online US convenience sample; final $n = 85$) to rate the extent to which a well-functioning instance of each of the 10 distinctive relationship dyads would be characterized by social closeness and interdependence. More specifically, for social closeness, we asked how much the partners would “deeply understand each other,” “accept and validate each other’s natures,” and “strive to care for and promote each other’s overall well-being” (taking the mean of these three items). For interdependence, we asked how “frequently,” “strongly,” and in “how many ways” each partner would affect the other’s thoughts, feelings, and behaviors across different situations (ditto). Genetic relatedness for each relationship was determined objectively.

We then entered genetic relatedness, social closeness, and interdependence ratings as predictors in separate linear mixed models similar to those described previously, regressing moral judgments on relational norms. We found no relationship between mean moral wrongness judgments and any of social closeness ($p = .11$, 95% CI $[-.04, .39]$), interdependence ($p = .14$, 95% CI $[-.05, .38]$), or genetic relatedness ($p = .78$, 95% CI $[-12.41, 9.27]$). By contrast, relational norms remained significantly predictive of moral wrongness judgments ($p < .001$, 95% CI $[.10, .15]$), even controlling for the other factors. In addition, measures of model fit suggest that the relational norms model (marginal $R^2 = .69$, AIC = 841) performed substantially better than any of the alternative models: social closeness (marginal $R^2 = .44$, AIC = 908.04), interdependence (marginal $R^2 = .44$, AIC = 908.00), and genetic relatedness (marginal $R^2 = .44$, AIC = 910.33). See *SI Appendix* for the full regression tables.

Discussion

Several scholars have stressed the importance of taking relational context into account in understanding our moral psychology. Yet the way that relational context has been conceptualized so far has suffered from certain limitations. Most commonly, relational context has been understood to vary in a one-dimensional way: for example in terms of the genetic relatedness of the interaction partners (35), or their social closeness or interdependence (36). A more promising approach, we think, is to conceptualize relationships in terms of the

distinctive cooperative functions they are normatively expected to rely upon for coordinating behavior in a given society (14,28,39). Although a number of authors have proposed various taxonomies of cooperative functions (16,37,38) that overlap theoretically with the set employed here (39), it has remained unclear how these functions actually are embedded in different types of relationships. Consequently, we undertook to measure relationship-specific patterns of prescribed cooperative functions (i.e., relational norms) in a U.S. cultural context and to demonstrate how these relational norms predict relationship-specific moral judgments.

To do this, we first measured prescribed cooperative functions for a large set of common dyadic relationships, yielding four-dimensional relational norm profiles for each relationship. Quantifying the distinctiveness of these relationships in terms of their relational norms revealed several distinct clusters of relationship types spanning the domains of care, hierarchy, mating, and reciprocity. Consistent with our predictions, such relational norms predicted out-of-sample moral judgments in relational context, and explained more relational variance in moral judgments than genetic relatedness, social closeness, or interdependence of relationship partners. This suggests that moral judgments of actions within a given relationship are guided by the extent to which the actions thwart or neglect prescribed cooperative functions for that relationship. Moreover, relationships with more similar relational norms showed more similar patterns of moral judgments (Figure 5). These findings reveal a robust underlying structure of expected relational obligations which shape our moral judgments.

Lewin (57) famously argued that behavior is a product of the person and the situation. In a similar spirit, our data confirm that judgments of moral behavior cannot be understood solely with reference to a given act or actor, but rather, must be interpreted in light of the situation, which, in this case, is the type of relationship existing between two individuals. Relationships in a given society can be characterized by distinctive profiles of cooperative norms. They will, therefore, typically be one of the most important situational factors in terms of explanatory power (58). Although relationship theorists have, for decades, worked to characterize the structural elements of various close relationships (59) and have sometimes categorized relationships in terms of cooperative functions necessary for human thriving (39,42), here we systematically described lay perceptions of the ideal functional make-up of a wide range of relationships as identified by ordinary language. Moreover, we were able to use this information to make accurate out-of-sample predictions of moral judgments concerning various actions. We hope that our approach will inspire further research in this vein, both theoretical and empirical, at the interface of relationship science and moral psychology. Ideally, such

research will help to integrate and enrich work in both domains, which has so far remained largely separate.

From a theoretical perspective, one aspect of our current account that requires further attention is the *reciprocity* function. In contrast with the other three functions considered, relationship-specific prescriptions for reciprocity did not significantly predict moral judgments for reciprocity violations. Why might this be so? One possibility is that the model we tested did not distinguish between two different types of reciprocity. In some relationships, such as those between strangers, acquaintances, or individuals doing business with one another (24), each party tracks the specific benefits contributed to, and received from, the other (60). In these relationships, reciprocity thus takes a tit-for-tat form in which benefits are offered and accepted on a highly contingent basis. This type of reciprocity is *transactional*, in that resources are provided, not in response to a real or perceived need on the part of the other, but rather, in response to the past or expected future provision of a similarly valued resource from the cooperation partner. In this, it relies on an explicit accounting of who owes what to whom, and is thus characteristic of so-called “exchange” relationships (42).

In other relationships, by contrast, such as those between friends, family members, or romantic partners – so-called “communal” relationships – reciprocity takes a different form: that of mutually expected responsiveness to one another’s needs. In this form of reciprocity, each party tracks the other’s needs (rather than specific benefits provided) (60) and strives to meet these needs to the best of their respective abilities, in proportion to the degree of responsibility each has assumed for the other’s welfare (42). Future work on moral judgments in relational context should distinguish between these two types of reciprocity: that is, mutual care-based reciprocity in communal relationships (when both partners have similar needs and abilities) and tit-for-tat reciprocity between “transactional” cooperation partners who have equal standing or claim on a resource.

A further limitation of the current studies is that they only concern moral *wrongness* judgments, based on actions that weaken one or more of the expected relational functions. What about judgments of moral rightness, goodness, or praiseworthiness as these relate to actions which *strengthen* one or more of the functions (61)? Will people be judged positively for “merely” meeting functional expectations, as when a parent-child relationship fulfills the care function, or will such judgments be reserved for so-called supererogatory behaviors, going above and beyond the call of duty (62)? Either way, we expect that praiseworthiness judgments for a given action will depend, among other things, on the relational context (functionally understood).

Much of the prior literature in moral psychology has focused on judgments of strangers involved in moral dilemmas that pit distinct ethical principles against one another: for example, a utilitarian imperative to maximize welfare, versus a deontological rule that forbids using individuals as a mere means to an end (5). A key tenet of utilitarianism is that welfare should be maximized *impartially*, rather than prioritizing the well-being of family members (for example) over distant strangers (63). Descriptive research on moral dilemmas shows that many people are not in fact impartial in this sense (64), consistent with our observations here that people have different cooperative expectations for different relationships, leading in turn to different moral judgments depending on relational context. One intriguing possibility is that individuals who more strongly endorse impartial beneficence will have more uniform prescriptions for cooperative functions across relationships, leading to more uniform moral judgments across relational contexts. This perspective also suggests possible antecedents of impartial beneficence. Because care is normative in close relationships (with family, friends, and romantic partners), caring for partners in these relationships does not typically elicit special approbation. Perhaps those who find a sense of purpose or belonging not in tending to close relationships, but in widely being admired (65), tend to “distribute” care across a broader set of relationships (thus showing relatively impartial beneficence).

We note that the generalizability of our findings may be limited in several ways. First, apart from relational role and gender (for mothers and fathers), we did not consider the possible impact of such target characteristics as race, religion, politics, age, and social class, on moral judgments. Each variable itself may impact moral judgments (1,66) and interact with relational context in systematic ways. Second, again apart from gender, we did not comprehensively evaluate how observer (i.e., participant) characteristics along those same demographic lines shape moral judgments, thereby impacting the correspondence between relational norms and moral judgments in relational context. Other individual differences among participants, for example in their relative tendency to engage in different styles of moral reasoning (64) will be important to assess in future research. We see our work as a starting point that may launch further investigations into how both target and observer relational qualities interact with each other and with other kinds of characteristics in shaping moral cognition.

Because we studied participants in the U.S., it also will be important to investigate whether our results generalize across different cultures (67). Although we expect that humans in all cultures form (or stand in) relationships which rely on one or more of the underlying cooperative functions we have highlighted, the patterning of relational norms likely will vary by

culture. Indeed, long-standing programs of research have documented such differences using alternative theoretical frameworks. Hindu research participants from the city of Mysore in southern India, for instance, expected care from a broader array of people -- from parents, friends, and even strangers -- than did research participants from the city of New Haven in the United States (68). The same difference applied to reciprocity (69). In another study, U.S. American wives felt that husbands should care more for them than for their mothers whereas the reverse held true for Egyptian wives (70). Future studies might also compare how “tight” (that is, lacking in variance across situations) relational norms are in each culture (71).

Our primary goal for this research has been simple: to investigate how relational context – in particular, the functional cooperative norms that prescriptively govern dyadic interactions of various kinds – shapes moral judgments. A secondary goal has been to push researchers studying human moral psychology to look at behaviors and associated judgments that are more characteristic of people’s day-to-day lives than heretofore has been the case. Much remains to be done, including more precise and sophisticated analyses of which cooperative functions apply to which relationships, how these functions relate to one another, and how they can be used to predict praiseworthiness judgments (not just judgments of moral wrongness as we have undertaken here). As we and others pursue work that places the study of morality in both geopolitical *and* relational cultural context, we anticipate the emergence of a more nuanced literature on human morality that becomes better integrated with broader and long-standing programs of research on relationships and prosocial behavior.

Materials and methods

All studies were reviewed and approved by the Yale University Institutional review board (protocol #20000022385); informed consent was obtained from participants in each instance prior to data collection. We have posted all study materials, pre-registration forms, raw data, and analysis code on the Open Science Framework (https://osf.io/zxjt6/?view_only=3ff630ea534a408fb9d8259be3db96f8). For complete study descriptions and supplementary findings, see the *SI Appendix*.

Stage 1. For Stage 1, the design, measures, sampling plan, and exclusion criteria were pre-registered at [aspredicted.org](https://aspredicted.org/#26400) (#26400). We used an online polling software (<https://www.nbrii.com/our-process/sample-size-calculator/>) to determine that at least 385 participants would be needed to obtain population estimates of functional expectations with a 5% margin of error and 95% confidence level. Anticipating participant exclusions, we over-sampled by about 15% and aimed to recruit 450 U.S. participants via the Prolific Academic platform (Prolific); 493 ultimately took the survey, each of whom was paid at a rate of \$7.25 per hour. Seventy (70) participants were excluded based on the pre-registered exclusion criteria, leaving us with a final sample of 423 participants (“Sample 1”) who completed an online survey. Participants were given descriptions and definitions of all five cooperative functions adapted from Bugental (57): care, coalition, hierarchy, mating, and reciprocity (see *SI*

Appendix for the full descriptions). To ensure that participants were thinking of the functions in the way we intended, participants were not able to advance to the main part of the study before passing multiple comprehension checks.

We then asked participants to indicate how much each of 20 common relationships should ideally serve each of the five cooperative functions, specifying: “if this kind of relationship was the best possible relationship of its kind it could be [i.e., according to general societal standards], how much should it serve each of those 5 relationship functions?” Participants rated each relationship type in random order. For each relationship type, we included a brief description (see *SI Appendix* for full descriptions). Then, for each combination of relationship and function, participants rated how much the relationship ideally should serve the given function on a sliding scale from -100 (definitely should not serve) to +100 (definitely should serve). Since every participant assessed all five functions for all 20 relationships, we obtained 100 data points per participant. Finally, we collected a battery of demographic measures (described next) as well as exploratory measures for future studies not included here.

In analyzing the demographic information, we first excluded the coalition ratings for the reasons described in the main text. We then used a linear mixed model to regress prescribed cooperative function scores on participant gender (female, male; 4 participants who marked ‘other’ were excluded) for each of the four remaining functions. Race (dichotomized to white and people of color, so we would not have to exclude participants from smaller demographic groups), reported annual income (“low” = \$35K or less, “high” = more than \$35K; split based on U.S. median income), religiosity (“high” versus “low” based on a mean split), and both social and economic political ideology (ditto) were entered into the model as categorical covariates. For each of the mating, care, and reciprocity functions, gender remained a significant predictor, as described in the main text. Other demographic predictors and full model summaries are in the *SI Appendix*.

Stage 2. For Stage 2, the hypothesis, design, measures, sampling plan, exclusion criteria, and analysis approach were pre-registered at [aspredicted.org](https://aspredicted.org/#31592) (#31592). Two main steps were involved: first, selection of a subset of relationships from Stage 1 plus the generation of “action items” to be rated for subsequent use; and second, the actual study, collecting ratings from a new sample (“Sample 2”). As before, the Sample 1 coalition data were excluded.

Using the (remaining) Sample 1 cooperative function scores for all 20 relationships, we performed an analysis that is conceptually similar to a representational similarity analysis (RSA), except that it relies on the Kolmogorov-Smirnov (K-S) distance statistic rather than distribution means. The goal of this analysis was to identify relationships with relatively dissimilar relational norm profiles, so that 10 of the least functionally redundant relationships could be used in the current Stage. For the RSA-like analysis, each relationship was compared to every other relationship on the dimensions of care, hierarchy, mating, and reciprocity. The mean of the four corresponding K-S distance statistics was used for this comparison. Next, we ranked each pair of relationships by its mean K-S distance, from least to most distant (that is, from most functionally redundant to least functionally redundant). We then dropped the relationship from each pair that had the lowest mean K-S distance from all other relationships in the set. Note: for theoretical reasons (i.e., to allow gender comparisons) we decided in advance to retain both the father-child and mother-child relationships in case they faced off. The final set of relationships identified by this procedure is shown in **Figure 2** of the main text.

We then created a set of 86 “action statements” describing behaviors that would plausibly weaken or violate specific cooperative functions based on their underlying logic (i.e., how each function solves its corresponding coordination problem). To determine the extent to which certain actions would characteristically weaken (or strengthen) each of the four dyadic cooperative functions, we had 15 trained judges rate all 86 action items in our set. These

judges were recruited among lab members and colleagues and were given extensive training, either in-person or using an online video conferencing platform, to ensure high quality ratings. They were instructed to consider *only* the functional implications of each action, setting any potential moral considerations strictly aside.

Ratings were obtained via an online survey. The survey included the same formal descriptions of cooperative functions used in Stage 1. Following multiple comprehension checks, the judges were shown the 86 action statements, in random order, in the format “Person A does X to Person B.” For each action and function combination, they made their judgment on a sliding scale ranging from “Would characteristically **weaken** [the function]” (-100) through “It depends / Would neither weaken nor strengthen” (0) to “Would characteristically **strengthen**” (+100).

Next, we created an algorithm to select 12 action items that were rated among the most characteristic in weakening each of the four cooperative functions (three statements per function). First, for each function, the algorithm ranked the actions, in ascending order, by their mean weakening “characteristicness” rating and randomly selected 3 out of the seven most characteristic actions. Second, it computed the mean rating across the three selected actions, yielding one mean score per function. Third, the algorithm computed the standard deviation of the four function-specific means generated in the previous step. Finally, steps one to three were repeated 10,000 times to find the combination of three action statements that yielded the lowest standard deviation of scores across functions. The second iteration of the algorithm was subjected to two further constraints so that we could ensure consistency with potential function-strengthening items planned for testing in future studies (‘strengthen’ set). The first constraint was that the minimum mean score in the ‘weaken’ set could not be lower than the minimum mean score in the ‘strengthen’ set. The second constraint was that the average of the final ‘weaken’ scores could not be more than one point lower than the average of the final ‘strengthen’ scores. So that future studies can be straightforwardly compared with the present study, we selected the ‘weaken’ action items so that they would weaken the cooperative functions to a similar degree as future ‘strengthen’ items would strengthen the cooperative functions. The reason for doing this is because we wanted to make sure that we identified a set of ‘weaken’ items that were not more extreme (in the ‘weaken’ direction) than future ‘strengthen’ items (in the ‘strengthen’ direction). This process resulted in a final set of 12 function-weakening action statements, with three per function, as shown in **Table 2**.

Proceeding to the second main part of Stage 2, a set of naive/lay participants (Sample 2) was recruited, this time on Amazon’s Mechanical Turk (MTurk). To power for the same confidence and margin of error as in Sample 1, but this time with a between-subjects design, it was determined that we would need ratings from 1,551 participants (see *SI Appendix* for the full rationale). Based on the Sample 1 exclusion rate, we over-recruited by about 10% and thus aimed to recruit 1,706 participants; 1,822 ultimately filled out at least part of the survey (not all finished), each of whom was paid \$1.00. Five hundred and two (502) participants were excluded based on pre-registered exclusion criteria, leaving us with a final sample of 1,320 participants. As in Stage 1, they were shown brief descriptions of their assigned relationship. They were told that they would be asked to rate the moral wrongness of various actions within the relationship. To orient them to the rating scale, we clarified that none of the actions they would see would be extreme (e.g., murder), but would rather all be actions that might plausibly occur within the course of day-to-day life.

After passing several attention and comprehension checks, participants were shown, in random order, all 12 action items, tailored to their assigned relationship. For instance, if they were assigned the romantic partner relationship, one of their items was: “Imagine that someone keeps checking their cell phone while their romantic partner tells a sad personal story. How morally wrong would that be, if at all?” Responses were recorded on a sliding scale

from “Not at all morally wrong” (0) to “Very morally wrong” (100). Finally, we collected exploratory data about how likely or unlikely it was that each of the rated actions would happen in real life, and administered the same battery of demographic measures as were used in Stage 1.

For the K-S distance analysis reported in the main text, please note that the functional ratings from Sample 1 were first z-scored to each Sample 1 participant in order to account for individual differences in scale use; for the moral wrongness ratings, no such z-scoring was performed because each Sample 2 participant made only 12 ratings (on account of the between-subjects design).

Stage 3. For Stage 3, we powered to have as many observations per distribution as were obtained in Stage 1. Given design differences between the studies, we determined that we would need ratings from 150 participants for the current study (see *SI Appendix* for the full rationale). This is the number we recruited on MTurk; ultimately, 149 participants completed the survey, each of whom was paid \$1.00 for their time. Sixty-four (64) participants were excluded based on the predetermined exclusion criteria, leaving us with a final sample of 85 participants. Participants were shown the same descriptions of relationships used in Stage 2 and asked to rate them along three dimensions each of social closeness and interdependence (see *SI Appendix* for the precise wording). Responses -- regarding the extent to which a well-functioning instance of each relationship would be characterized by each dimension of both constructs -- were recorded on a sliding scale from 0 to 100, labelled appropriately for each dimension. Similar demographic measures to those used in the previous studies were administered. Please note that, given the unexpectedly large proportion of excluded participants in this study, we performed a sensitivity/robustness analysis with no exclusions (see *SI Appendix* for details), and the results remain substantively the same.

Data availability

All original data (anonymized) is available open access at <https://osf.io/zxjt6/>. Source data for figures can be found in the folder labeled “Source Data - Figures.”

Acknowledgments

Thank you to Billy Brady, Elena Khusainova, Henry Glick, Kevin Anderson, Clara Colombatto, and Hongbo Yu for statistical advice and coding assistance on this project. Thank you to Rachel Calcott, Aden Goolsbee, Nell Mermin-Bunnell, Yuri Munir, Lillian Yuan, Vivan Fung, Vanessa Copeland, Heeral McGhee, Alan Presburger, Samar Allibhoy, Elena DeBre, Isobel Munday, Gargi Singh, Stephanie Brown, Ryan Cox, Brian Bink, Qihe Sun, and Daniel Do, for help in drafting stimuli or serving as judges to rate items used in the study. Finally, thank you to the members of the Crockett, Clark/Bargh, Bloom, and Knobe labs at Yale University for feedback on this research and/or comments on earlier drafts of this manuscript.

References

1. Hester N, Gray K. The moral psychology of raceless genderless strangers. *Perspect Psychol Sci.* 2020;15(2):16–230.
2. Batson CD, Kobrynowicz D, Dinnerstein JL, Kampf HC, Wilson AD. In a very different voice: unmasking moral hypocrisy. *J Pers Soc Psychol.* 1997;72(6):1335–48.
3. Bowles S. Policies designed for self-interested citizens may undermine “the moral sentiments”: evidence from economic experiments. *Science.* 2008;320(5883):1605–9.

4. Crockett MJ, Kurth-Nelson Z, Siegel JZ, Dayan P, Dolan RJ. Harm to others outweighs harm to self in moral decision making. *PNAS*. 2014;111(48):17320–5.
5. Conway P, Goldstein-Greenwood J, Polacek D, Greene JD. Sacrificial utilitarian judgments do reflect concern for the greater good: Clarification via process dissociation and the judgments of philosophers. *Cognition*. 2018;179:241–65.
6. Kahane G, Everett JAC, Earp BD, Farias M, Savulescu J. ‘Utilitarian’ judgments in sacrificial moral dilemmas do not reflect impartial concern for the greater good. *Cognition*. 2015;134:193–209.
7. Mikhail J. Universal moral grammar: theory, evidence and the future. *Trends in Cognitive Sciences*. 2007;11(4):143–52.
8. Greene JD. The secret joke of Kant’s soul. In: Sinnott-Armstrong W, Miller CB, editors. *Moral Psychology: The Neuroscience of Morality: Emotion, Brain Disorders, and Development*. Cambridge, MA: MIT Press; 2008. p. 35–79.
9. Sterri AB, Moen OM. The ethics of emergencies. *Philos Stud*. 2020;online ahead of print.
10. Clark MS, Boothby E, Clark-Polner E, Reis H. Understanding prosocial behavior requires understanding relational context. In: Schroeder DA, Graziano WG, editors. *The Oxford Handbook of Prosocial Behavior*. Oxford University Press; 2015.
11. Clark MS, Boothby E. A strange(r) analysis of morality: a consideration of relational context and the broader literature is needed. *Behavioral and Brain Sciences*. 2013;36(1):85–6.
12. Bloom P. Family, community, trolley problems, and the crisis in moral psychology. *The Yale Review*. 2011;99(2):26–43.
13. Ko A, Pick CM, Kwon JY, Barlev M, Krems JA, Varnum MEW, et al. Family matters: rethinking the psychology of human social motivation. *Perspect Psychol Sci*. 2020;15(1):173–201.
14. Clark MS, Earp BD, Crockett MJ. Who are “we” and why are we cooperating? Insights from social psychology. *Behav Brain Sci*. 2020;43(e66):21–3.
15. Schein C. The importance of context in moral judgments. *Perspect Psychol Sci*. 2020;15(2):207–15.
16. Rai TS, Fiske AP. Moral psychology is relationship regulation: moral motives for unity, hierarchy, equality, and proportionality. *Psychological Review*. 2011;118(1):57–75.
17. Isern-Mas C, Gomila A. Naturalizing Darwall’s second person standpoint. *Integr Psych Behav*. 2020;online ahead of print.
18. Tomasello M. The moral psychology of obligation. *Behav Brain Sci*. 2020;43(e56):1–58.
19. Lee J, Holyoak KJ. “But he’s my brother”: the impact of family obligation on moral judgments and decisions. *Mem Cogn*. 2020;48(1):158–70.
20. McManus RM, Kleiman-Weiner M, Young L. What we owe to family: the impact of special obligations on moral judgment. *Psychological Science*. 2020;in press.
21. Everett JAC, Faber NS, Savulescu J, Crockett MJ. The costs of being consequentialist: social inference from instrumental harm and impartial beneficence. *J Exp Soc Psychol*. 2018;79:200–16.
22. Koleva S, Selterman D, Iyer R, Ditto P, Graham J. The moral compass of insecurity: anxious and avoidant attachment predict moral judgment. *Social Psychological and Personality Science*. 2014;5(2):185–94.
23. Marshall J, Mermin-Bunnell K, Bloom P. Developing judgments about peers’ obligation to intervene. *Cognition*. 2020;201:104215.
24. McGraw AP, Tetlock PE. Taboo trade-offs, relational framing, and the acceptability of exchanges. *Journal of Consumer Psychology*. 2005;15(1):2–15.
25. Selterman D, Koleva S. Moral judgment of close relationship behaviors. *Journal of Social and Personal Relationships*. 2015;32(7):922–45.
26. Selterman D, Moors AC, Koleva S. Moral judgment toward relationship betrayals and those who commit them. *Pers Relationship*. 2018;25(1):65–86.

27. Simpson A, Laham SM. Individual differences in relational construal are associated with variability in moral judgment. *Personality and Individual Differences*. 2015;74:49–54.
28. Simpson A, Laham SM, Fiske AP. Wrongness in different relationships: relational context effects on moral judgment. *J Soc Psychol*. 2016;156(6):594–609.
29. Sunar D, Cesur S, Piyale ZE, Tepe B, Biten AF, Hill CT, et al. People respond with different moral emotions to violations in different relational models: a cross-cultural comparison. *Emotion*. 2020;online ahead of print.
30. Tepe B, Aydınli-Karakulak A. Beyond harmfulness and impurity: moral wrongness as a violation of relational motivations. *Journal of Personality and Social Psychology*. 2019;117(2):310–37.
31. Waytz A, Dungan J, Young L. The whistleblower's dilemma and the fairness–loyalty tradeoff. *J Exp Soc Psychol*. 2013;49(6):1027–33.
32. Weidman AC, Sowden WJ, Berg MK, Kross E. Punish or protect? How close relationships shape responses to moral violations. *Pers Soc Psychol Bull*. 2020;46(5):693–708.
33. Rowe SJ, Vonasch AJ, Turp M-J. Unjustifiably irresponsible: the effects of social roles on attributions of intent. *Social Psychological and Personality Science*. 2020;online ahead of print.
34. Mammen M, Köymen B, Tomasello M. Young children's moral judgments depend on the social relationship between agents. *Cog Dev*. in press;
35. McManus RM, Mason JE, Young L. To whom are we most obligated? Moral values and genetic relatedness structure perceived obligations and their impact on moral judgment. *PsyArXiv*. 2020;pre-print.
36. Gilead M, David YB, Ecker Y. Not our fault: judgments of apathy versus harm toward socially proximal versus distant others. *Soc Psychol Pers Sci*. 2018;9(5):568–75.
37. Curry OS, Mullins DA, Whitehouse H. Is it good to cooperate? Testing the theory of morality-as-cooperation in 60 societies. *Current Anthropology*. 2019;60(1):47–69.
38. Haidt J, Joseph C. The moral mind: How five sets of innate intuitions guide the development of many culture-specific virtues, and perhaps even modules. In: Carruthers P, Laurence S, Stich S, editors. *The Innate Mind*. New York: Oxford University Press; 2007. p. 367–91.
39. Bugental DB. Acquisition of the algorithms of social life: a domain-based approach. *Psychological Bulletin*. 2000;126(2):187–219.
40. Baumard N, André J-B, Sperber D. A mutualistic approach to morality: The evolution of fairness by partner choice. *Behavioral and Brain Sciences*. 2013;36(01):59–78.
41. Sznycer D, Lukaszewski AW. The emotion–valuation constellation: multiple emotions are governed by a common grammar of social valuation. *Evol Hum Behav*. 2019;40(4):395–404.
42. Clark MS, Mills J. The difference between communal and exchange relationships: What it is and is not. *Pers Soc Psychol Bull*. 1993;19(6):684–91.
43. Amanatullah ET, Morris MW, Curhan JR. Negotiators who give too much: unmitigated communion, relational anxieties, and economic costs in distributive and integrative bargaining. *J Pers Soc Psychol*. 2008;95(3):723–38.
44. Helgeson VS. Relation of agency and communion to well-being: evidence and potential explanations. *Psychol Bull*. 1994;116(3):412–28.
45. Jones DC. Friendship satisfaction and gender: an examination of sex differences in contributors to friendship satisfaction. *J Soc Pers Relat*. 1991;8(2):167–85.
46. Le BM, Impett EA. When holding back helps: suppressing negative emotions during sacrifice feels authentic and is beneficial for highly interdependent people. *Psychol Sci*. 2013;24(9):1809–15.
47. Jonason PK, Hatfield E, Boler VM. Who engages in serious and casual sex relationships? An individual differences perspective. *Pers Individ Differ*. 2015;75:205–9.
48. Mark KP, Garcia JR, Fisher HE. Perceived emotional and sexual satisfaction across sexual relationship contexts: gender and sexual orientation differences and similarities. *Can J Hum Sex*. 2015;24(2):120–30.

49. Kriegeskorte N, Mur M, Bandettini PA. Representational similarity analysis - connecting the branches of systems neuroscience. *Front Syst Neurosci.* 2008;2(4):1–28.
50. Voorhees EM. The cluster hypothesis revisited. In: *Proceedings of the 8th annual international ACM SIGIR conference on Research and development in information retrieval.* Montreal, Quebec, Canada: Association for Computing Machinery; 1985. p. 188–96. (SIGIR '85).
51. Turiel E. Thought about actions in social domains: morality, social conventions, and social interactions. *Cognitive Development.* 23(1):136–54.
52. Nakagawa S, Parker TH. Replicating research in ecology and evolution: feasibility, incentives, and the cost-benefit conundrum. *BMC Biology.* 2015 Oct 28;13(88):1–6.
53. Lindström B, Jangard S, Selbing I, Olsson A. The role of a “common is moral” heuristic in the stability and change of moral norms. *Journal of Experimental Psychology: General.* 2018;147(2):228–42.
54. Burnstein E, Crandall C, Kitayama S. Some neo-Darwinian decision rules for altruism: Weighing cues for inclusive fitness as a function of the biological importance of the decision. *J Pers Soc Psychol.* 1994;67(5):773–89.
55. Mills J, Clark MS, Ford TE, Johnson M. Measurement of communal strength. *Pers Relationship.* 2004;11(2):213–30.
56. Berscheid E, Snyder M, Omoto A. The Relationship Closeness Inventory: assessing the closeness of interpersonal relationships. *J Pers Soc Psychol.* 1989;57(5):792–807.
57. Lewin K. *Field Theory in Social Science.* New York: Harper; 1951.
58. Clark MS, Lemay EP, Reis HT. Other people as situations: relational context shapes psychological phenomena. In: Rauthmann JF, Sherman R, Funder DC, editors. *The Oxford Handbook of Psychological Situations.* Oxford University Press; 2018. p. online ahead of print.
59. Kelley HH, Holmes JG, Kerr NL, Reis HT, Rusbult CE, Van Lange PAM. *An Atlas of Interpersonal Situations.* Cambridge: Cambridge University Press; 2003.
60. Clark MS, Mills JR, Corcoran DM. Keeping track of needs and inputs of friends and strangers. *Personality & Social Psychology Bulletin.* 1989;15(4):533–42.
61. Anderson RA, Crockett MJ, Pizarro DA. A theory of moral praise. *Trends Cog Sci.* 2020;online ahead of print.
62. Archer A. Are acts of supererogation always praiseworthy? *Theoria.* 2016;82(3):238–55.
63. Mill JS. *Utilitarianism.* London: Parker, Son and Bourn; 1863. 122 p.
64. Kahane G, Everett JAC, Earp BD, Caviola L, Faber NS, Crockett MJ, et al. Beyond sacrificial harm: a two-dimensional model of utilitarian psychology. *Psychological Review.* 2018;125(2):131–64.
65. Hirsch JL, Clark MS. Multiple paths to belonging that we should study together. *Perspect Psychol Sci.* 2019;14(2):238–55.
66. Jones EF, Parker BL, Joyner MH, Ulku-Steiner B. The influences of behavior valence and actor race on black and white children’s moral and liking judgments. *J Psychol.* 1999;133(2):194–204.
67. Awad E, Dsouza S, Kim R, Schulz J, Henrich J, Shariff A, et al. The moral machine experiment. *Science.* 2018;563(7729):59–64.
68. Miller JG, Bersoff DM, Harwood RL. Perceptions of social responsibilities in India and in the United States: Moral imperatives or personal decisions? *J Pers Soc Psychol.* 1990;58(1):33–47.
69. Miller JG, Bersoff DM. Cultural influences on the moral status of reciprocity and the discounting of endogenous motivation: *Pers Soc Psychol Bull.* 2016;20(5):592–602.
70. Pataki SP, Fathelbab S, Clark MS, Malinowski CH. Communal strength norms in the United States and Egypt. *Interpersona.* 2013;7(1):77–87.
71. Gelfand M, Raver JL, Nishii L, Leslie L, Lun J, Lim BC, et al. Differences between tight and loose cultures: a 33-nation study. *Science.* 2011;332(6033):1100–4.

Supplementary information for:

How Social Relationships Shape Moral Wrongness Judgments

Stage 1

Pre-registration #26400 on aspredicted.org. Full materials, raw data, and code available at https://osf.io/zxjt6/?view_only=66c1211300974dd68e97b88269fec4a3.

Participants.

Using an [online polling software](#), we calculated that 385 participants would be required to have a representative U.S. sample with a 5% margin of error and 95% confidence level. Based on exclusion rates of previous studies conducted in our lab, we over-recruited by about 15% and recruited 450 U.S. participants representative for age, race, and gender via the Prolific Academic platform (Prolific). Ultimately 493 participants took some portion of the survey; not all of them finished. Participants were paid \$0.60 to complete a training session in which they learned about the relational functions of care, coalition, hierarchy, reciprocity, and mating, plus a \$2.00 bonus for completing the rest of the survey (rating each relationship on the extent to which it should ideally serve or not serve each function). To ensure high quality data, we included several attention, comprehension, and bot checks. Seventy (70) participants were excluded based on pre-registered exclusion criteria (see [Table S1](#)). This left us with a final sample of 423 participants (217 female, 201 male, 4 other/nonbinary) ranging in age from 18 to 79 ($M_{\text{age}} = 44.25$, $SD_{\text{age}} = 15.67$); see [Table S2](#) for complete demographic information.

Table S1: Summary of pre-registered exclusion criteria for Stage 1

Exclusion criteria met	Type of check	Excluded <i>N</i>
Did not reach main portion of survey (i.e., did not pass training)	Comprehension check	10
Did not type in the word 'FRIDAY'	Bot check	6
Did not move slider to (at least) 1 of 2 specified positions	Attention check	54

Note: some participants met more than one criterion.

Table S2: Demographics of Sample 1 participants

Age	N (%)	Race	N (%)	Gender	N (%)
18 - 27	82 (19.39%)	White	296 (69.98%)	Female	217 (51.30%)
28 - 37	78 (18.44%)	Black/ African-American	60 (14.18%)	Male	201 (47.52%)
38 - 47	70 (16.55%)	Asian	29 (6.86%)	Other/ Non-binary	4 (0.95%)
48 - 57	75 (17.73%)	Hispanic/Latinx	21 (4.96%)		
58+	117 (27.66%)	Other	10 (2.36%)		
Missing	1 (0.24%)	American Indian/ Alaska Native	4 (0.95%)		
		Hawaiian/ Pacific Islander	2 (0.47%)		

Procedure.

Participants completed a brief online survey through the Qualtrics interface. Before getting to the main part of the survey, participants were shown the full descriptions of each of the five main cooperative functions, as shown in [Table S3](#).

Table S3: Descriptions of relationship functions

Function	Description
Care	<p>Full version: The main purpose of this kind of relationship is to make sure that a person's basic well-being is secure, without any strings attached to the giving or receiving of support (like expecting compensation or favors in return, or feeling a debt). In other words, it is to make sure that people have someone in their corner on whom they can truly count for care and support, in good times and bad.</p> <p>Note that there are two roles assumed by this relationship: the caregiving role (the person that can be truly counted on), and the care-seeking role (the person who may need unconditional support).</p> <p>Brief version: the function of giving or receiving unconditional support</p>

Hierarchy	<p>Full version: The main purpose of this kind of relationship is to help coordinate behavior between people with different status (often they have unequal power or responsibility). In many situations, it is most effective for one person to be the 'leader' or have final say about what happens.</p> <p>So, hierarchical relationships involve assigning people to different roles based on their status or power in a given situation, to help coordinate behavior and accomplish goals.</p> <p>There are two main roles in such relationships: the leader role and the follower role. The person in the leader role has 'final say' over what happens, while the person in the follower role ultimately must go along with what the leader decides.</p> <p>Brief version: the function of coordinating behavior between people of different status, power, or responsibility</p>
Mating	<p>Full version: The main purpose of this kind of relationship is to find and maintain a sexual partner. For our ancestors, the ultimate point of mating was to produce healthy offspring, so that we could pass on our genes and continue as a species.</p> <p>Of course, today we have birth control, and people often have sexual relationships without consciously planning to have children. But the underlying "logic" of the mating relationship -- in terms of the feelings and motivations it tends to inspire -- remains the same: to attract and secure a mate and stay with that person long enough to at least potentially have children together.</p> <p>Brief version: the function of establishing and maintaining a sexual partnership</p>
Reciprocity	<p>Full version: The main purpose of this kind of relationship is to coordinate behavior between people who can mutually benefit, where they each have equal say in a given situation or activity.</p> <p>So, this can be a kind of 'tit-for-tat' arrangement between people, where each one says, "I'll scratch your back if you scratch mine." It can also be a way to keep things fair between people exchanging favors, goods, or services. It can even be a way to coordinate activities for mutual enjoyment, such as playing a game together. In any case, it involves making sure that the scales between people don't get too far out of balance.</p> <p>Brief version: the function of coordinating behavior between people with equal say in a situation and keeping things fair</p>
Coalition	<p>Full version: The main purpose of this kind of relationship is to form and maintain a group identity, so group members can work toward a common goal. People in the same group look out for each other and try to promote their own group's interests over the interests of competing groups: "us versus them."</p> <p>This involves having shared expectations for what's normal or appropriate behavior and potentially making sacrifices for the good of the group, especially when it's in danger.</p> <p>Brief version: the function of forming and maintaining a group identity for a common goal: us versus them</p>

To ensure that participants were paying attention and were thinking of the functions in the way we wanted them to, each description was followed by a multiple-choice question about the definitions of the functions. Participants were not allowed to advance to the main part of the survey if they failed to answer this check correctly.

Once we had introduced the participants to the five cooperative functions, we gave them instructions for the main task of the survey. For each of 20 relationships, we asked participants how much the relationship ideally should serve each of the five cooperative functions. We specified that, by ‘ideally,’ we meant that “if this kind of relationship was the best possible relationship of its kind it could be,” how much should it serve each of those five functions?

Participants were then presented 20 blocks of questions, one for each relationship, in random order. For each relationship, we included a specific description of what we meant by that relationship – see [Table S4](#) for the descriptions. Then, for each combination of relationship and cooperative function, participants rated how much the relationship ideally should serve the given function, with the function now presented in its brief form as a reminder (see Table S3). For instance, if the relationship and cooperative function pair was siblings/care, participants would be asked: “To what extent should the relationship between siblings ideally serve the function of **giving or receiving unconditional support?** (Care function.)” Responses were recorded on a sliding scale ranging from ‘Definitely SHOULD NOT’ (-100) through ‘Neutral’ (0) to ‘Definitely SHOULD’ (+100). Each of the recruited participants responded to questions about all five functions for all 20 relationships, yielding 100 data points per participant. Finally, we collected a battery of demographic measures: gender, age, race, ethnicity, income, level of education, English fluency, political leanings on social and economic issues, and religiosity.

Table S4: Descriptions of relationships

Relationship	Description
Siblings	This refers to brothers and sisters. It includes adoptive as well as biological brothers and sisters.
Long-term romantic partners	This refers to romantic partners that have a commitment to each other, meaning they intend to remain together for the long term. It includes married partners like spouses, but also long-term romantic partners who aren’t married but still have a commitment.
Close friends	This refers to people who seek one another out to spend time together, and who refer to each other as best friends, close friends, or good friends. It includes two people of any gender combination who are committed to remaining friends for the long term, and whose interest in each other is NOT primarily romantic.
Work colleagues or classmates	This refers to people who interact with each other on a regular basis at work/school or in work/school related activities. It includes only people who are on the same career or schooling level as each other.
Boss and employee	This refers to any workplace relationship in which one person directly supervises the other and that same person has some decision-making control over the other’s activities and outcomes. It includes only people who interact with each other in person in a workplace.

Teacher and student	This refers to relationships in which someone with more experience guides someone with less experience, such as in school, sports or other extracurricular activity. It includes teacher-student or coach-player relationships, for example.
Doctor and patient	This refers to relationships in which one person provides expert medical or therapeutic services to the other person. It includes medical doctor-patient relationships or therapist-client relationships, for example.
Extended family members	This refers to relationships with relatives who are not in the immediate family (so, not parents or siblings). It includes relationships with a cousin, aunt, uncle, grandfather or grandmother.
Teammates	This refers to a relationship between two people who are in a clearly defined group together working towards a common goal. It includes teammates on a sports team or members of a theater troupe, for example.
Customer and seller	This refers to any relationship in which someone sells something and another person buys that thing, and they interact with each other directly. It includes a local baker and the baker's customer, or a house cleaner and the person who owns the house, for example.
Neighbors	This refers to people who live in the same section of an apartment building, or within the same block of houses on a street. It includes only people who know each other and at least occasionally interact.
Friends with benefits	This refers to people who know each other and interact sexually on a somewhat regular basis, but who refer to each other as friends rather than romantic partners. It includes friends who interact sexually without committing to a monogamous romantic relationship.
Acquaintances	This refers to people who know each other, and interact now and then, but don't consider one another to be friends. It includes people one might know from work, school, or the neighborhood.
Members of a political party	This refers to people who are active, registered members of the same political party. They likely hold similar ideological positions and tend to vote for the same candidates in elections. They may not necessarily know each other for any other reason than their common political activities and party affiliation.
Roommates/housemates	This refers to anyone who lives together who are not family members or in a romantic relationship. This includes college room- or suite mates, apartment mates, or housemates.
Father and child (under 18)	This refers to fathers and their non-adult (meaning under age 18) children specifically. It includes both biological and adoptive fathers/children.
Mother and child (under 18)	This refers to mothers and their non-adult (meaning under age 18) children specifically. It includes both biological and adoptive parents/children.
Father and child (over 18)	This refers to fathers and their adult (meaning over age 18) children specifically. It includes both biological and adoptive fathers/children.
Mother and child (over 18)	This refers to mothers and their adult (meaning over age 18) children specifically. It includes both biological and adoptive mothers/children.
Strangers	This refers to people who encounter each other in any setting for the first time. It includes people who don't know each other from another context and don't anticipate interacting again in the future.

Data preparation and analysis details

Raw data files (.csv) were prepared and analyzed using Python, within a Jupyter Notebook environment. Primary packages used: numpy, scipy, statsmodels, matplotlib, seaborn, pandas. For data files and all coding scripts, see the OSF link above.

Supplementary results

Table S5. Most to least functionally *polarized* relationships.

	Functional Expectations				Across all functions
	Care	Hier.	Mate.	Recip.	
	<i>M</i>	<i>M</i>	<i>M</i>	<i>M</i>	<i>SD</i>
Siblings	84.06	-0.83	-98.47	74.96	74.86
Mother/child (under 18)	95.53	65.39	-98.22	9.05	74.10
Father/child (under 18)	93.99	67.39	-98.37	12.93	73.99
Father/child (over 18)	86.58	35.97	-97.59	53.62	70.32
Mother/child (over 18)	88.05	28.02	-97.69	49.75	69.79
Boss and employee	7.86	84.75	-92.17	29.14	65.45
Extended family members	64.65	15.22	-96.40	57.26	65.00
Teacher and student	42.97	72.77	-95.61	24.49	64.30
Teammates	50.53	31.86	-73.00	75.43	63.72
Doctor and patient	53.75	41.63	-95.31	30.40	61.21
Close friends	79.39	-20.32	-50.31	79.96	59.18
Members of a political party	7.60	45.24	-66.04	63.82	55.39
Customer and seller	-18.37	22.04	-81.77	60.39	53.16
Colleagues or classmates	17.38	18.92	-60.00	77.29	50.93
Roommates or housemates	24.90	-4.48	-52.39	87.30	50.71
Neighbors	13.08	-16.25	-57.99	67.57	46.35
Romantic partners	92.43	-14.07	95.12	84.95	45.80
Friends with benefits	28.13	-30.36	58.43	59.87	39.87
Acquaintances	-2.74	-3.54	-46.82	51.38	34.85
Strangers	-26.62	-11.00	-55.87	34.95	33.46

Table S6. Most to least functionally *specific* relationships.

	Functional Expectations							
	Care	Hier.	Mate.	Recip.	Highest <i>M</i> Function	Highest <i>M</i> Value	Other Sum	Max Other Difference
	<i>M</i>	<i>M</i>	<i>M</i>	<i>M</i>		<i>M</i>		
Boss and employee	7.86	84.75	-92.17	29.14	Hierarchy	84.75	-55.17	139.92
Customer and seller	-18.37	22.04	-81.77	60.39	Reciprocity	60.39	-78.09	138.49

Neighbors	13.08	-16.25	-57.99	67.57	Reciprocity	67.57	-61.15	128.72
Strangers	-26.62	-11.00	-55.87	34.95	Reciprocity	34.95	-93.49	128.44
Mother/child (under 18)	95.53	65.39	-98.22	9.05	Care	95.53	-23.78	119.31
Roommates or housemates	24.90	-4.48	-52.39	87.30	Reciprocity	87.30	-31.96	119.25
Father/child (under 18)	93.99	67.39	-98.37	12.93	Care	93.99	-18.05	112.04
Siblings	84.06	-0.83	-98.47	74.96	Care	84.06	-24.34	108.40
Mother/child (over 18)	88.05	28.02	-97.69	49.75	Care	88.05	-19.91	107.96
Acquaintances	-2.74	-3.54	-46.82	51.38	Reciprocity	51.38	-53.09	104.47
Colleagues or classmates	17.38	18.92	-60.00	77.29	Reciprocity	77.29	-23.70	100.99
Teacher and student	42.97	72.77	-95.61	24.49	Hierarchy	72.77	-28.14	100.91
Father/child (over 18)	86.58	35.97	-97.59	53.62	Care	86.58	-8.00	94.59
Extended family members	64.65	15.22	-96.40	57.26	Care	64.65	-23.92	88.57
Doctor and patient	53.75	41.63	-95.31	30.40	Care	53.75	-23.29	77.04
Members of a political party	7.60	45.24	-66.04	63.82	Reciprocity	63.82	-13.20	77.02
Close friends	79.39	-20.32	-50.31	79.96	Reciprocity	79.96	8.77	71.19
Teammates	50.53	31.86	-73.00	75.43	Reciprocity	75.43	9.38	66.05
Friends with benefits	28.13	-30.36	58.43	59.87	Reciprocity	59.87	56.20	3.68
Romantic partners	92.43	-14.07	95.12	84.95	Mating	95.12	163.31	-68.18

Figure S1. Coalition ratings. Kernel density plot of functional expectations for coalition only for 20 common relationships. Dot represents the mean, with 95% confidence intervals. The height of the curve represents density: the likely proportions of scores (relative to each function) that fall within the given range along the x-axis.

Figure S2. “Functional profile” violin plots for all 20 relationships.

Tables S7a-S7e. Complete demographic analyses for Stage 1/Sample 1: regression tables. Note that participants who reported gender as ‘other’ were excluded from the analysis as this was not a big enough sub-group to make statistical comparisons possible. Table a = overall model; b = gender and other demographic effects on functional expectations for mating; c = same for care; d = same for reciprocity; e = same for hierarchy.

a. Overall

Predictor	β	Std. Error	p
(Intercept)	-.02	.01	.08
Gender (Female)	< -.001	.01	.95
Ethnicity (Non-White)	.01	.01	.52
Income (High)	-.003	.01	.79
Economic Ideology (Liberal)	-.01	.02	.44
Social Ideology (Liberal)	.003	.02	.83
Religiosity (Very Religious)	-.01	.01	.44

Note. These are the results of fitting a fixed effects only model. A mixed-effects model failed to significantly improve upon these results.

b. Mating

Predictor	β	Std. Error	95% CI [LL, UL]	p	Fit
(Intercept)	-1.09	.03	[-1.15, -1.02]	< .001	
Gender (Female)	-.07	.03	[-.12, -.02]	< .01	
Ethnicity (Non-White)	-.02	.03	[-.07, .04]	.49	
Income	-.02	.03	[-.07, .03]	.37	

(High)

Economic Ideology (Liberal)	-.05	.04	[-.12, .02]	.16
Social Ideology (Liberal)	0.06	.04	[-.01, .13]	.12
Religiosity (Very Religious)	-.05	.03	[-.10, .00]	.07

Conditional $R^2 = .88$

Note. *LL* and *UL* indicate the lower and upper limits of a 95% confidence interval, respectively.

c. Care

Predictor	β	Std. Error	95% CI [LL, UL]	<i>p</i>	Fit
(Intercept)	.37	.03	[.31, .42]	< .001	
Gender (Female)	.06	.02	[.01, .11]	.012	
Ethnicity (Non-White)	.004	.02	[-.04, .05]	.87	
Income (High)	-.003	.02	[-.05, .04]	.91	
Economic Ideology (Liberal)	-.01	.03	[-.08, .05]	.67	
Social Ideology (Liberal)	-.02	.03	[-.08, .04]	.51	
Religiosity (Very Religious)	.05	.02	[-.001, .10]	.06	
					Conditional $R^2 = .88$

d. Reciprocity

Predictor	β	Std. Error	95% CI [LL, UL]	<i>p</i>	Fit
(Intercept)	.53	.02	[.48, .58]	< .001	
Gender	.03	.02	[-.01, .07]	.12	

(Female)				
Ethnicity (Non-White)	< .001	.02	[-.04, .04]	1
Income (High)	-.02	.02	[-.06 .01]	.28
Economic Ideology (Liberal)	.03	.03	[-.02, .08]	.27
Social Ideology (Liberal)	-.01	.03	[-.07, .04]	.59
Religiosity (Very Religious)	< .001	.02	[-.04, .04]	.98
Conditional R ² = .88				

e. Hierarchy

Predictor	β	Std. Error	95% CI [LL, UL]	<i>p</i>	Fit
(Intercept)	.09	.04	[.48, .58]	.014	
Gender (Female)	-.02	.03	[-.01, .07]	.48	
Ethnicity (Non-White)	.05	.03	[-.04, .04]	.14	
Income (High)	.04	.03	[-.06 .01]	.23	
Economic Ideology (Liberal)	-.02	.04	[-.02, .08]	.69	
Social Ideology (Liberal)	-.01	.04	[-.07, .04]	.89	
Religiosity (Very Religious)	.04	.03	[-.04, .04]	.26	
Conditional R ² = .88					

Stage 2

Pre-registration #31592 on aspredicted.org. Full materials, raw data, and code available at https://osf.io/zxjt6/?view_only=66c1211300974dd68e97b88269fec4a3.

Materials creation

Selection of relationships

In Stage 1, we collected data for 20 different relationships. Our analysis of the Kolmogorov-Smirnov (K-S) distance between relationships (see the main text) revealed that several of these relationships were highly similar in terms of their prescribed cooperative functions (i.e., relational norm profiles). To avoid redundancy in Stage 2, we decided to use a subset of 10 relationships for the following study that were among the most functionally dissimilar to each other. Our procedure for selecting these relationships is described next.

As noted in the main manuscript, we first excluded the data pertaining to the coalition function. We then calculated the mean K-S distance score for every pairing of relationships across the four remaining functions. We set up 10 slots for "face-offs" between pairs of relationships with low K-S distances. For each face-off, we dropped the relationship with the lowest mean K-S distance from all other relationships (i.e., the more redundant of the two relationships considered in the context of the entire set). To enable gender comparisons, however, we first instituted a rule such that, if a father/mother relationship faced off, each was retained. Then, we moved sequentially from the lowest K-S score pairings to the highest K-S score pairings until all 10 slots were filled, dropping relationships along the way according to the first rule (if a relationship faced off with another relationship that had already been eliminated, it was retained by default). The final set of relationships identified by this method were: long-term romantic partners, friends with benefits, boss and employee, colleagues or classmates, mother/father and under-18 child, siblings, close friends, roommates or housemates, teammates, and strangers.

Selecting action statements

In this part of the study, we sought to test the hypothesis that the relational norm profile of a given relationship (based on the prescribed cooperative functions identified in Stage 1) would predict moral judgments for violations of associated functions in the context of that relationship. To this end, we created a set of 86 actions that we thought would plausibly weaken or impair one or more cooperative

functions. To determine the extent to which each action would characteristically weaken (or strengthen) each of the four dyadic cooperative functions, we had 15 trained judges rate each of the 86 actions in our set. These judges were recruited among lab members and colleagues and were given extensive training either in person or over Skype to ensure that their ratings reflected only the functional implications of each action (i.e., according to the cooperative logic of the functions) rather than being about moral judgments of any kind.

The trained judges completed an online survey, which included the same descriptions of cooperative functions that we used in Stage 1. After reading these descriptions and completing multiple comprehension checks, the judges were shown the 86 action statements, all of which were of the form “Person A does X to Person B” (see original materials at the OSF link above for the full list of actions). For each action, the judges were asked about the extent to which the action would weaken or strengthen each of the five functions. Their responses were recorded on a sliding scale, ranging from “Would characteristically **weaken**” (-100) through “It depends / Would neither weaken nor strengthen” (0) to “Would characteristically **strengthen**” (100). Thus, we obtained a mean rating between -100 and 100 for each action-function pair.

Based on the judges’ ratings, we used an algorithm (described in the main text)¹ to identify 12 actions (three for each function) that were most characteristic as individual function-weakensers, while maintaining roughly equal “characteristicness” of items across functions.² This process resulted in a final set of 12 function-weakening action statements, depicted in [Table S8](#).

Table S8: Action statements used in Stage 2

Function	Action	Mean rating
Care	<ul style="list-style-type: none"> - Person A sees Person B crying and walks away from them - Person A keeps checking their cellphone while Person B tells a sad personal story - Person A watches passively while Person B carries several heavy boxes up the stairs, even though they could easily help 	-79.0
Hierarchy	<ul style="list-style-type: none"> - Person A refuses to follow a reasonable order from Person B - Person A repeatedly interrupts Person B while they are speaking - Person A decides to skip a meeting scheduled with Person B without a good excuse 	-77.0

¹ MATLAB code is available at <https://osf.io/j435q/>.

² Three items from the 86-action set were excluded before running the algorithm due to having misleading or ambiguous wording. After running the code, we also dropped the item “Person A mocks Person B for a poor performance, when Person B tried their best,” as we realized that the word “performance” could have different meanings across relationships (e.g., job performance versus an artistic performance).

Mating	<ul style="list-style-type: none"> - Person A refuses to have sex with Person B - Person A repeatedly turns down Person B's offer to go on a romantic date - Person A invests time and energy in a romantic relationship with someone other than Person B 	-82.3
Reciprocity	<ul style="list-style-type: none"> - Person A decides not to pay Person B back, hoping Person B won't remember - Person A decides not to return Person B's nice favor - Person A charges Person B \$50 for an item worth \$25. 	-78.9

Participants.

In Stage 1, we powered for 95% confidence in a 5% margin of error for a nationally representative sample across age, race, and gender. This required that we have 385 observations per distribution (with each participant rating all 20 original relationships on all 5 original functions). To ensure that our Sample 2 distributions would be comparable to those from Sample 1, we powered for the same confidence in the same margin of error. In the first sample, every participant gave one rating per function for all of 20 relationships; in this study, Sample 2 participants would give three ratings for just one relationship out of a smaller set of 13 relationships (13 because questions regarding the non-symmetrical relationships, e.g., boss-employee, were asked in both directions). To achieve parity, then, we multiplied the previous target sample of 385 (per distribution) by 13 (accounting for the switch to a between-subjects design) and divided by three (accounting for three ratings per function in the current sample, compared to just one in Sample 1), yielding a required sample of 1,551. As in Stage 1, we over-recruited by about 10%, aiming for 1,706 participants for Stage 2. Ultimately, 1,822 participants took at least part of the survey (not all finished).

Once again, the participants were recruited online³ and were paid were paid \$1.00 to complete the survey. Five hundred and two (502) participants were excluded based on pre-registered exclusion criteria (see [Table S9](#) for criteria), leaving us with a final sample of 1,320 participants (553 female, 758 male, 6 other/non-binary) ranging in age from 18 to 73 ($M_{\text{age}} = 35.32$, $SD_{\text{age}} = 10.56$). See [Table S10](#) for complete demographic information.

Table S9: Summary of exclusion criteria for Stage 2

Exclusion criteria met	Type of check	Excluded <i>N</i>
Failed question about survey instructions OR answered multiple-choice question	Comprehension/attention check	246

³ Because Prolific did not allow for samples larger than 1,500, we instead relied on Amazon's Mechanical Turk platform (MTurk).

incorrectly (correct answer includes the word “embarrassment”)

Did not move slider to (at least) 1 of 2 specified position	Attention check	171
Failed CAPTCHA test	Bot check	0
Failed text-entry test	Bot check	53
Being younger than 18	Demographic check	2
Not fluent English speaker	Demographic check	12
Finished survey in < 4 min.	Quality check	73

Note: some participants met more than one criterion.

Table S10: Demographics of Sample 2 participants

Age	N (%)	Race	N (%)	Gender	N (%)
18 - 27	329 (24.96%)	White	929 (70.49%)	Female	553 (41.96%)
28-37	561 (42.56%)	Black/ African-American	175 (13.28%)	Male	758 (57.51%)
38-47	230 (17.45%)	Asian	92 (6.98%)	Other/ Non-binary	6 (0.46%)
48-57	135 (10.24%)	Hispanic/ Latinx	88 (6.68%)		
58+	63 (4.78%)	Other	19 (1.44%)		
Missing	0 (0.00%)	American Indian/ Alaska Native	12 (0.91%)		
		Hawaiian/ Pacific Islander	1 (0.08%)		

Procedure.

Each participant was assigned to one of 13 relationship-pairs (13 because questions about the non-symmetrical relationships, e.g., boss-employee, were asked separately in both directions), and were shown a brief description of their assigned relationship (see [Table S4](#)). We informed participants that they would be asked to consider various actions in the context of their assigned relationship and to answer how morally wrong each of those actions would be. To orient them to the rating scale, we clarified that none of the actions they would see would be extreme (e.g., murder), but rather would all be actions that might plausibly occur within the course of day-to-day life. We then ‘anchored’ their expectations by showing them a list of actions comparable to the ones included in the task.

Following instructions and attention checks, participants were shown the 12 (three for each of the four functions) hypothetical actions selected by our algorithmic approach described above. For instance, for the sibling relationship, participants were asked, e.g., “Imagine that someone keeps checking their cellphone while their sibling tells a sad personal story. How morally wrong would that be, if at all?” Responses were recorded on a sliding scale from “Not at all morally wrong” (0) to “Very morally wrong” (100). Finally, we collected exploratory data about how likely it is that each action would occur in real life, plus the same demographic measures that we collected from Sample 1 participants.

Data preparation and analysis details

Raw data files (.csv) were prepared and analyzed using Python, within a Jupyter Notebook environment. Primary packages used: numpy, scipy, statsmodels, matplotlib, seaborn, pandas. For data files and all coding scripts, see the OSF link above.

Supplementary results

Tables S11a-S11d. Complete demographic analyses for Stage 2/Sample 2: regression tables. Note that participants who reported gender as ‘other’ were excluded from the analysis as this was not a big enough sub-group to make statistical comparisons possible. Table a = overall model; b = gender and other demographic effects moral wrongness judgments for mating; c = same for care; d = same for reciprocity; e = same for hierarchy.

a. Overall

Predictor	β	Std. Error	95% CI [LL, UL]	<i>p</i>	Fit
(Intercept)	41.61	1.00	[39.64, .43.57]	< .001	

Gender (Female)	-.51	.85	[-2.18, 1.16]	.55
Ethnicity (Non-White)	4.39	.93	[2.56, .6.22]	< .001
Income (High)	.99	.85	[-.68, 2.67]	.24
Economic Ideology (Liberal)	1.76	.10	[-.39, 3.91]	.11
Social Ideology (Liberal)	-1.31	.10	[-3.48, .85]	.24
Religiosity (Very Religious)	8.28	.87	[6.58, 9.98]	< .001

Conditional $R^2 = .14$ **b. Mating**

Predictor	β	Std. Error	95% CI [LL, UL]	<i>p</i>	Fit
(Intercept)	5.83	1.61	[2.68, .8.98]	< .001	
Gender (Female)	-2.66	1.37	[-5.33, 0.02]	0.052	
Ethnicity (Non-White)	5.48	1.50	[2.55, 8.42]	< .001	
Income (High)	4.79	1.37	[2.11, 7.47]	< .001	
Economic Ideology (Liberal)	1.71	1.76	[-1.73, 5.16]	.33	
Social Ideology (Liberal)	1.29	1.77	[-2.18, 4.75]	.47	
Religiosity (Very Religious)	1.81	1.39	[12.09, 17.53]	< .001	

Conditional $R^2 = .61$ **c. Care**

Predictor	β	Std. Error	95% CI [LL, UL]	<i>p</i>	Fit
(Intercept)	55.72	1.60	[52.6, 58.86]	< .001	
Gender (Female)	2.60	1.34	[-.03, 5.24]	.053	
Ethnicity (Non-White)	3.77	1.48	[.86, 6.67]	.01	
Income (High)	-1.04	1.35	[-3.68, 1.61]	.44	
Economic Ideology (Liberal)	2.66	1.79	[-.86, 6.17]	.14	
Social Ideology (Liberal)	-3.00	1.80	[-6.52, .53]	.1	
Religiosity (Very Religious)	6.45	1.37	[3.77, 9.13]	< .001	
					Conditional $R^2 = .49$

d. Reciprocity

Predictor	β	Std. Error	95% CI [LL, UL]	<i>p</i>	Fit
(Intercept)	62.92	1.12	[60.73, 65.12]	< .001	
Gender (Female)	.15	.95	[-1.72, 2.02]	.88	
Ethnicity (Non-White)	2.48	1.04	[.43, 4.52]	.02	
Income (High)	-.67	.95	[-2.54, 1.20]	.48	
Economic Ideology (Liberal)	2.82	1.22	[.42, 5.23]	.02	
Social Ideology (Liberal)	-2.40	1.23	[-4.82, .01]	.052	
Religiosity	3.34	.97	[1.45, 5.24]	< .001	

(Very Religious)

Conditional $R^2 = .07$ **e. Hierarchy**

Predictor	β	Std. Error	95% CI [LL, UL]	p	Fit
(Intercept)	47.76	1.60	[44.62, 50.89]	< .001	
Gender (Female)	-2.58	1.35	[-5.23, .07]	.06	
Ethnicity (Non-White)	5.93	1.50	[2.98, 8.87]	< .001	
Income (High)	.34	1.36	[-2.32, 3.01]	.80	
Economic Ideology (Liberal)	1.38	1.70	[-1.95, 4.72]	.42	
Social Ideology (Liberal)	-3.34	1.72	[-6.70, .03]	.052	
Religiosity (Very Religious)	7.05	1.38	[4.34, 9.75]	< .001	

Conditional $R^2 = .46$ **S11f.** Full regression table for main analysis, controlling for demographic information.

Predictor	β	Std. Error	95% CI [LL, UL]	p	Fit
(Intercept)	76.03	1.31	[73.45, 78.60]	< .001	
Relational Norms	16.26	0.32	[15.64, 16.88]	< .001	
Action Likelihood	-.20	.01	[-.21, -.18]	< .001	
Target Specificity	0.37	.01	[.34, .40]	< .001	
Gender (Female)	-.81	.81	[-2.39, .77]	.31	
Ethnicity (Non-White)	4.96	.88	[3.23, 6.69]	< .001	

Income (High)	.54	.81	[-1.04, 2.13]	.50
Economic Ideology (Liberal)	2.40	1.04	[.37, 4.43]	.02
Social Ideology (Liberal)	-2.09	1.04	[-4.13, -.04]	.05
Religiosity (Very Religious)	8.47	.82	[6.86, 10.07]	< .001

Conditional R² = .63

Figures S3a-S3d. Scatterplots of function-specific correlations between KS distance scores in relational norm and moral judgment space respectively: a = care, b = mating, c = hierarchy, d = reciprocity. Spearman’s r is the reported value.

b.

c.

d.

Stage 3

Participants

For Stage 3, we aimed to have as many observations per distribution as we had in Stage 1. For Stage 1, which was powered for a nationally representative sample across age, race, and gender, we recruited for 450 observations with a final sample of 423 (Sample 1). In the current study, each participant would have to make three ratings for each measured construct rather than one (as in Stage 1), so we determined to recruit a third of the number of target participants as we had in our first sample to achieve similar statistical power. Accordingly, we recruited 150 participants, with 167 participants ultimately taking at least part of the survey (not all finished). Each participant who completed the survey⁴ was paid \$1.00. Sixty-four participants (64) were excluded on the basis of the pre-registered exclusion criteria (see [Table S12](#)), leaving us with a final sample of $N = 85$ (38 female, 46 male, 1 other/non-binary) ranging in age from 21 to 69 ($M_{\text{age}} = 34.82$, $SD_{\text{age}} = 10.66$); see [Table S13](#) for complete demographic information.

⁴ 18 “participants” in the raw data file had no location information recorded and the data posted to Qualtrics nearly a week after the completion of the survey. We assumed these were bots and deleted those 18 lines of data prior to applying the exclusion criterion. Hence the $n = 149$ “completed” participants reported in the manuscript: $167 - 18 = 149$.

Table S12: Summary of exclusion criteria for Stage 3

Exclusion criteria met	Type of check	Excluded <i>N</i>
Did not move sliders to (at least) 1 of 2 specified positions	Attention check	64

Table S13: Demographics of Sample 3 participants

Age	N (%)	Race	N (%)	Gender	N (%)
18 - 27	23 (27.06%)	White	59 (69.41%)	Female	38 (44.71%)
28 - 37	37 (43.53%)	Black/ African-American	12 (14.12%)	Male	46 (54.12%)
38 - 47	14 (16.47%)	Asian	5 (5.88%)	Other/ Non-binary	1 (1.18%)
48 - 57	6 (7.06%)	Hispanic/ Latinx	4 (4.71%)		
58+	5 (5.88%)	Multiracial	3 (3.53%)		
		Native American	2 (2.35%)		

Procedure

We informed participants that they would be shown a set of relationships and asked some questions about each. Then, participants were shown, in random order, each of the 10 relationship pairs that were used in Stage 2 along with a brief description. Because this study was conducted during the COVID-19 pandemic, we were concerned that participants' judgments about relationships might be affected by the unprecedented circumstances. To address this concern, we included a note asking participants to think of the relationships as they would be under more 'normal' circumstances.

For each of the 10 relationship pairs, participants were asked to rate it along three dimensions of social closeness and three dimensions of interdependence: see [Table S14](#) for the exact descriptions of each. As an example, for the sibling relationship and one of the social closeness dimensions, participants were asked, e.g., "In an ideal, well-functioning relationship between siblings, to what extent would the relationship be characterized by **deeply understanding each other**?" Responses were recorded on a sliding scale from 0 to 100 (see endpoint labels in [Table S14](#) below). Finally, we included similar individual demographic measures as in the previous samples.

Table S14: Dimensions of social closeness and interdependence

Construct	Dimension
Social Closeness	- In an ideal, well-functioning relationship between [relationship pair], to what extent would the relationship be characterized by deeply understanding each other ? (0 = Not at all; 100 = A great deal)
	- In an ideal, well-functioning relationship between [relationship pair], to what extent would the relationship be characterized by accepting and validating each other's natures ? (0 = Not at all; 100 = A great deal)
	- In an ideal, well-functioning relationship between [relationship pair], to what extent would the relationship be characterized by striving to care for and promote each other's overall well-being ? (0 = Not at all; 100 = A great deal)
Inter-dependence	- In an ideal, well-functioning relationship between [relationship pair], how frequently would they affect each other's thoughts, feelings, and behaviors? (0 = Not at all frequently; 100 = Very frequently)
	- In an ideal, well-functioning relationship between [relationship pair], in how many different ways would they affect each other's thoughts, feelings, and behaviors across different situations? (0 = Very few ways; 100 = A great variety of ways)
	- In an ideal, well-functioning relationship between [relationship pair], how strongly would they affect each other's thoughts, feelings, and behaviors? (0 = Not at all strongly; 100 = Very strongly)

Data preparation and analysis details

Raw data files (.csv) were prepared and analyzed using Python, within a Jupyter Notebook environment. Primary packages used: numpy, scipy, statsmodels, matplotlib, seaborn, pandas. For data files and all coding scripts, see the OSF link above.

Supplementary results

Table S15: Full-exclusion regression model for Stage 3

Predictor	β	Std. Error	95% CI [LL, UL]	p	Fit
(Intercept)	103.12	13.15	[44.62, 50.89]	< .001	
Relational Norms	.12	.01	[-5.23, .07]	< .001	
Social Closeness	.24	.35	[2.98, 8.87]	.48	
Interdependency	-.20	.33	[-2.32, 3.01]	.55	
Genetic Relatedness	1.13	3.52	[-1.95, 4.72]	.75	
Action Likelihood	-.51	.09	[-6.70, .03]	< .001	

Target Specificity .52 .18 [4.34, 9.75] < .01

Conditional $R^2 = .89$

Please note that due to the large proportion of exclusions, we re-ran the regression with no exclusions (as a sensitivity/robustness test), and the results are substantially the same. See the below table for the regression results with no exclusions:

Table S16: No exclusions regression model for Stage 3

Predictor	β	Std. Error	95% CI [LL, UL]	p	Fit
(Intercept)	104.38	15.54	[74.49, 134.26]	< .001	
Relational Norms	9.36	.96	[7.51, 11.21]	< .001	
Social Closeness	.44	.73	[-.96, 1.84]	.55	
Interdependency	-.36	.70	[-1.71, .1]	.61	
Genetic Relatedness	.89	3.41	[-5.67, 7.46]	.79	
Action Likelihood	-.53	.09	[-.70, -.36]	< .001	
Target Specificity	.54	.18	[.19, .88]	< .01	
					Conditional $R^2 = .89$

Figure S4: Study 3 results. Kernel density plot of expectations for social closeness and interdependency for 10 common relationships. Dot represents the mean, with 95% confidence intervals. The height of the curve represents density: the likely proportions of scores (relative to each function) that fall within the given range along the x-axis.