

Journal of Personality and Social Psychology

Generalized Morality Culturally Evolves as an Adaptive Heuristic in Large Social Networks

Joshua Conrad Jackson, Jamin Halberstadt, Masanori Takezawa, Kongmeng Liew, Kristopher Smith, Coren Apicella, and Kurt Gray

Online First Publication, September 21, 2023. <https://dx.doi.org/10.1037/pspa0000358>

CITATION

Jackson, J. C., Halberstadt, J., Takezawa, M., Liew, K., Smith, K., Apicella, C., & Gray, K. (2023, September 21). Generalized Morality Culturally Evolves as an Adaptive Heuristic in Large Social Networks. *Journal of Personality and Social Psychology*. Advance online publication. <https://dx.doi.org/10.1037/pspa0000358>

Generalized Morality Culturally Evolves as an Adaptive Heuristic in Large Social Networks

Joshua Conrad Jackson¹, Jamin Halberstadt², Masanori Takezawa³, Kongmeng Liew^{4, 5},
Kristopher Smith⁶, Coren Apicella⁷, and Kurt Gray⁸

¹ Department of Behavioral Science, Booth School of Business, University of Chicago

² Department of Psychology, University of Otago

³ Department of Behavioral Science, Hokkaido University

⁴ Graduate School of Science and Technology, Nara Institute of Science and Technology

⁵ School of Psychology, Speech and Hearing, University of Canterbury

⁶ Department of Anthropology, Washington State University

⁷ Department of Psychology, University of Pennsylvania

⁸ Department of Psychology, University of North Carolina at Chapel Hill

Why do people assume that a generous person should also be honest? Why do we even use words like “moral” and “immoral”? We explore these questions with a new model of how people perceive moral character. We propose that people vary in the extent to which they perceive moral character as “localized” (varying along many contextually embedded dimensions) versus “generalized” (varying along a single dimension from morally bad to morally good). This variation might be partly the product of cultural evolutionary adaptations to different kinds of social networks. As networks grow larger, perceptions of generalized morality are increasingly valuable for predicting cooperation during partner selection, especially in novel contexts. Our studies show that social network size correlates with perceptions of generalized morality in United States and international samples (Study 1) and that East African hunter-gatherers with greater exposure outside their local region perceive morality as more generalized compared to those who have remained in their local region (Study 2). We support the adaptive value of generalized morality in large and unfamiliar social networks with an agent-based model (Study 3), and in experiments where we manipulate partner unfamiliarity (Study 4). Our final study shows that perceptions of morality have become more generalized over the last 200 years of English-language history, which suggests that it may be coevolving with rising social complexity and anonymity in the English-speaking world (Study 5). We discuss the implications of this theory for the cultural evolution of political systems, religion, and taxonomical theories of morality.

Keywords: moral psychology, cultural evolution, generalized morality, natural language processing, field study

Supplemental materials: <https://doi.org/10.1037/pspa0000358.supp>

Plato and Aristotle did not agree on much, but their most famous disagreement may have concerned the nature of morality. The philosophers’ debate is famously captured by Raphael’s fresco “School of Athens” (see Figure 1). Plato, holding a copy of his *Timaeus*, points upward toward the heavens. His gesture represents his theory of nonphysical forms, including the belief that there is a

generalized “form of the good” that participates in everything that is virtuous. Aristotle, holding his copy of *Ethics*, points down to the earth to signal his belief that virtue must be judged based on the behavior of particular people in particular situations. In the *Ethics*, Aristotle wrote, “Good is said in many ways,” rejecting the idea of generalized goodness.

Joshua Conrad Jackson  <https://orcid.org/0000-0002-2947-9815>


The authors thank Ideliya Khismatova, Victoria Maghalli, Jasmine Zheng, and Ryan Drabble for their research assistance. They also thank Keith Payne, Katherine Gates, Kristen Lindquist, and Nava Caluori for their valuable comments on an earlier draft of this article.

Joshua Conrad Jackson played a lead role in conceptualization, data curation, formal analysis, investigation, methodology, project administration, visualization, and writing—original draft and a supporting role in funding acquisition. Jamin Halberstadt played a supporting role in conceptualization and an equal role in writing—review and editing. Masanori

Takezawa played a supporting role in methodology and writing—review and editing. Kongmeng Liew played a supporting role in methodology and writing—review and editing. Kristopher Smith played a supporting role in methodology and writing—review and editing. Coren Apicella played a supporting role in methodology and writing—review and editing. Kurt Gray played a lead role in funding acquisition, a supporting role in conceptualization, and an equal role in writing—review and editing.

Correspondence concerning this article should be addressed to Joshua Conrad Jackson, Department of Behavioral Science, Booth School of Business, University of Chicago, 5807 South Woodlawn Avenue, Chicago, IL 60637, United States. Email: joshua.jackson@chicagobooth.edu

Figure 1
Raphael's *School of Athens* Fresco



Note. Plato (center left) points to the heavens to support his theory of forms, whereas Aristotle (center right) points to the earth to signify his emphasis on concrete particular. Image credit: “The School of Athens” by Raphael (1509–1511). In the public domain. See the online article for the color version of this figure.

Which philosopher was right? Thousands of years have passed, yet people may still disagree. This article is devoted to understanding and predicting why people still hold diverse views of morality. We explore people’s different perceptions of *moral character complexity*. We suggest that some people, like Plato, may see moral character as low in complexity and endorse *generalized morality*. They view others as possessing a holistic moral character: an overall capacity for good or evil that transcends any specific behavior or situation. They invoke this quality when describing people as generally “good,” “moral,” “evil,” or “immoral” (see Supplemental Study 1). Perceiving generalized morality implies that someone who once acted immorally will be untrustworthy in the future and should be avoided or imprisoned.

At the other end of the spectrum are people who see moral character as complex and who endorse *localized morality*. Like Aristotle, they view morality as specific to particular people and particular situations. They use phrases like “generous with money” and “a responsible caregiver” to highlight context-specific elements of morality and “cowardly in battle” and “self-centered on vacation” to highlight context-specific elements of immorality. Perceiving localized morality implies that someone who once acted immorally could be trustworthy in the future, at least in a different context.

To take an example, consider learning that a billionaire evaded taxes using offshore bank accounts. Someone who subscribes to localized morality could believe that this billionaire is selfish with taxes but may act differently in other contexts, perhaps being kind toward his children or giving big tips at restaurants. Someone who subscribes to generalized morality could believe that their taxation-related behavior implies that this billionaire is immoral across all contexts—a tax evader is also a bad parent and a bad tipper. Likewise, when reports emerge of a politician sending their staffers

heartfelt “thank you” cards after a long campaign, people can either infer that this politician is specifically a thoughtful employer (perceiving localized morality) or that they are overall a moral person (perceiving generalized morality).

Few people should have a view of moral character that is absolutely generalized or completely localized. It is possible to endorse an intermediate level of moral character complexity by using words like “honest,” “stingy,” and “responsible” to describe behavior in a bounded range of situations. It is also possible to have different views of the structure of moral character for different people. People may be happy to say that a ruthless dictator is pure evil but more likely to say that their sister can act good or bad depending on the situation. Nevertheless, we suggest that there is person-level and even culture-level variance in the degree to which people assume moral character is generalized versus localized, and that this variance is meaningful and underappreciated in moral psychology.

People often ignore or fail to recognize this continuous spectrum from localized to generalized morality. The same dictionaries will define the concept of “morality” as both “moral quality or character”—signifying a singular generalized morality—and “conformity to the rules of right conduct”—signifying multiple rules that make up localized dimensions of morality (Dictionary.com, 2023). This blurring of localized and generalized morality resembles how the single word “intelligence” can communicate both localized forms of aptitude, such as math versus verbal proficiency, as well as a single underlying “general” intelligence (Spearman, 1904). While scholars continue to debate the best level at which to conceptualize intelligence (Kovacs & Conway, 2019), we focus here on how people conceptualize moral character and what might predict the extent to which people view moral character as generalized or localized.

We claim that assumptions about moral character are neither genetically innate nor randomly distributed, but that they are meaningfully influenced by cultural evolution in light of changing natural and social environments. We focus specifically on the nature of people's social networks. Drawing from theories of partner choice and cultural evolution (Apicella & Silk, 2019; Smaldino, 2019), we propose that perceptions of moral character become more generalized as societies get larger and more anonymous.

Changing perceptions of moral character due to social network size may be at least partly adaptive. In small social networks, complex perceptions of moral character allow people to use highly localized attributes (e.g., "courageous in conflict") as priors to predict someone's likelihood of cooperation in particular situations (e.g., whether someone will enlist in a draft). But in large social networks composed of strangers, which have become more common over human history, people do not receive enough information about their partners to develop many localized priors (Dunbar, 2009). Generalized morality may culturally evolve in these contexts as an adaptive heuristic that allows people to make moderately accurate predictions about their social partners' likelihood of cooperating in any context with limited information.

We test our model using five multimethod studies. Our first two studies show that social network size is correlated with perceptions of generalized morality in large correlational surveys of people living in the United States, Singapore, Germany, and Brazil, as well as in field data from East African hunter-gatherers. Our following two studies use an agent-based model and an experiment to show that generalized morality is adaptive for predicting cooperation in large social networks, especially when interacting with unfamiliar individuals in novel contexts. Our final study is a historical natural language processing (NLP) analysis, which suggests that views of moral character have become more generalized over history. Each study suggests that the moral domain has a dynamic structure; as human societies have grown larger and more complex, perceptions of moral character have become simpler and more generalized.

Character Complexity and the Challenge of Choosing Cooperative Partners

Theories of moral psychology differ in many ways, but most agree that moral judgment is closely tied to effective group living (Curry et al., 2019; Graham et al., 2013), helping especially to maintain group-based cooperation (Curry et al., 2019; Rai & Fiske, 2011). Like all animals who live in groups, humans evolved to live together without cheating each other out of resources, and sometimes even sacrificing personal gains for the good of the group (Axelrod & Hamilton, 1981; Nowak, 2006). Many nonhuman animals have evolved to cooperate in groups through kin selection; if group members share a high proportion of genes, it becomes adaptive (from the gene's standpoint) for individuals to sacrifice themselves for the sake of their relatives (Dawkins, 2016; Dunford, 1977). But human groups are unique because our groups are too large and genetically diverse to support cooperation through kin selection alone (Richerson & Boyd, 2008). This means that humans must develop other mechanisms to keep each other honest.

One of these supplemental mechanisms for sustaining human cooperation is partner choice. All animals can choose their interaction partners to varying extents (Cheney & Seyfarth, 1988; Hauser & Marler, 1993), but humans have a unique ability to

identify and interact with cooperative partners while avoiding defectors. If cooperators can accurately find one another using information about moral reputation, this will not only maximize their own personal gains but also benefit the group. When cooperators match up, defectors are forced to interact with each other, and their selfishness will create low joint outcomes, leading them to die off or convert to more cooperative strategies over time (Apicella & Silk, 2019; Barclay & Willer, 2007). Using moral reputation in partner choice has proven to maximize cooperation in laboratory studies with American and British students (Barclay & Willer, 2007; Sylwester & Roberts, 2010) and in field studies of horticulturalists and foraging communities (Bird et al., 2012; Macfarlan et al., 2012). Of course, cooperation is not the only factor involved in partner choice. People may also seek out peers who are highly competent, such as good hunters in hunter-gatherer communities (Apicella et al., 2012), but cooperation is nevertheless an important dimension of partner choice, especially in social networks with high relational mobility where defectors are more difficult to detect (Smith et al., 2022).

Yet there is a hidden challenge to partner choice: Cooperation happens in specific situations, and people are more cooperative in some situations than others. One's spouse might be a generous donor to charities while refusing to share their food at dinner. The teenager down the street might be a responsible babysitter, even though they cheat on school exams. This situational variation could be as simple as a change in interaction partner: A friend could be loyal to their company but disloyal to their spouse. In Supplemental Studies 2a and b, we empirically illustrate this variation in cooperation across situations using self-report vignettes and real-world behaviors. We find that there is a significant but low correlation in cooperation across contexts. Our pilot studies are not the first empirical evidence that behavior varies across situations—this has been an assumption in social psychology at least since Mischel's (2013) *Personality and Assessment*. Studies using single economic games like the prisoner's dilemma do not model this variation, but it is a significant challenge for real-world partner choice.

How should people overcome this challenge and learn to optimize their partner choices in the face of contextual variability in moral behavior? A "localized morality" approach could involve distinguishing and relying upon moral attributes that are localized to specific contexts. For example, perceiving someone as a "responsible caregiver" lends confidence that this neighbor can be trusted to take care of one's children for an evening without making assumptions about their behavior in other situations. Similarly, labeling someone as a "loyal employee" does not necessarily imply that they are also a loyal spouse. At the other extreme, a "generalized morality" approach could treat contextual variability as noise, averaging across this noise with generalized perceptions of some "moral" people as cooperative in all situations and other "immoral" people as uncooperative. At first glance, generalized morality seems limited because it ignores meaningful situational variability. But we suggest that, in sufficiently large and unfamiliar groups, this limitation becomes a strength, and generalized morality can actually optimize partner choice.

Generalized Morality as an Adaptive Heuristic in Large Social Networks

We propose that people's inferences about moral character may become more generalized in large and unfamiliar social networks,

and that this process of generalization may even be adaptive. Using the framework initially developed by Herbert Simon (Simon & Newell, 1958) and extended by Gigerenzer et al. (2011), we view partner choice using generalized morality as a “fast and frugal” cognitive process that “ignores information” (situational variability in cooperation) in order to “exploit the structure of the environment” (large social networks full of unfamiliar people).

Generalized morality, like other adaptive heuristics, should only be effective under certain conditions. Cooperation must correlate meaningfully across situations, and people must know very little about their interaction partners. Under these conditions, people can use the limited knowledge they have about a partner to infer whether or not this person will cooperate in a range of other contexts. For example, upon learning that a coworker regularly donates blood, one can infer that they will also be more likely to cooperate in situations where you have no information about their behavior, like giving to charitable causes, telling the truth, and contributing food to a potluck. These inferences will not always be accurate because cooperation only correlates moderately across situations. But they will be more accurate than random guessing about someone’s likelihood of cooperating in these novel situations.

In small social networks, generalized inferences may not be very useful because people are highly familiar with each other (Dunbar, 2003). When someone in a small hunter–gatherer group is looking for a hunting partner, they can find someone whom they know has a reputation as a cooperative hunter. When looking for someone to help with a building project, they can find someone who has shared building materials in the past. By learning about prior behavior in a broad range of specific contexts, people can develop localized priors that accurately predict how partners will behave across situations.

As groups get larger, however, imperfect but generalized tools for inference become more valuable (Smaldino, 2019). In a large organization or city, personal experience and gossip can only supply limited information about social partners’ previous behaviors (Dunbar, 2003). Most people in these social networks are virtual strangers—you may have some limited information about their past behavior (e.g., you know someone regularly donates blood), but no information about how they behave in most situations (e.g., charitable giving, telling the truth). Generalized morality could be a useful tool in these social environments because it will give you priors for predicting partner behavior in any situation based on very limited input. These priors will sometimes be wrong (e.g., someone who regularly gives blood may in fact be a chronic liar), but they will be right at a greater-than-chance rate because cooperation does correlate across situations (Davidson et al., 2015). This means that cross-domain inferences are better than randomly guessing at whether someone will cooperate in a novel dilemma. The cost of imperfect predictions is offset by the value of better-than-chance predictions in any given situation.

The adaptive value of generalized morality in large social networks could explain curious results from previous cross-cultural research. For example, cross-cultural field studies have found that hunter–gatherers such as the Hadza or Yasawa are less likely to explain actions in terms of abstract mental states or dispositions (e.g., my campmate stole because he is selfish) compared to respondents in large cities like Los Angeles (Barrett et al., 2016; Curtin et al., 2020; Gendron et al., 2020). One explanation for this pattern is that dispositional moral traits like “selfish” are generalized inference tools—they help people in large and anonymous social

networks infer cooperation in novel situations based on behavior in previous situations. In a city like Los Angeles, a character attribution like selfishness can help predict someone’s behavior in future cooperation dilemmas. But in a small hunter–gatherer camp, these predictions will be less useful because people already have a deep knowledge about how their social partners behave in different contexts. Our theory is also consistent with the recent finding that declines in kinship intensity due to Church prohibitions on cousin marriage may have preceded rises in generalized trust in Medieval Europe (Schulz et al., 2019). Drops in kinship intensity increase social network size because people must create bonds with more partners outside of their family unit. As people’s social networks grew more mobile and interaction partners became more anonymous, generalized morality may have become more functional in partner selection.

In the general discussion, we integrate our theory with other models of culture and cognition and distinguish it from theories of culture and the self (Markus & Kitayama, 1991; Shweder et al., 1984). We also describe the ways that generalized morality resembles and differs from making internal attributions about moral character, citing recent and highly relevant work in this space (Lammers et al., 2018).

An adaptive heuristic perspective also suggests that people’s views of moral character may also vary regionally and historically within societies. The United States of America as a whole is a large and urbanized country, but it also contains dozens of small, rural, and homogeneous communities that are culturally differentiated (Harrington & Gelfand, 2014). Perceptions of moral character may be more prevalent among larger social networks and localized morality should be more prevalent among smaller social networks. There may also be historical variation in perceptions of moral character. If societies grow larger and more anonymous over time, perceptions of moral character may become more generalized.

The Cultural Evolution of Generalized Morality

The final piece of our theory considers historical changes in how people have perceived character complexity. We argue that rising society size and urbanization throughout the Holocene has facilitated increasingly generalized perceptions of moral character. We can understand the mechanisms behind these dynamics using dual inheritance theories of cultural evolution, and we can track these trends using changes in natural language.

Human societies are larger, denser, and more anonymous than ever before (Murdock & Provost, 1973; Turchin et al., 2022). The scaling up of human groups is visible across any number of metrics: capital cities are more populous, infrastructure is more developed, governments control more territory and larger populations, and societies include more ethnic groups (Turchin et al., 2018). Escalating social complexity can be traced back to the end of the Ice Age when human groups were able to more easily live in large sedentary communities that grew their own food (Gupta, 2004). Social networks have become even larger and less familiar in the last five centuries as human groups have become more relationally mobile and less organized by kin ties (Blanc, 2020; Henrich, 2020; Newson et al., 2005; Schulz et al., 2019).

Theoretical models of cultural evolution suggest that human behavior and cognition have adaptively coevolved with these changes in social structure (Boyd & Richerson, 1988; Cavalli-Sforza & Feldman, 1981;

Richerson & Boyd, 2008). Models of cultural niche construction show that changes in social structure create adaptive pressures that select for human behaviors (Laland et al., 2001). Dual inheritance models of payoff-biased and prestige-biased transmission show that humans will preferentially copy behaviors of people who are successful or prestigious, which facilitates the spread of adaptive behavior throughout human groups (Chudek et al., 2012; Henrich & Gil-White, 2001; Kendal et al., 2009). The theory of “cognitive gadgets” argues that cognition—in addition to behavior and technology—can change via these cultural evolutionary pathways (C. Heyes, 2018; C. M. Heyes & Frith, 2014). These cultural evolutionary models make it plausible that individual people’s beliefs about moral character may have coevolved with social structure over human history and that generalized morality may have become more prevalent if it is an adaptive heuristic in large social networks.

Cultural evolutionary models of psychological change are usually hypothetical because of methodological limitations: It is impossible to ask people from the 18th century about their perceptions of moral character. However, innovations in language analysis now make it possible to infer historical changes in psychology (Jackson et al., 2020; Muthukrishna et al., 2021). Language offers a window into the mind since people use words to communicate their thoughts and feelings, and scholars have already used text analysis to track historical variation in stereotypes (Charlesworth et al., 2022; Garg et al., 2018), emotions (Jackson et al., 2019; Morin & Acerbi, 2017), and religious beliefs (Caluori et al., 2020; Jackson et al., 2021).

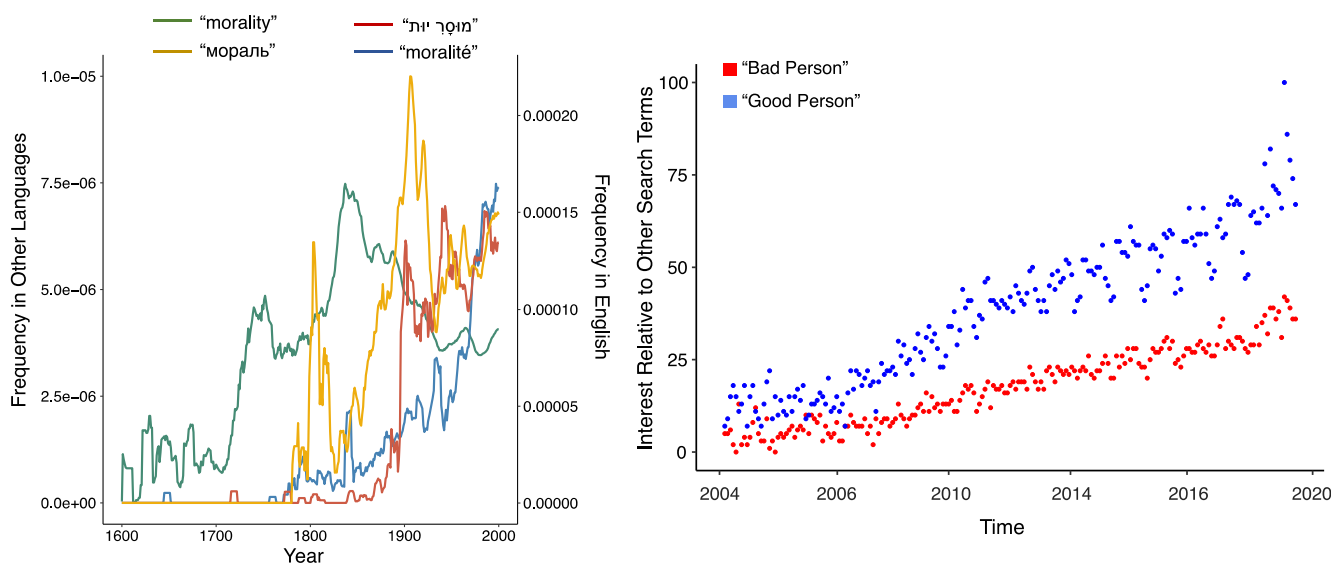
Language may also offer a window into the historical rise of generalized morality. Basic data show that words like “moral” and “morality,” which communicate generalized morality, spread in

multiple languages during the 16th and 17th centuries as European cities exploded in population. Today, as people’s online social networks explode in size, secular words connoting generalized morality like “good person” and “bad person” are growing in frequency (Figure 2). Both of these trends suggest that generalized morality may have coevolved with social network size, at least in the Western world.

These trends are not sufficient to prove the rise of generalized morality. Words rise and fall in prevalence for many different reasons. A more nuanced approach to historically tracking generalized morality in language could estimate the semantic association between different moral attributes over time. The English vocabulary contains dozens of words connoting localized moral attributes. For example, moral foundations theory (MFT) highlights differences between words like “respectful” and “lawful,” which connote cooperation with authority figures, from words like “wholesome” and “virtuous,” which connote following purity norms (Graham et al., 2009).

If perceptions of moral character have grown more generalized over time, localized moral words may become more interchangeable over time. This possibility is consistent with studies showing that people view a range of moral attributes as interchangeable (Landy et al., 2016, Supplemental Materials). In our Supplemental Materials, we present a similar study asking participants to make character judgments of their peers using attributes representing different categories of words, such as “honest,” “principled,” “responsible,” “fair,” and “loyal” (see Supplemental Study 3). Rather than using three, four, five, or six dimensions of moral judgment, which correspond to sets of cooperative contexts, people’s ratings reflected a single dimension ranging from morally good to morally bad. This perception of generalized morality may reflect the culmination of a

Figure 2
Generalized Morality in Language



Note. (Left) The rise of words communicating the concept of morality in language. Frequency represents the rate of each word as a proportion of all words in books. Supplemental Figure S4 shows that this rise has not characterized other localized moral attributes. (Right) Interest in the phrases “good person” and “bad person” over a time period where people’s social networks have expanded further due to social media. Data come from Google Trends and represent search frequency for each term within the United States, scaled so that a score of 100 represents the highest frequency data point. See the online article for the color version of this figure.

historical erosion of moral character complexity, and it may be possible to track this erosion using trends in natural language.

Contextualizing Our Research Among Other Theories of Moral Psychology

Discussions about the meaning of morality and moral character have a long history, dating back at least to debates between Plato and Aristotle about the nature of virtue. But contemporary moral psychologists have focused less on this debate and more on Plato and Aristotle's taxonomies of virtue. Despite their disagreements, both philosophers wrote similar lists of "cardinal virtues." Plato's *Republic* cited "wisdom," "temperance," "bravery," and "justice," and Aristotle's *Rhetoric* further added "magnificence," "magnanimity," "liberality," "gentleness," and "prudence."

Modern theories of modern psychology have renovated these taxonomical lists with an added assumption that humans evolved to reward these virtues because they increased cooperation in early human groups. The most popular of these adaptive moral taxonomies is MFT (Graham et al., 2009; Haidt & Joseph, 2004). MFT argues that the human mind contains a set of innate and psychologically distinct mechanisms that prepare people to moralize a specific set of values, including harm, fairness, loyalty, authority, and purity, that are argued to make group living more successful. Although there is little evidence that these concerns stem from distinct psychological mechanisms—all judgments of moral acts seem to revolve around a template of perceived harm (Schein & Gray, 2018)—MFT was an important development because it deconstructed the monolith of "morality" into descriptively different acts that extended beyond classical questions of rights and justice (Graham et al., 2013).

Another theory that resembles MFT is the morality as cooperation hypothesis (MAC), developed by anthropologists (Curry et al., 2019) to explain the connection between the moralization of different acts and the evolution of cooperation with societies. MAC outlines seven different moral concerns and explicitly connects them to domains of cooperation (e.g., a moral concern for "bravery" facilitates cooperation during intergroup conflicts; a moral concern for "reciprocity" facilitates cooperation during trade). This theory outlines how these moral concerns could inform partner choice judgments that foster social cooperation.

We follow MAC in suggesting that moral acts are important for helping to facilitate future cooperation but diverge from both MAC and MFT in two ways. First, we more explicitly distinguish the structure of judgments about moral character and judgments about cooperative acts. Although these two are obviously related, as people use behaviors to infer moral character, person-level moral character judgments seem most consequential for predicting people's future behavior (Hartman et al., 2022; Pizarro & Tannenbaum, 2012).

Second, we argue that moral character judgments should not fit a specific number of dimensions. Even if scientific taxonomies might be useful by carving up and labeling the complexity of morality into three (Rozin et al., 1999), four (Haidt & Joseph, 2004), five (Graham et al., 2009), six (Iyer et al., 2012), seven (Curry et al., 2019), or even 10 categories (Schwartz, 2012), we suggest that the minds of everyday people around the world are unlikely to hew closely to these researcher-made divisions when judging the moral character of others. Serious cross-cultural efforts have already shown that

there is no single way of carving up the nomological network of morality across people and cultures (Atari et al., 2022).

We instead predict that people use moral information dynamically in the service of selecting future cooperation partners, balancing the predictive power of localized morality with the cognition-saving ease of generalized morality. Whether there might be a way to take moral complexity from across people and societies and find an average number of moral concerns is an interesting question, but here we are less concerned with any potential *average* than with the *variance*. Consistent with metascientific calls to document variability in psychological processes (Yarkoni, 2022), we focus on exploring variability in how people perceive the structure of moral character and testing whether that variability is tied to social network size.

Summary of Research Program

Here we have described a new theoretical model of how people perceive moral character. In this model, people vary in their perceptions of character complexity, and this variance is sensitive to cultural evolutionary pressures. In large social networks, generalized morality is an adaptive heuristic that may spread through payoff-biased transmission because it increases the accuracy of partner choice predictions. However, localized morality will be a more effective and prevalent partner choice strategy in small social networks. As a result, individuals and cultural groups should vary widely in their moral character complexity across time and space.

Our empirical studies test this theory in terms of three hypotheses:

Hypothesis 1: Social network size correlates with perceptions of generalized morality.

Hypothesis 2: Perceptions of generalized morality improve cooperation predictions in partner choice dilemmas in large and unfamiliar (vs. small and familiar) networks.

Hypothesis 3: Perceptions of moral character have become more generalized over human history, at least within the English-speaking world.

Studies 1 and 2 test Hypothesis 1. Studies 1a and b are cross-regional and cross-national surveys which show that social network size is associated with perceptions of generalized morality. Study 2 conceptually replicates this finding in Hadza hunter-gatherers, showing that level of external exposure outside of Hadzaland explains whether Hadza individuals perceive the morality of their campmates as generalized.

Studies 3 and 4 test Hypothesis 2. Study 3 is an agent-based model which shows that, given plausible assumptions, generalized morality becomes increasingly valuable as social networks grow larger and less familiar. This model also shows that perceptions of moral character should become more generalized in large social networks if agents learn through payoff-biased transmission. Study 4 is an experiment that shows that generalized morality is particularly valuable when people interact with unfamiliar partners in novel situations.

Study 5 examines the historical dimension of our theory, testing Hypothesis 3. We use word embeddings trained on massive corpora of English-language text to show that different moral attributes (e.g., fair, loyal, caring) have become more semantically generalizable

over the last 200 years of human history. The semantic convergence of these attributes is consistent with our prediction that perceptions of moral character have become increasingly generalized and complements our cross-sectional evidence in Studies 1 and 2.

Open Science and Ethics Statement

The design, procedure, and analysis strategy for all studies except for Studies 2 and 5 were presented and approved during a dissertation proposal. Studies 2 and 5 were added after the dissertation proposal and were preregistered before data compilation and analysis. All data, code, and preregistrations can be found at https://osf.io/d3eyt/?view_only=b08cd944086147338f18462d145ae418. Our Supplemental Materials contain a preregistered direct replication of Study 4 and a replication of Study 5 using a different corpus. We received internal review board approval prior to conducting this research (Protocol No. 18-3040).

Hypothesis 1: Generalized Morality Is Prevalent in Large Social Networks

Our first claim is that there is meaningful cross-cultural variation in whether people perceive moral character as generalized or localized. In smaller and more familiar social networks, people should be more likely to view moral character as multidimensional and localized. However, people in larger social networks should view moral character as more unidimensional and generalized. Studies 1 and 2 test this hypothesis using large online surveys (Study 1) and field data among Hadza hunter-gatherers (Study 2).

Studies 1a and b: Generalized Moral Judgment Correlates With Social Network Size

Studies 1a and b tested whether social network size was associated with perceptions of generalized morality using surveys of people around the world (Study 1a) and a representative sample of Americans (Study 1b). In these surveys, we collected two individual-level indicators of social network size and correlated these indicators with self-reported perceptions of generalized morality. Our statistical models controlled for sociodemographic variables like age, gender, religiosity, education, and socioeconomic status (SES). The studies were direct replications, and we present their methods and results together.

One strength of these studies was that we modeled within-nation variation. In some previous research, researchers have measured variables like relational mobility at the nation level, suggesting, for example, that France has greater relational mobility than Morocco (Thomson et al., 2018). But modern-day nations are complex and contain just as much within-group variation in social network characteristics as between-group variation. To maximize the precision of our measures and avoid these kinds of nation comparisons, we collected individual-level data on participants' social network characteristics.

Method

Participants. For Study 1a, we advertised for 1,000 participants from four nations (The United States, Brazil, Singapore, and Germany) using Qualtrics panels. We chose these nations because

they cover different world regions and they vary in their cultural tightness: the strictness of cultural norms (Gelfand et al., 2011). Since people are more likely to make moralized judgments in tight cultures (Jackson et al., 2021), we sought to test whether the relationship between moral beliefs and social network size was robust in a sample that included people from both tight and loose cultures. We recruited 1,000 participants because this was the largest sample that we could afford given Qualtrics panels pricing. The United States and Singapore are English-speaking countries, but Brazil and Germany are not. Therefore, native speakers translated the survey from English into Portuguese and German for these speakers using standard translation and back-translation procedures. In total, 1,044 participants (484 men, 560 women; $M_{\text{age}} = 44.45$, $SD_{\text{age}} = 16.04$; 267 from Singapore, 256 from Brazil, 260 from Germany, and 261 from the United States) completed the survey.

For Study 1b, we advertised for 2,000 American participants using the Qualtrics panels service. We determined sample size by recruiting as many participants as possible given the cost constraints of the panel service. Participants were pseudorepresentative in that they were recruited to be nationally representative on the key dimensions of age, political party affiliation, race, and region of the country (South, Northeast, Midwest, West). In total, 2011 participants (504 men, 1,501 women, six nonbinary; $M_{\text{age}} = 50.49$, $SD_{\text{age}} = 16.40$) completed the survey.

Measures.

Social Network Size. We operationalized social network size using two key metrics: participants' self-reported number of daily face-to-face interaction partners (as a proxy for the size of their in-person social network) and participants' self-reported number of friends on Facebook (as a proxy for the size of their virtual social network). Given these measures, owning a Facebook account was a prerequisite for participating in the surveys. Participants were excluded from analyses if they listed more than 5,000 friends on Facebook since Facebook does not allow more than 5,000 friends. Participants were also excluded if they reported interacting with more than 1,000 unique people per day, which would require a unique interaction every 43.2 s over 12 hr. This procedure excluded 25 participants from the Study 1a sample and 11 participants from the Study 1b sample. Results are essentially identical regardless of these exclusions.

Generalized Morality. We designed a six-item scale to assess participants' perception of generalized morality. Items 1–3 pertained to participants' belief in generalized morality (“At their core, people are either morally good or morally evil”; “Every person has a basic good or evil moral character”; “All forms of cooperation or noncooperation can be traced to people’s underlying moral character”), and Items 4–6 pertained to participants' tendency to use generalized morality in their social interactions (“I often think about people’s underlying moral character when I interact with them”; “When deciding whether to trust someone, I try to gauge their underlying moral character”; “I rarely, if ever, need to gauge someone’s fundamental underlying moral character”—reverse coded). Participants responded to each item on a 1–100 scale anchored at 1 (*strongly disagree*) and 100 (*strongly agree*).

The scale had a Cronbach's α value of 0.66 in Study 1a and 0.71 in Study 1b. These reliability values were relatively low because of the reverse-scored item, which had low item-total correlations in both Studies 1a ($r = 0.19$) and 1b ($r = 0.39$). When we removed the reverse-scored item, the scale showed higher Cronbach's α values in

both studies ($\alpha > 0.75$). It is not uncommon for reverse-scored items to decrease scale reliability (Gelfand et al., 2011). We ultimately ran our analyses with and without the reverse-scored item. We found that every significant result was replicated with or without the item. Here we present the results with the item included, and we have made our code and data publicly accessible so that readers can see how results change when the item is excluded.

Sociodemographic Controls. Variables like someone's number of Facebook friends represent much more than social network size: They also signal a variety of sociodemographic characteristics like age, gender, and SES. To control for potential confounds associated with these sociodemographic variables, we measured age, self-identified gender, SES, religiosity, and education in both Studies 1a and 1b. We measured SES using people's responses to the McArthur Ladder item, which asks people to rate themselves higher or lower on a 10-rung ladder, where higher values represent people with the most money, highest education, and best jobs. We measured religiosity using the six-item Supernatural Beliefs Scale (Bluemke et al., 2016). We measured education level with a dummy-coded measure of whether people completed a 4-year college degree.

Analytic Plan. Both the number of Facebook friends and the number of daily interaction partners were positively skewed, so we log-transformed them prior to analysis. We began testing hypotheses with zero-order correlations between perception of generalized morality and each metric of social network size. We then replicated these analyses with multiple regressions that controlled for demographic characteristics. In all studies presented in this article, coefficients associated with lower case *b* are unstandardized, and those associated with β are standardized.

Results

Was generalized morality perception higher among people who inhabit large social networks? In support of our hypothesis,

zero-order correlations in the Study 1a data set showed that generalized morality perception was correlated with both the number of Facebook friends, $r(1011) = 0.17, p < .001$, and the number of everyday interaction partners, $r(1011) = 0.11, p < .001$. These correlations persisted after controlling for SES, age, gender, religiosity, education, and fixed effects representing the four countries in our analysis (see Table 1).

We found the same pattern of results in the Study 1b data set, although the associations were smaller in the single-nation survey. The association between generalized morality perception and number of Facebook friends remained statistically significant, $r = 0.10, p < .001$, even with covariates (see Table 2), but the association with number of everyday interaction partners was small, $r = 0.06, p = .008$, and failed to reach statistical significance with covariates (see Table 2). Nevertheless, the fact that three of the four expected relationships were robust across two diverse data sets with a variety of control variables suggests that the significant results were unlikely to be false positives.

One concern is that our questionnaire did not measure participants' perceptions of moral character specifically but captured their interest in people more generally. More sociable people might have larger social networks and might also be more adept at using moral character to differentiate their social partners.

To address this concern, we replicated our regressions with an alternate index that just included Items 1–3, concerning people's beliefs about morality. Controlling for the same covariates we summarize in Tables 1 and 2, this alternative index of generalized morality perception was significantly associated with number of Facebook friends, $b = 1.48, SE = 0.71, \beta = 0.08, t = 2.11, p = .04$, 95% CIs [0.10, 2.87], and everyday interaction partners, $b = 3.29, SE = 1.27, \beta = 0.08, t = 2.59, p = .009$, 95% CIs [0.80, 5.77], in Study 1a, and was significantly associated with Facebook friends, $b = 1.51, SE = 0.55, \beta = 0.07, t = 2.76, p = .006$, 95% CIs [0.44, 2.59], but not with everyday interaction partners, $b = 0.47$,

Table 1
Correlates of Generalized Morality Endorsement in Study 1a

Outcome predictor	DFs	<i>b</i> (<i>SE</i>)	<i>t</i>	<i>p</i>	95% CIs
Model 1	1,003				
Facebook friends		1.30 (0.54)	2.42	.02	[0.25, 2.35]
SES		0.68 (0.22)	3.04	.002	[-0.10, 0.05]
Age		-0.03 (0.04)	-0.69	.49	[-0.10, 0.05]
Gender		-1.37 (0.94)	-1.45	.15	[-3.22, 0.49]
Religiosity		1.79 (0.31)	5.85	<.001	[1.19, 2.39]
Education		-1.28 (1.12)	-1.15	.25	[-3.47, 0.91]
Singapore		-0.74 (1.38)	-0.53	.59	[-3.44, 1.97]
Brazil		4.18 (1.51)	2.77	.006	[1.22, 7.15]
Germany		-2.86 (1.43)	-2.00	.046	[-5.66, -0.06]
Model 2	1,003				
Everyday interaction partners		2.66 (0.97)	2.75	.006	[0.76, 4.55]
SES		0.69 (0.22)	3.13	.002	[0.26, 1.13]
Age		-0.04 (0.03)	-1.17	.24	[-0.11, 0.03]
Gender		-1.04 (0.94)	-1.10	.27	[-2.90, 0.82]
Religiosity		1.82 (0.30)	5.99	<.001	[1.22, 2.42]
Education		-1.46 (1.12)	-1.31	.19	[-3.65, 0.73]
Singapore		-0.42 (1.38)	-0.30	.76	[-3.12, 2.29]
Brazil		3.99 (1.51)	2.64	.009	[1.02, 6.96]
Germany		-3.94 (1.43)	-2.76	.006	[-6.75, -1.13]

Note. Both social network size proxies have been log-transformed. Country fixed effects are contrasted against the United States in this model. DFs = degrees of freedom; SE = standard error; CI = confidence interval; SES = socioeconomic status.

Table 2
Correlates of Generalized Morality Perception in Study 1b

Outcome predictor	DFs	<i>b</i> (<i>SE</i>)	<i>t</i>	<i>p</i>	95% CIs
Model 1	1993				
Facebook friends		1.44 (0.43)	3.35	<.001	[0.60, 2.29]
SES		-0.04 (0.20)	-0.22	.83	[-0.43, 0.34]
Age		-0.08 (0.03)	-3.36	<.001	[-0.14, -0.04]
Gender		-2.36 (0.90)	-2.62	.009	[-4.13, -0.60]
Religiosity		2.19 (0.27)	8.26	<.001	[1.67, 2.70]
Education		-1.92 (0.89)	-2.17	.03	[-3.66, -0.18]
Model 2	1993				
Everyday interaction partners		1.56 (0.81)	1.94	.053	[-0.02, 3.15]
SES		-0.06 (0.20)	-0.30	.76	[-0.45, 0.33]
Age		-0.10 (0.03)	-4.02	<.001	[-0.15, -0.05]
Gender		-1.97 (0.90)	-2.19	.03	[-3.73, -0.21]
Religiosity		2.25 (0.26)	8.51	<.001	[1.73, 2.77]
Education		-2.02 (0.90)	-2.26	.02	[-3.79, -0.27]

Note. Both social network size proxies have been log-transformed. DFs = degrees of freedom; *SE* = standard error; CI = confidence interval; SES = socioeconomic status.

SE = 1.03, β = 0.01, *t* = 0.45, *p* = .65, 95% CIs [-1.55, 2.48], in Study 1b.

Discussion

Studies 1a and b found a correlation between social network size and perceptions of generalized morality. This association replicated regardless of whether we operationalized social network size via people's online social networks (i.e., number of Facebook friends) or offline social networks (e.g., number of everyday face-to-face interaction partners). Moreover, the relationships remained statistically significant when adding a variety of sociodemographic controls in all but one test. The association was similar across an international sample (Study 1a) and a sample of American participants (Study 1b). These studies offer early evidence that people living in larger social networks hold a more generalized conception of morality.

A limitation of these studies is that we asked directly about generalized morality using a self-report questionnaire, which could confound beliefs about the structure of moral character with beliefs about the importance of moral character. People could agree with statements such as "Every person has a basic good or evil moral character" because they view moral character along a single dimension, but they could also plausibly agree with these statements because they ascribe a great deal of importance to moral character. Another limitation of our scale was that all the items described generalized morality. None of the items described localized morality.

In Study 2, we addressed these limitations with a measure that captured people's assumptions about the covariance of moral attributes such as "generous" and "honest" without confounding people's subjective beliefs about the importance of these attributes. This study also took place in a different cultural context: an East African hunter-gatherer group in the midst of rapid cultural change.

Study 2: External Exposure and Morality in a Hunter-Gatherer Society

Study 2 was a preregistered field study of moral character perceptions among the Hadza, a nomadic hunter-gatherer group

living along the Central Rift Valley in northern Tanzania. The Hadza historically lived in bands of about 30 children and adults (Marlowe, 2010). However, Hadza life is rapidly changing because of the growing encroachment of outside society via a rising number of aid workers, missionaries, and ethnotourists (Apicella, 2018; Apicella et al., 2014; Crittenden et al., 2017). Many Hadza hunter-gatherers have also begun working in nearby cities, which has created a social ecology where some Hadza have had substantially more external exposure outside of their local region than others. In a group of Hadza hunter-gatherers sampled in 2019, 40% reported living outside of Hadzaland at some point, 25% reported having held a job that pays money, and nearly 60% claimed to have heard of the former U.S. President Barack Obama (Smith & Apicella, 2020b).

Cultural, economic, and political change in Hadzaland has led to changes to diet (Crittenden et al., 2017) and foraging strategies (Pollom et al., 2021), but it could also have changed people's beliefs about morality (Workman et al., 2022). In particular, our model would suggest that Hadza individuals' perceptions of moral character have become more generalized as their social networks have grown larger and that current-day Hadza who have developed larger and more unfamiliar social networks should perceive moral character as more generalized than those whose social networks have remained small.

We tested this account by estimating external exposure in a sample of Hadza hunter-gatherers. We measured several indicators measuring whether participants had exposure outside their local social network, including whether they had lived outside Hadzaland, whether they had worked in jobs, and whether they had learned about Western culture. Greater external exposure expands the network of potential cooperative partners who do not have the same food-sharing norms as the Hadza, which may motivate greater interest in attending to potential partners willingness to cooperate (Smith et al., 2022).

We also measured Hadza participants' perceptions of generalized versus localized morality via their beliefs about the moral attributes of people in their camp. Participants ranked their campmates on five different moral attributes. We operationalized perceptions of generalized morality via whether participants ranked their campmates consistently or inconsistently across these attributes. We

predicted that Hadza participants with more external exposure would be more likely to rank campmates as consistently moral or immoral across attributes with less ranking variation across attributes, displaying generalized morality.

Method

Preregistration. The data for this study were originally collected by Smith and Apicella (2020a, 2020b) in 2019, and the measures and sampling procedures are thoroughly described in that article. The first author of this article wrote a detailed preregistration to test the hypotheses using this data set before conducting any of the analyses that we report here. Our preregistration included the sample size, measures involved, and configuration of our statistical models.

Sample. Data were collected using a snowball sampling procedure. A Tanzanian research assistant collected data by visiting a camp and interviewing each member of the camp individually in Swahili. Members of that camp would then direct the interviewer to another camp until they could not identify any more camps. The number of adults in each camp was relatively small, ranging from 11 to 20. Eighty-seven participants originally completed the study. However, one participant was subsequently excluded because he refused to evaluate his campmates on hard work, saying that no one in the camp worked to gather food. Another participant was excluded because he had been living in a local village until recently and claimed that he was not sufficiently familiar with any of his campmates to rate their moral character. The final sample involved 85 participants who rated 91 campmates (673 total observations). The sample included 41 women and 44 men ($M_{\text{age}} = 36.33$, $SD_{\text{age}} = 13.87$) from six different camps.

Moral Perceptions Measure. Participants ranked eight individuals in their camp on five attributes that were relevant to morality: generosity (“Who is the most generous?”), honesty (“Who is the most honest?”), effort toward foraging (“Who works the hardest to get food?”), partner choice preference (“who would you most like to live with if you were to move camp tomorrow?”), and having a good heart (“Who has a good heart?”). These rankings were done using cards in one-on-one interviews with research assistants, and the full details of the procedure are summarized in Smith and Apicella (2020a).

We used these rankings to create two measures of generalized morality. The primary measure represented the standard deviation of the rankings across targets. For example, someone who ranked a given campmate k as top-ranked across generosity, honesty, and so forth, would receive a higher generalized morality score than someone who assigned campmate k different generosity and honesty rankings. The logic behind this measure was that homogeneous rankings would reflect an assumption that the different moral attributes were semantically interchangeable, whereas variable rankings would not reflect this assumption. Figure 3 illustrates the ranking procedure and stylized illustrations of “generalized morality” and “localized morality” in rankings of moral attributes.

We also created a secondary measure of generalized morality perception focused on the “good heart” attribute. To our knowledge, and according to previous dictionaries, there is no single word in Hadza language that resembles generalized morality like “moral” or “immoral” in English (see Supplemental Study 1, for formal analyses of how people interpret the meaning of “moral”); “good heart” was the closest term to approximate the concept of

generalized morality (Miller et al., 2013; Purzycki et al., 2018; Smith & Apicella, 2020a). Since an English-speaking participant who endorses generalized morality would agree that someone’s morality is strongly predictive of other cooperative attributes (see Studies 1a and b), we theorized that a Hadza participant who endorses generalized morality would view “good heart” as strongly related to all the attributes that they viewed in the study. This measure was therefore the magnitude of the correlation between the good heart ranking and the other rankings. Results using the primary and secondary measures were similar, so we present the analyses of the primary measure in the main text and the analyses of the secondary measure in the Supplemental Materials.

External Exposure. We measured external exposure via 10 different indicators: Years of school (log-transformed), whether participants could count to 10 in Swahili, whether participants had held a job outside of Hadzaland, whether participants could identify the president of Tanzania, whether participants could identify Barack Obama, whether participants could identify Nelson Mandela, whether participants could identify Mahatma Gandhi, whether participants had lived outside of Hadzaland, and whether participants had lived in the neighboring city to Hadzaland.

An exploratory factor analysis identified two factors within this exposure measure. One factor (eigenvalue = 1.41) contained the three items about knowledge of foreign figures (Gandhi, Obama, and Mandela). The other factor (eigenvalue = 3.79) contained the remaining items. The first factor appeared to be tapping knowledge of foreign culture, whereas the second factor appeared to be tapping personal exposure via how much time participants had spent outside of Hadzaland working and in school. These factors were not independent, with a robust positive correlation ($r = 0.40$, $p < .001$). Given this correlation, we began by fitting a single composite index and then analyzing the two factors separately.

Analytic Plan. We tested these predictions using cross-classified multilevel models with Gaussian estimation in which observations were nested within judge (the person doing the ratings) and subject (the person being rated). All analyses controlled for preregistered variables that we identified as potential confounds: judge age, judge sex, subject age, subject sex, and whether the subject was the judge’s spouse. All of our findings were replicated without these controls.

Results

Results supported our prediction that Hadza hunter-gatherers with more external exposure would have greater perceptions of generalized morality. External exposure was significantly associated with perceptions of generalized moral character, measured via lower standard deviations of moral attribute rankings, $b = -0.48$, $SE = 0.14$, $\beta = -0.21$, $t = -3.53$, $p < .001$, 95% CIs $[-0.75, -0.21]$. Table 3 shows that this link remained significant controlling for judge age and sex, subject age and sex, and whether the subject was the judge’s spouse. In Figure 4, we show the pairwise correlation between each attribute for participants below the median on external exposure and above the median on external exposure. The plot illustrates how participants’ rankings were always correlated to a certain extent (when participants gave someone a high “honesty” rank, they also tended to give them a high “generosity” rank). However, these correlations were higher for participants with higher exposure.

In a second model, we separated external exposure into the “foreign knowledge” and “personal experience” indices. As we

Figure 3
The Study 2 Moral Perception Measure

Generalized Morality				Localized Morality			
	A ₁ Honest	A ₂ Generous	A ₃ Effortful		A ₁ Honest	A ₂ Generous	A ₃ Effortful
C ₁	4	4	4	C ₁	1	6	3
C ₂	1	1	1	C ₂	4	2	7
C ₃	6	6	6	C ₃	8	1	8
C ₄	7	7	7	C ₄	2	8	6
C ₅	5	5	5	C ₅	7	5	2



Note. The top panel illustrates stylized grids showing how we measured perceptions of generalized morality in Study 2. Numbers in these grids represent rankings of campmates (C) on moral attributes (A), and shading represents higher rankings (more moral) versus lower rankings (less moral). This 5 × 3 grid in the figure is truncated (participants actually ranked eight campmates on five attributes). Generalized morality was operationalized through more consistent rankings of moral attributes (top left), whereas localized morality was operationalized through more inconsistent rankings of moral attributes (top right). The bottom panel shows the research assistant administering the moral perceptions measure. See the online article for the color version of this figure.

mention in the methods, the two factors correlated highly ($r = 0.40$), and neither factor was a dominant predictor when they were modeled together. Foreign knowledge was statistically significant, whereas personal experience was not. However, neither effect was as large in magnitude as the combined exposure index. Furthermore, these effects were not significantly different from each other in their magnitude, suggesting that knowledge of foreign cultures and personal opportunities to broaden one's social network may be intricately linked in this sample.

Discussion

This study offered further support for our hypothesis that social network size correlates with perceptions of generalized morality (Hypothesis 1). Hadza hunter-gatherers' level of external exposure was associated with their perception of generalized morality when they ranked their campmates on five moral attributes. Participants with greater external exposure were more likely to rank campmates as less variable across moral attributes compared to participants who had less foreign knowledge and less experience outside of Hadzaland.

This study complements Studies 1a and b. The limitation of Study 2 was that we did not measure social network size directly, instead relying on external exposure. Study 1 compensated for this limitation with a direct measure of social network size. The limitation of Study 1 was that it used a self-report measure of generalized morality, which could be confounded with moral importance. Study 2 compensated for this limitation with a less demand-laden measure of generalized morality, which measured participants' assumptions about the structure of moral character more purely.

In exploratory analyses, we also separated our measure of external exposure into two indices measuring foreign knowledge and personal experiences (e.g., in schools and nearby cities). This analysis allowed us to estimate whether generalized morality was more tied to knowledge of Western culture or personal experiences in large social networks. Our model accommodates both mechanisms, as we suggest that people can update their beliefs about moral character through social learning (supporting the role of cultural transmission of generalized morality) but also that adopting a perception of generalized morality should be functional for people living in large and unfamiliar social networks (supporting personal experience). Separating the factors yielded inconclusive results; the

Table 3
Correlates of Generalized Morality in Study 2

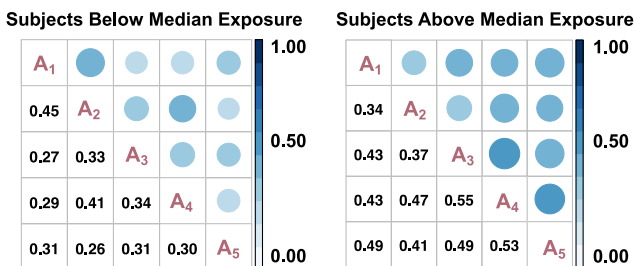
Outcome predictor	<i>N</i>	<i>b</i> (<i>SE</i>)	<i>t</i>	<i>p</i>	95% CIs
Model 1					
External exposure	673	−0.51 (0.15)	−3.42	<.001	[−0.80, −0.22]
Judge age		−0.003 (.003)	−0.75	.46	[−0.008, 0.003]
Judge male		0.05 (0.09)	0.55	.59	[−0.13, 0.23]
Participant age		−0.0002 (0.002)	−0.13	.90	[−0.004, 0.003]
Participant male		−0.05 (0.05)	−1.00	.32	[−0.16, 0.05]
Participant spouse		0.08 (0.14)	0.58	.56	[−0.20, 0.36]
Model 2					
Foreign knowledge	673	−0.27 (0.13)	−2.08	.04	[−0.53, −0.02]
Personal experience		−0.02 (0.02)	−0.95	.34	[−0.07, 0.02]
Judge age		−0.002 (0.003)	−0.61	.55	[−0.008, 0.004]
Judge male		0.01 (0.09)	−0.11	.92	[−0.17, 0.19]
Participant age		−0.0003 (0.002)	−0.14	.89	[−0.004, 0.003]
Participant male		−0.05 (0.05)	−0.98	.33	[−0.16, 0.05]
Participant spouse		0.09 (0.14)	0.61	.54	[−0.20, 0.37]

Note. The *N* term represents the total number of rankings in the analysis. *SE* = standard error; *CI* = confidence interval.

two factors correlated highly with each other, and neither factor had a particularly strong relationship with perceptions of generalized morality when they were modeled together. Knowledge of foreign culture and local experience with large social networks may be too difficult to differentiate in this sample.

A limitation of both Studies 1 and 2 is that they are correlational. They do not test whether increases in social network size actually produce perceptions of generalized morality. This leaves open the possibility of reverse causality: People with generalized morality might be more attractive social interaction partners and may accrue larger social networks. We consider this possibility interesting and plausible but not sufficient to account for why generalized morality spreads in large social networks. We therefore designed Studies 3 and 4 to formally examine how social network size and partner anonymity could cause greater prevalence of generalized morality as an adaptive heuristic.

Figure 4
Pairwise Correlation of Moral Attribute Rankings by Exposure



Note. The left panel shows the pairwise correlations between moral attribute rankings made by participants below the median of external exposure. The right panel shows that pairwise correlations between moral attribute rankings made by participants above the median of external exposure. The numerical correlation coefficient is displayed in the lower left quadrant, and the top-right quadrant shows circles sized and shaded according to the size of the correlation. See the online article for the color version of this figure.

Hypothesis 2: Generalized Morality Is an Adaptive Heuristic in Large Social Networks

Studies 3 and 4 test whether generalized morality is an adaptive heuristic in large social networks (Hypothesis 2) using more tightly controlled and internally valid studies than the correlational Studies 1 and 2. In Study 3, we present an agent-based model that explores whether this adaptive function is plausible given reasonable assumptions (e.g., that people learn using “payoff-biased transmission,” in which they copy the behavior of successful partners). Study 4 is a more focused experiment using human subjects. This study shows that generalized morality can improve people’s ability to predict the cooperation of unfamiliar partners in novel situations. These studies do not imply that generalized morality is prevalent in large social networks due to function alone. But they do suggest that function could be an important reason why generalized morality is common in large social networks.

Study 3: Agent-Based Model

Our agent-based model¹ was designed to test whether unfamiliarity and social network size would increase the adaptiveness of generalized morality. We simulated a series of artificial social

¹ *Note about agent-based models.* Agent-based models involve artificial worlds in which computerized “agents” interact in specific rule-based ways. They are not typically meant to test hypotheses about individual people’s behavior but rather hypotheses about how population-level patterns can emerge over time given plausible individual-level behaviors (Jackson et al., 2017). A significant strength of these models is that they can track behavior over long periods of time with no missing data, accommodate replications under the same set of conditions, and manipulate features of social structure across entire populations. For these reasons, they are ideal for showing the logical plausibility and necessary assumptions of a theory. It is inappropriate to analyze agent-based models using inferential statistics. Interpreting agent-based modeling results using inferential statistics is inappropriate because there is no distinction between the sample and population in an agent-based model. In an agent-based model, the sample *is* the population and can be arbitrarily large (Jackson et al., 2017). For this reason, it is better to interpret model results using illustrative visualizations and by demonstrating consistency across simulation runs.

networks. Agents in these networks engaged in partner choice dilemmas in which they needed to predict their partner's likelihood of cooperating in a given context using inferences about moral character. Agents' initial strategies for predicting cooperation varied randomly and continuously from exclusively localized (predictions were completely based on partner behavior in a specific context) to exclusively generalized (predictions were completely context-independent), but agents could update their strategies over time through payoff-biased transmission (Kendal et al., 2009). This meant that functional partner choice strategies would evolve to be prevalent.

We hypothesized that generalized morality would evolve to be prevalent in large social networks because partners are more anonymous in these networks. We tested this hypothesis in two ways using two versions of our model. In the first version, we manipulated the size and connectivity of the social network across runs of the model. We focused on size and connectivity because people will have more social ties in larger (vs. smaller) and more connected (vs. less connected) networks and should be more unfamiliar with their partners, on average, in these kinds of networks. Our *a priori* hypothesis was that generalized morality might therefore be more prevalent in large and highly connected networks.

In our second version of the model, we manipulated partner anonymity directly: Readers can think of this as a computational version of an experiment that manipulates a mediator. We introduced "social memory shocks," in which agents "forgot" information about individuals in their social network. The advantage of this model was that it allowed us to test whether a sudden rise in partner unfamiliarity would increase the prevalence of generalized morality with strong causal inference, but the disadvantage of this model was that it was unrealistic—outside of science fiction, entire societies do not have their social memories wiped.

Method

For symbolically minded readers, we have included a more formal description of the model dynamics with equations, tables, and visualizations in the Supplemental Materials. We present a plain text version here, which is accessible for a broader readership.

Plain Text Description of Model Dynamics. Imagine that you are dropped into a community of strangers. In this community, as in real life, you must navigate social dilemmas that involve cooperation. For example, you may be traveling for a weekend and looking for a responsible babysitter or hoping to find an honest investment partner with whom to create a business. These dilemmas will involve a wide variety of people in a wide variety of situations. Your objective as you navigate this community is to form collaborative bonds with people who will reciprocate cooperation and to avoid placing trust in people who will try to exploit you.

You initially have no insight into who will be more or less likely to cooperate in these sorts of situations, but over time you form impressions of each person's moral character using previous experiences. As you become more familiar with your partners, you can use your impressions of moral character to predict someone's future likelihood of cooperating in a given situation and calibrate your trust in that person accordingly. For example, you might choose to create a business with someone whom you strongly believe will be an honest partner. On the other hand, you might

shorten your weekend trip to an overnight stay if you have doubts about your babysitter's level of responsibility. One complicating factor is that people are more likely to cooperate in some situations than others (e.g., your honest business partner might be an irresponsible babysitter).

When you predict someone's likelihood of cooperation in a given situation, you must balance two sources of information about moral character: your general impression of that person's moral character from the sum of your previous interactions (generalized morality) and information about that person within the context at hand (localized morality). There is no obvious way to decide which of these factors is more important, but you can infer which might be more important by watching how successful people in your social network choose their partners. Some of your peers might emphasize moral character as generalized across all situations, whereas others may tend to emphasize a more contextually nuanced perception of morality. By copying these successful individuals, your strategy for predicting cooperation might evolve over time. The central question for our model is: "How will your strategy evolve?"

Model Versions. Model 1 was a "cross-sectional" model with 200 rounds in which we manipulated parameters across runs. We manipulated the size and connectivity of the social network and the correlation of cooperativeness across situations. Supplemental Table S6 summarizes the breakdown of these parameters across simulation runs. We predicted that generalized morality would be initially adaptive in all models, and that it would be replaced by localized morality as agents became more familiar with their partners. However, we predicted that generalized morality would be adaptive for longer in larger networks and denser (i.e., more connected) networks since it would take longer to become familiar with connected partners in these networks.

Manipulating the covariance α between cooperation in different contexts provided an additional robustness test for our hypothesis tests. As this coefficient approaches 1.00, generalized morality develops equivalent predictive power to localized morality because cooperativeness does not vary across situations. However, as the coefficient approaches 0.00, generalized morality becomes equivalent to randomly guessing cooperativeness because cooperativeness in one situation does not relate meaningfully to cooperativeness in other situations. We were therefore interested in testing whether our predictions would hold if α was moderate, or whether they would only hold given the assumption that cooperativeness correlates highly across situations (with a value of α approaching 1.00).

Model 2 was a "longitudinal" model in which we manipulated a key parameter across time within runs. We defined a fixed network size (10), network connectivity (0.20), and cross-context correlation in cooperation (0.40) for all runs. These runs also had 1,000 runs each, making them longer than the runs in Model 1, so that we could better illustrate changes over time. We then introduced an "anonymity shock" every 200 rounds of each simulation run, in which agents lost all previously learned information about their partners. In other words, C was reset to a matrix wholly composed of undefined values. This formulation allowed us to test whether generalized morality increased following anonymity shocks and then decreased between anonymity shocks as agents rerecognized information about their partners' behavior. In other words, we could manipulate our theorized mediator (partner familiarity) of the relationship between social network size and generalized morality.

Results

Cross-Sectional Findings. We first analyzed the cross-sectional model in which we manipulated social network size, connectivity, and cross-situation covariance in cooperativeness across simulation runs.

We began by visualizing how generalized morality changed throughout the model based on variance in social network size. Figure 5 illustrates these dynamics. In this figure, line width captures variation in the results across runs, meaning that thicker lines illustrate greater between-run noise than thinner lines. The results are also broken down by cross-situation covariance in cooperativeness. The left panel of Figure 5 shows that generalized morality quickly declined in favor of localized morality when cooperativeness did not covary at all across situations. Regardless of social network size, agents adopted a more situation-specific approach to predicting cooperation as they became more familiar with their partners. This is not surprising. When cooperation does not correlate from one situation to another, it is not diagnostic to use previous cooperation in one situation to predict future cooperation in a different situation.

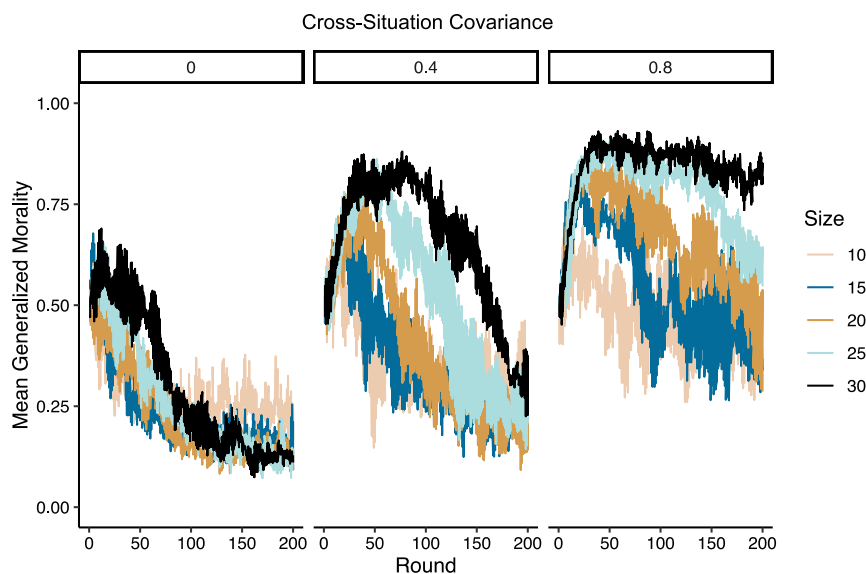
The middle panel of Figure 5 shows that, when cooperativeness covaries moderately across situations, generalized morality rises in prevalence in the early stages of the model—when agents are mostly unfamiliar with one another—and then declines in prevalence as agents become more familiar. This decline is most rapid in small social networks and slowest in large social networks. At Rounds 100 and 150, generalized morality was substantially more prevalent in larger versus smaller social networks in all runs of the model, although this variation had mostly collapsed by Round 200. Finally,

the right panel of Figure 5 shows that when cooperativeness covaries highly across situations, generalized morality grows to be more prevalent in larger (vs. smaller) social networks and remains prevalent in large social networks over time in all runs of the model. This effect was most pronounced in the largest social networks. In large social networks where agents behaved similarly across situations, generalized morality persevered as an adaptive strategy for predicting cooperation, even when agents became more familiar with one another.

Running the same analyses for connectivity revealed substantially noisier results, as illustrated by the thicker lines in Figure 6 versus Figure 5. This is probably because greater connectivity did not only increase anonymity in social networks; it also made the model dynamics more stochastic. Since agents in highly connected networks simultaneously learned from all their social ties (see Equation 2 in the Supplemental Materials), slight shifts in agent behavior could affect the entire population's behavior. In contrast, large social networks with relatively low connectivity (the black time series in Figure 5) are more stable because the network is more modular, meaning that random shifts in one agent's behavior should only affect a small number of connected agents during the social learning stage. These large and modular social networks also better resemble real-world social networks; completely connected social networks are extremely rare in the real world (D. J. Watts & Strogatz, 1998), and so we assign more weight to the network size manipulations when interpreting our results.

In sum, our cross-sectional models showed that social network size—which increases unfamiliarity in social networks—facilitated generalized morality. We next analyzed whether we could produce the same effects by directly manipulating anonymity.

Figure 5
Model Results by Social Network Size and Cross-Situation Covariance in Cooperativeness



Note. Line width captures variability in the results across runs, such that thicker lines illustrate greater variability across runs than thinner lines. Colors represent social network size across runs (these networks had two dimensions, so the number of agents is equal to size squared), and panels represent cross-situation covariance in cooperativeness across runs. Each run contained 200 rounds, which are plotted on the x -axis. The y -axis plots mean generalized morality (β_1) across all agents in the run over time. See the online article for the color version of this figure.

Longitudinal Findings. We next analyzed the longitudinal model, which contained social memory shocks. Figure 7 depicts the results of this model across each of the 10 runs (bottom pane) and aggregated across the runs (top panel). The red time series in Figure 7 represent the mean anonymity—measured as the percent of defined values in agents’ prediction matrices. There are spikes in anonymity every 200 rounds, which correspond to the anonymity shocks.

The black time series in Figure 7 represents the mean generalized morality parameter (β_1) across all agents in the model. A score of 0.50 would mean that agents, on average, assign equal weight to generalized morality and localized morality when they predict cooperation in partner choice dilemmas. A score of 0.00 would mean that agents rely completely on localized morality when they predict cooperation. Neither generalized nor localized morality approach complete dominance (1.00), which suggests that it was adaptive for agents to integrate both generalized and situation-specific information when predicting cooperation. However, generalized morality rises noticeably following each anonymity shock. This is most visible in the aggregated top panel of Figure 7 but also visible in the noisier individual runs. This longitudinal finding offers causal evidence that perceptions of generalized morality are adaptive in unfamiliar social networks given moderate covariance in cooperativeness across situations.

Discussion

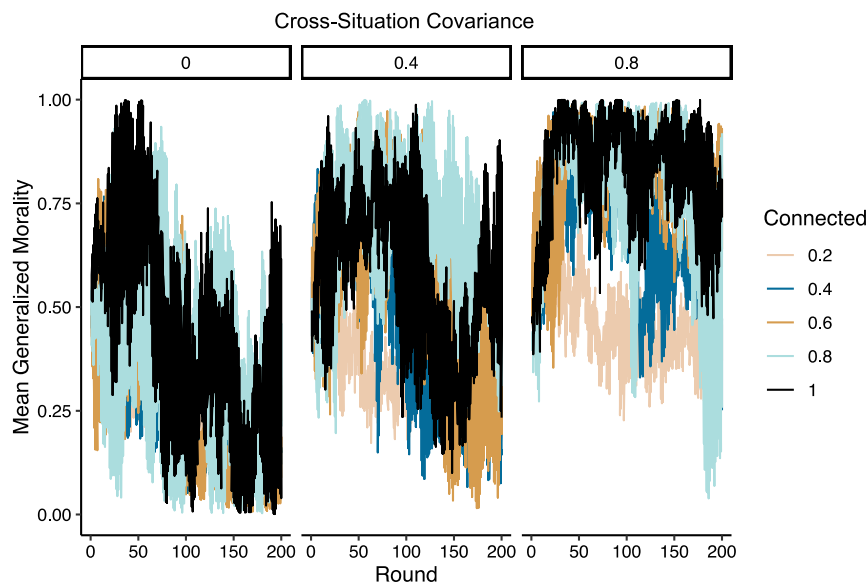
We constructed an agent-based model that supported the adaptiveness of generalized morality in large and unfamiliar social networks. In this model, agents engaged in repeated partner choice dilemmas where they had to predict the cooperation of their partners

using a combination of localized morality (relying on situation-specific information about their partners) and generalized morality (generalizing across situations to form impressions of their partners). In a modified trust game, agents who made better predictions earned more resources than agents who made worse predictions. Agents could adapt their weighting of generalized versus localized morality by observing which strategies were successful among their social ties. In this framework, we tested whether greater unfamiliarity would increase the adaptive value—and by extension, the prevalence—of generalized morality.

We supported our hypothesis in two ways. In one version of the model, we found that “anonymity shocks,” which increased unfamiliarity, led to reliable surges of generalized morality in the population of agents. In another version of the model, we found that agents interacting in larger social networks—which are characterized by greater unfamiliarity—relied more on generalized morality than agents interacting in smaller social networks. We also found some evidence that denser social networks increased reliance on generalized morality, but this evidence was far less clear than our social network size findings, largely because results were less consistent across simulation runs. In sum, the results suggest that generalized morality should culturally evolve in large social networks if agents rely on payoff-biased transmission.


We also conducted robustness tests to explore how these findings changed based on cross-situation covariance in cooperation. In some runs of the model, agents’ behavior was completely independent across situations—an individual’s likelihood of cooperating in one situation had no relationship with their likelihood of cooperating in a different situation. In other runs of the model, agents’ behavior covaried moderately (0.40) or highly (0.80) across situations. These robustness

Figure 6
Model Results by Connectivity and Cross-Situation Covariance in Cooperativeness



Note. Line width captures variability in the results across runs, such that thicker lines illustrate greater variability across runs than thinner lines. Colors represent connectivity across runs, and panels represent cross-situation covariance in cooperativeness across runs. Each run contained 200 rounds, which are plotted on the *x*-axis. The *y*-axis plots mean generalized morality (β_1) across all agents in the run over time. See the online article for the color version of this figure.

Figure 7
Results of Longitudinal Model Formulation



Note. The bottom panel displays results of 10 individual runs. The top panel aggregates the results of the individual runs. The red time series is mean anonymity, measured as the percent of defined values in agents' prediction matrices. The black time series is the mean generalized morality parameter (β_1) across all agents. See the online article for the color version of this figure.

tests showed that generalized morality is generally maladaptive when cooperativeness does not covary across situations and generally adaptive—especially in large social networks—when cooperativeness covaries highly across situations. When cross-situation covariance is moderate, generalized morality is adaptive, but only when partners are unfamiliar with each other. Building on this finding, our next study experimentally tested whether generalized morality is adaptive in partner choice dilemmas involving highly unfamiliar partners.

Study 4: Experimental Test

Our agent-based model holistically tested whether generalized morality is functional in large social networks, and social networks composed of mostly unfamiliar partners. But what kinds of interactions specifically make generalized morality valuable when partners are unfamiliar? Study 4 tested the possibility that generalized morality is most valuable when people have some prior information about an interaction partner but do not have information about a partner's cooperativeness in a particular context. In these cases, perceptions of generalized morality may help people to fill in the gaps and make an educated guess at cooperation.

For example, someone might decide to trust a financial advisor with their savings because they heard she regularly volunteers at a soup kitchen. Although soup kitchen volunteering will not always translate to financial ethics, it provides a rough approximation of whether someone will invest your money responsibly. On the other hand, generalized morality might be maladaptive when a partner's previous behavior in some situation is known (soup kitchen volunteering should be discarded when you know that a financial advisor misleads investors to prop up her own fund).

We tested this hypothesis by asking participants to complete a series of partner choice dilemmas in a hypothetical community in which they either had access to situation-specific information about cooperativeness (familiar condition) or did not have situation-specific information about cooperativeness (unfamiliar condition). We also manipulated their conception of morality to be localized or generalized. We then measured participants' ability to predict cooperation. We hypothesized that generalized morality would improve predicting cooperation in the unfamiliar condition but would impair predicting cooperation in the familiar condition. We also replicated these results with a preregistered study, which we present in the supplemental materials as Supplemental Study 4.

Method

Sample. We advertised the study for 1,000 participants on Amazon Mechanical Turk through CloudResearch. Only 795 participants signed up before our 24-hr recruitment window expired, and only 692 participants (326 men, 362 women, four nonbinary; $M_{\text{age}} = 40.44$, $SD_{\text{age}} = 12.24$) completed all the measures and passed the attention check (participants were instructed to list “gardening” from a series of hobbies if they were paying attention). We suspect that this low response rate may have been due to the fact that the study was quite long (~15 min), and participants may generally prefer to participate in shorter studies.

Paradigm. Participants read that they would see a series of dilemmas within a hypothetical community where they would need to predict a partner’s likelihood of cooperating. Participants were also told that the accuracy of their predictions would determine the size of a bonus that they received after the study, with the worst possible set of predictions resulting in a bonus of \$0.00 and the best possible set of predictions resulting in a bonus of \$1.00.

Participants from both conditions then completed 14 trials in which they predicted partner cooperation. In each trial, participants were presented with a profile containing a person’s name and some information about different moral attributes rated from 0 to 100. Beside the person’s profile, participants read a cooperative dilemma involving the partner. For example, Figure 8 summarizes a dilemma where participants needed to decide whether the partner could be trusted to responsibly look after the participant’s children for several days. After reading the dilemma, participants rated their level of trust that the person would cooperate in the situation from 0 (“I would not trust them at all”) to 100 (“I would trust them a great deal”).

Several aspects of this design were chosen intentionally. We selected the seven different attributes (“caregiving ability,” “group loyalty,” “reciprocity,” “interpersonal respect,” “allocation fairness,” and “resource generosity”) based on the seven different moral intuitions summarized in the MAC model (Curry et al., 2019). We also pretested the cooperative dilemmas such that the context of each dilemma would be particularly relevant to one of the moral attributes. The dilemma described in Figure 8 was pretested as most relevant to “caregiving ability.” We provide the full list of vignettes and the moral attributes that were matched to each vignette in Supplemental Study 1.

We measured “prediction error” in this paradigm as the absolute value of the difference between participants’ 1–100 trust estimate and the partner’s real score on the attribute that was most relevant to the dilemma. For example, in a trial where the scenario was pretested to match caregiving ability, the partner had a caregiving ability of 70, and the participant indicated a trust level of 50, the participant would receive a prediction error of 20 (the absolute value of 50–70). We adapted this paradigm directly from the mechanics in our agent-based model.

We also note that we simulated the attribute scores so that they varied according to real-world behavior. In Supplemental Study 2a, we had participants select their likelihood of cooperating in each of the dilemmas, and we used the covariance in self-reported cooperativeness across dilemmas to simulate the covariance across the moral attributes in the Study 4 profiles (see Supplemental Study 2a, for more details).

Familiarity Manipulation. In the “familiar” condition, all profiles in the cooperation dilemmas had visible scores. In the “unfamiliar” condition, however, participants lacked information

about half of the moral attributes, including the attribute that was most relevant to the dilemma. For instance, in the scenario where participants needed to trust a partner to look after their children, information about the partner’s caregiving ability would be missing (see Figure 8, for an illustration). This design emulated the conditions under which generalized morality would be most functional according to our model—when agents lacked information about their partner’s likelihood of cooperating in a given situation but could develop an approximate prior using information about the partner’s cooperation in other situations.

Generalized Morality Manipulation and Measurement. In the “generalized morality” condition, participants were told that “people in this community have a fundamental underlying moral character, which determines how they behave in a range of situations in everyday life. People’s individual traits give you hints about what their moral character might be.” In the “localized morality” condition, participants were told that “people in this community have no single underlying moral character, which means that they may be more cooperative in some situations than others.” This manipulation was intended to manipulate whether participants used a generalized conception of morality (aggregating all of the different moral attributes into a single “good–bad” view of each partner) or represented their partners’ attributes as separate predictors of cooperation in different situations.

In addition to this manipulation, we also measured people’s self-reported strategy for estimating cooperation. After participants completed these dilemmas, they responded to an item asking: “How did you determine whether to cooperate with people?” Participants could choose either: (a) “I tried to get a general impression of the person’s moral character by averaging across their traits” (our measure of generalized morality), (b) “I tried to remember specific traits so I could get a sense of people’s behavior in different contexts” (our measure of localized morality), or (c) “I tried a different strategy (please specify).” We excluded 30 participants who tried a different strategy since we could not be sure whether they favored a more localized approach or generalized approach to predicting cooperation.

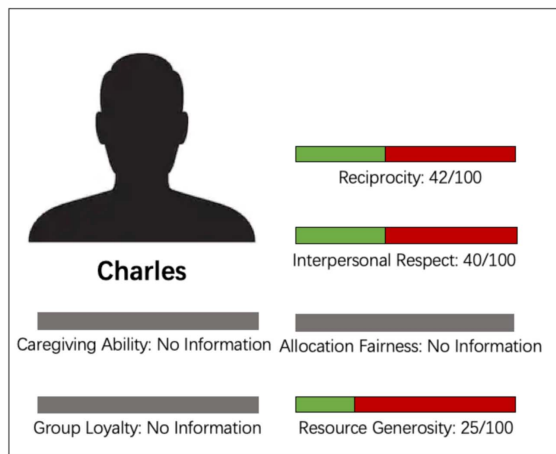
Analytic Plan. Our first analysis was a manipulation check in which we fit a generalized linear model using logistical regression to estimate whether participants in the generalized morality condition reported using the generalized morality strategy. We next fit two multilevel models with observations nested in participants to test whether familiarity interacted with generalized morality to predict prediction error. In one of these models, we interacted familiarity with the manipulation of generalized morality. In the other model, we interacted familiarity with participants’ self-reported reliance on generalized morality.

Results

Note that, when interpreting models of prediction error, higher (i.e., more positive) estimates represent *greater* prediction error and therefore *worse* performance.

Manipulation Check. Participants in the generalized morality condition were more likely to report using generalized morality (42.55%) than participants in the localized morality condition (27.93%), confirming that our manipulation was successful, $b = 0.65$, $OR = 1.91$, $SE = 0.17$, $t = 3.92$, $p < .001$, 95% CIs [0.99, 0.31].

Figure 8
A Cooperation Dilemma in the “Unfamiliar” Condition of Study 4



Consider the following scenario:

You are leaving for a trip that will last several days. You have two small children, and you want them to be safe, but you don't have any close family members who will take care of them while you are gone. You have an arrangement with your neighbor to take care of each other's children when you are out of town, but you always wonder whether they are doing their part responsibly.

Based on this person's attributes (displayed below), how much would you trust them in this scenario?



Note. The profile is lacking information about the key attribute that participants need to estimate their partner's likelihood of cooperating in the dilemma. See the online article for the color version of this figure.

Generalized Morality Manipulation and Familiarity Condition. As predicted, there was a significant interaction between the familiarity condition and the generalized morality condition, $b = -2.62$, $SE = 0.89$, $\beta = -0.04$, $t(688.14) = -2.95$, $p = .003$, 95% CIs $[-4.36, -0.89]$, such that participants using the generalized morality strategy made significantly better cooperation predictions when they were in the unfamiliarity condition, $b = -1.82$, $SE = 0.63$, $\beta = -0.06$, $t = -2.89$, $p = .004$, 95% CIs $[-3.05, -0.59]$, but similar cooperation predictions in the familiarity condition, $b = 0.81$, $SE = 0.62$, $\beta = 0.02$, $t = 1.28$, $p = .20$, 95% CIs $[-0.43, 2.04]$ (see Figure 9). In other words, participants in the generalized morality condition were better than participants in the localized morality condition at predicting partner cooperation when they did not have access to their partner's situation-relevant moral attribute, presumably because they used less relevant moral attributes to inform their predictions.

Generalized Morality Measurement and Familiarity Condition. We found similar results using participants' self-reports of generalized moral judgment as we had found using the manipulation. There was a significant interaction between generalized moral judgment and familiarity on prediction error, $b = -3.32$, $SE = 0.93$, $\beta = -0.10$, $t = -3.60$, $p < .001$, 95% CIs $[-5.12, -1.51]$. However, unlike the model where we included condition, both simple slopes were statistically significant in the model where we included self-reports. Participants who self-reported using generalized moral judgment to predict cooperation performed significantly worse in the familiar condition compared to participants who self-reported using complex moral judgment, $b = 2.04$, $SE = 0.66$, $\beta = 0.06$, $t = 3.08$, $p = .002$, 95% CIs $[0.74, 3.34]$, but significantly better in the unfamiliar condition, $b = -1.28$, $SE = 0.64$, $\beta = 0.04$, $t = -2.00$, $p = .047$, 95% CIs $[-2.53, -0.02]$. These effects, illustrated in Figure 9, are meaningful because they simultaneously illustrate (a) the cost of using generalized moral judgment when one can rely on a localized

attribute that is more relevant to a specific cooperation context and (b) the benefit of using generalized moral judgment when these context-specific attributes are not available.

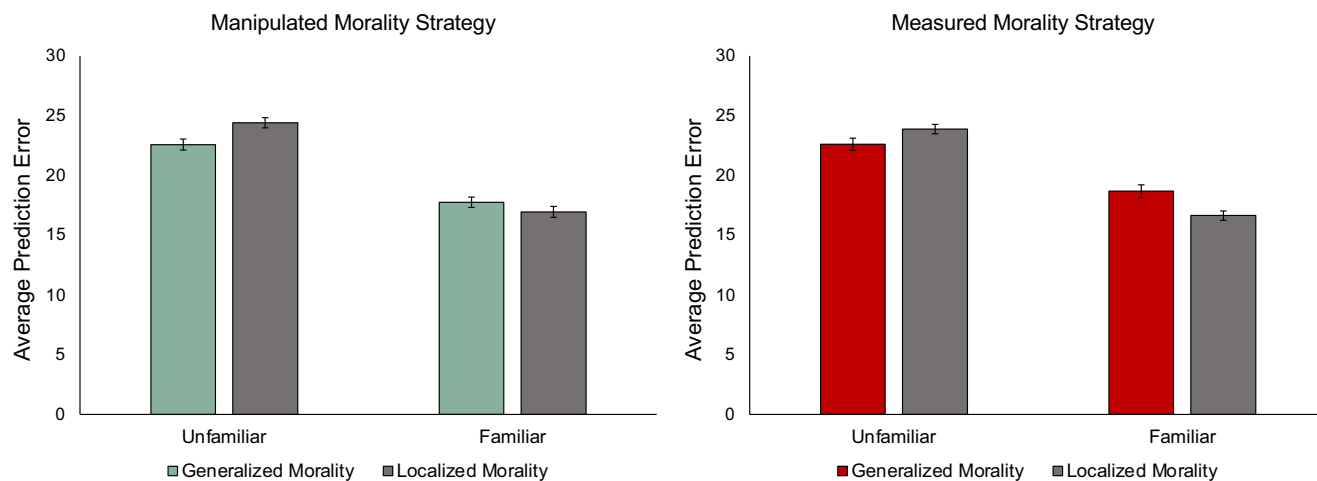
Discussion

In Study 4, we used an experiment with human subjects to show that generalized morality is adaptive when people lack information about the likelihood of partner cooperation in a particular situation but have information about a partner's previous cooperativeness in other situations. We also found that using generalized morality can be disadvantageous when people have information about a partner's situation-relevant moral attributes. We replicate this interaction in Supplemental Study 4, which is a preregistered replication of Study 4.

One strength of Study 4 was that it allowed us to test a critical finding of our agent-based model with human subjects. In our model, we assumed that perceiving generalized morality would encourage people to recruit information about cooperativeness across situations and found that this strategy improved agents' ability to predict cooperation in unfamiliar contexts but hindered their ability to predict cooperation in familiar contexts. Here we provide evidence supporting that assumption and the associated result in a paradigm that closely mirrored the structure of our agent-based model.

One limitation of Study 4 was that we did not truly manipulate participants' beliefs about the nature of morality. We instead encouraged them to think about morality as generalized or localized in a particular hypothetical community. People's assumptions about morality are deep-seated, and we reasoned that trying to manipulate these beliefs would be difficult, if not impossible. Furthermore, Studies 1 and 2 compensate for this limitation by providing more ecologically valid measures of generalized morality and showing

Figure 9
Results of Study 4



Note. Error bars represent standard error. See the online article for the color version of this figure.

that these measures reliably correlate with social network size and unfamiliarity in the real world.

Hypothesis 3: Generalized Morality Has Become More Prevalent Over Human History

Our final study examined the historical dimension of our theory, testing our hypothesis that perceptions of morality have become increasingly generalized over history. We apply a method known as word embeddings, in which words are assigned to high-dimensional vectors and the geometry of the vectors captures semantic relations between words. Past research has applied word embeddings to study how some words change in their meaning faster than others (Hamilton et al., 2018), and how stereotypes change based on the strength of semantic associations involving social identities (e.g., “woman”—“warm”; Charlesworth et al., 2021). We applied word embeddings to study whether words representing different moral values have historically become more semantically interchangeable.

We predicted that words from different moral foundations (e.g., “fair” vs. “loyal”) would have faster rates of semantic convergence than words from the same moral foundation (e.g., “fair,” “impartial”) from 1800 to 1999, since this pattern would be consistent with the breakdown of moral specificity and the rise of generalized morality for English-language speakers. Our study period (1800–1999) only captures the last 200 years of human history, but several accelerants of human social network size—such as urbanization, relational mobility, and residential mobility—nevertheless rose in the Western world during this time (Henrich, 2020).

Study 5: Historical Word Embeddings Analysis

Method

Primer on Word Embeddings. Word embeddings refer to a set of techniques that model semantic space based on how words co-occur in text. In the process of training a shallow neural network model to predict word co-occurrence, words are gradually mapped

to vectors (embedded), which represent their coordinates in a high-dimensional semantic space. These are commonly referred to as word embedding models and come in several forms (e.g., “common bag of words [CBOW]” and “skip-gram with negative sampling [SGNS]” models, which vary in the specific neural architecture used to predict word cooccurrence). However, across all models, words are always embedded in multidimensional spaces and can thus be interpreted in a similar way (Mikolov et al., 2013).

The logic of word embeddings models rests on the well-supported distributional hypothesis from linguistics—that words which occur in the same contexts will have similar meanings (Harris, 1954). That is, words that are closer together in a multidimensional embedding space have more similar meaning than words that are further apart. The distance between these words is often quantified through cosine similarity (Turney & Pantel, 2010), which is bounded between 1 (perfect semantic redundancy) and -1 (perfect semantic divergence).

Sample and Data Set. We sampled our data from the Google Ngram corpus, which contains more than 150 billion words from 1800 to 1999. The Google Ngram corpus is the most reliable of the Google Books corpora because it is the largest, and corpus size is correlated with the accuracy of word embeddings models. One justified critique of the Google Ngram corpus is that it has changed in content over time due to the influx of scientific literature and the rise of self-publishing (Pechenick et al., 2015). For this reason, some studies prefer the smaller Google English Fiction corpus since the content type is comparatively stable over time.

We tested our hypotheses in both corpora and found the same result. We present our analyses of the Google Ngram corpus in the main text because we view the English Fiction corpus as less reliable due to extremely sparse word vectors throughout much of the 19th century, but we nevertheless summarize our results using the English Fiction corpus in the Supplemental Materials. We only used embeddings models up until 2000 CE, which is common practice since content changes accelerated in the 21st century (Charlesworth et al., 2022; Grossmann & Varnum, 2015).

We used word embeddings models that were pretrained on the Google Ngram corpus by Hamilton et al. (2018). These “Histword”

models are among the most commonly used English-language embeddings, and they have been used to study stereotypes in past research (Charlesworth et al., 2022; Garg et al., 2018). Following these past studies, we used the SGNS (word2vec) embeddings because they are best at accommodating sparse data sets in which many words are only rarely used and are best suited for comparisons across decades (Hamilton et al., 2018). The models were trained separately for each decade using several different approaches, and their corresponding vector spaces were aligned to facilitate comparisons across time points. Indeed, one of the main innovations of the Histword models was that they were specifically designed to capture historical semantic change.

Stimuli. To capture moral language, we used words from the Linguistic Inquiry and Word Count moral foundations dictionary (Graham et al., 2009). The dictionary is highly heterogeneous in terms of word length and part of speech, and some words are highly specific to modern Western culture (e.g., the word “bourgeoisie” is included in the “Authority virtue” subdictionary). To create a sample of words that was more linguistically balanced and culturally generalizable, we extracted 12 positive adjectives and 12 negative adjectives from the dictionary. We selected two positive and two negative words for each moral foundation, with a further two words to represent generalized moral goodness and generalized moral badness. All stimuli are listed in Table 4.

Analytic Strategy. We began by extracting the word vectors for each of our stimuli. All but one word (“Allegiant”) appeared in the Histwords embeddings for at least two decades. “Allegiant” did not appear in any model, most likely because it was not used frequently enough in text. Figure 10 displays these word vectors in a condensed two-dimensional space, which we determined by applying a t-distributed stochastic neighbor embedding (tSNE) algorithm to the 300 dimensions in the original Histwords data set.

Table 4
Stimuli Used in Study 5

Word	Valence	Foundation
Compassionate	Positive	Care
Generous	Positive	Care
Fair	Positive	Fairness
Impartial	Positive	Fairness
Loyal	Positive	Loyalty
Allegiant	Positive	Loyalty
Honorable	Positive	Authority
Respectful	Positive	Authority
Wholesome	Positive	Purity
Dignified	Positive	Purity
Moral	Positive	Generalized
Good	Positive	Generalized
Cruel	Negative	Care
Malevolent	Negative	Care
Dishonest	Negative	Fairness
Fraudulent	Negative	Fairness
Unfaithful	Negative	Loyalty
Treacherous	Negative	Loyalty
Dishonorable	Negative	Authority
Disrespectful	Negative	Authority
Degrading	Negative	Purity
Disgusting	Negative	Purity
Immoral	Negative	Generalized
Bad	Negative	Generalized

After extracting the word vectors, we calculated the pairwise cosine similarity between the moral words in each pretrained word embeddings model, which resulted in a series of 24×24 matrices for each historical decade. We then melted these pairwise similarities into single vectors and then bound these vectors into a data set that contained variables representing (a) the decade, (b) word m in a pairwise comparison, (c) word n in a pairwise comparison, and (c) the cosine similarity between m and n . We analyzed this data set after removing duplicate comparisons (e.g., “good–bad” and “bad–good”). For the main analysis, we also removed cosine similarity values of 0 since these represent a pairwise comparison that did not have sufficient decade-specific data to generate a cosine similarity value. Key results replicate regardless of whether or not we remove these values.

Our central hypothesis was that moral words would semantically converge over time, and this convergence would be most rapid for words from different moral foundations compared to words from the same moral foundation. We tested this hypothesis by first fitting the relationship between decade and cosine similarity to estimate general semantic convergence and then interacting decade with a dichotomous “foundation match” variable that represented whether two moral words were from the same moral foundation (coded as 1) or different moral foundations (coded as 0). Since the data points in these analyses were not independent (e.g., the same words appear multiple times in the word m and word n columns of the data set), we fit multilevel models with intercepts and slopes randomly varying across words. We also controlled for whether comparisons were between words of the same positive or negative valence (a dichotomous “valence match” variable) and whether at least one of the words was positive (a dichotomous “positive” variable) to account for the fact that negative words may have greater semantic drift than positive words (Jackson et al., 2023).


Results

How have moral semantics changed over recent history? Our initial model (see Table 5) found that decade was significantly and positively related to cosine similarity, such that moral words have become more semantically interchangeable over time, $b = 0.01$, $SE = 0.001$, $\beta = 0.09$, $t = 7.90$, $p < .001$, 95% CIs [0.007, 0.01]. This initial finding is interesting. However, it should be taken with caution because the 20th-century word embeddings models are less sparse than the 19th-century models, and sparsity may decrease the accuracy of cosine similarity estimates. The SGNS approach and the Histwords models more generally were trained specifically to enable comparisons across decades, but the sparsity of the 19th-century models could nevertheless bias estimates. Therefore, our key test was whether the rate of semantic convergence would vary across words from the same versus different moral foundations. This test was important because it held the overall rate of semantic convergence constant and examined the variability in this rate across different pairs of words.

This follow-up model (see Table 5, Model 2) found a significant interaction between decade and foundation match. There has been no significant semantic convergence among words from the same moral foundation, $b = 0.0002$, $SE = 0.004$, $\beta = 0.01$, $t = 0.07$, $p = .95$, 95% CIs [−0.006, 0.006], but a robust semantic convergence across words from different moral foundations, $b = 0.01$, $SE = 0.001$, $\beta = 0.10$, $t = 8.37$, $p < .001$, 95% CIs [0.008, 0.01]. In the

Figure 10

The Semantic Space of Our Moral Stimuli in the 1800–1810 Model and the 1990–2000 NGram Model



Note. Points in space were derived from a tSNE dimension reduction algorithm applied to the 300-dimensional vectors, which are compressed to a two-dimensional space. Nodes of the same color connote words belonging to the same foundation, according to the LIWC dictionary. tSNE = t-distributed stochastic neighbor embedding. LIWC = Linguistic Inquiry and Word Count. See the online article for the color version of this figure.

1800–1810 data set, words from the same moral foundations had far more semantic similarity than words from different foundations, $b = 0.10$, $SE = 0.02$, $\beta = 0.32$, $t = 5.67$, $p < .001$, 95% CIs [0.07, 0.14]. However, this gap had shrunk considerably by the 1990–2000 data set, $b = 0.07$, $SE = 0.02$, $\beta = 0.22$, $t = 3.76$, $p = .001$, 95% CIs [0.03, 0.10]. We did not find any accompanying interaction with words from the same versus different valence. All statistics from these

models are presented in Table 5. Figure 11 illustrates the cosine similarity rates for groups of words from the same versus different moral foundations.

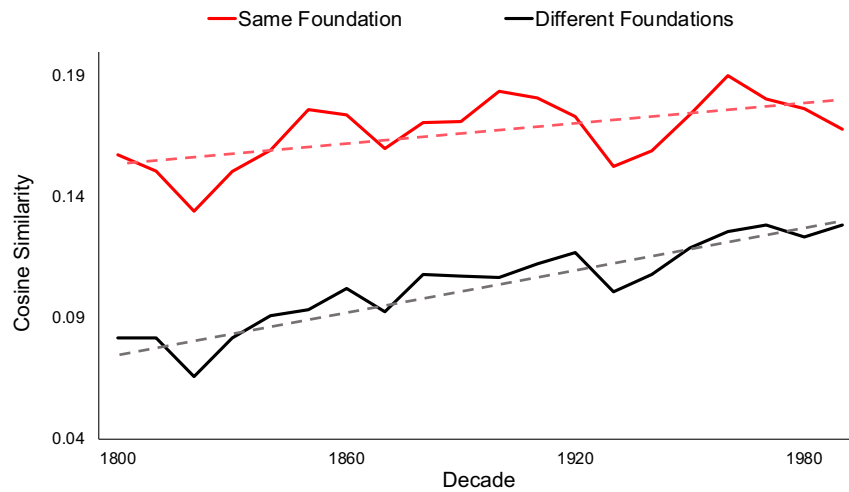
One alternative explanation for this result was that there was a ceiling effect of cosine similarity for moral attributes within the same moral foundation (e.g., “fair” vs. “impartial”). However, this possibility seems highly unlikely because the average cosine

Table 5
Models Predicting Historical Cosine Similarity in Study 5

Model variable	<i>N</i>	<i>b</i> (<i>SE</i>)	<i>t</i>	<i>p</i>	95% CIs
Model 1	4,341				
Decade		0.01 (0.001)	7.90	<.001	[0.007, 0.01]
Positive		−0.04 (0.02)	−1.84	.08	[−0.08, 0.005]
Affect match		0.07 (0.01)	5.24	<.001	[0.04, 0.10]
Foundation match		0.08 (0.02)	4.94	<.001	[0.05, 0.12]
Model 2	4,341				
Decade		0.01 (0.001)	7.83	<.001	[0.007, 0.01]
Positive		−0.04 (0.02)	−1.91	.07	[−0.08, 0.04]
Affect match		0.07 (0.01)	5.23	<.001	[0.04, 0.10]
Foundation match		0.09 (0.02)	5.00	<.001	[0.05, 0.12]
Decade × Positive		0.02 (0.004)	5.23	<.001	[0.01, 0.03]
Decade × Affect Match		0.0008 (0.003)	0.29	.77	[−0.005, 0.007]
Decade × Foundation Match		−0.02 (0.004)	−2.94	.003	[−0.02, −0.003]

Note. The *N* term represents the number of pairwise comparisons included in the analysis. *SE* = standard error; *CI* = confidence interval.

Figure 11
Moral Semantics Over 200 Years of History



Note. Differential rates of semantic convergence for attributes representing the same versus different “foundations.” See the online article for the color version of this figure.

similarity for words within the same foundation was 0.14, which was significantly higher than words from different foundations (0.07, $p < .001$), but far from the ceiling value of 1.00. It is not uncommon to see cosine similarity values approaching 1.00 for synonyms.

In our Supplemental Materials, we present another analysis that shows that this form of semantic convergence is at least somewhat unique to words about morality. In this analysis, we examine the historical semantics of emotion words (e.g., “rage,” “fury”), which are either theorized to represent the same prototypical emotion (anger) versus emotion words (e.g., “rage,” “terror”), which are theorized to represent different prototypical emotions (anger vs. fear; Shaver et al., 1987). In this sample of emotion words, we do not see a corresponding pattern of semantic convergence, which is greater for emotion words from different prototypical emotions versus the same prototypical emotions. In other words, not all aspects of human experience are growing more semantically generalized. We revisit this idea in the general discussion.

Discussion

Study 5 supported the historical rise of generalized morality. An analysis of English-language moral attributes from 1800 to 1999 found that words representing different moral attributes have become more semantically interchangeable throughout history. The boundaries between different attributes such as “fair” and “loyal” were clearer in 1800 than they were in 1999. The high covariance of moral attributes in recent history resembles how Hadza hunter-gatherers with more external exposure perceived “honest” and “generous” to be highly correlated when they ranked their partners.

The main limitation of Study 5 is that we cannot measure the cause of change in moral semantics. According to our theory, perceptions of morality are becoming more generalized, at least partly because people’s social networks are growing larger and more anonymous, but we cannot be sure of this from analyzing the Google

Books corpus. Perceptions of morality could be growing increasingly generalized for any number of other reasons. This limitation is partly offset by our cross-sectional studies (Studies 1 and 2), which show that social network characteristics are directly linked with perceptions of generalized morality, even controlling for sociodemographic variables. However, we encourage future research that analyzes trends in moral language within a more specific time period and cultural context in which it is possible to test whether social network size coevolves with generalized morality. Using data sets of language on social media websites like Reddit and Twitter may be a good context for this kind of research since the size of social communities (e.g., subreddits) can easily be quantified.

General Discussion

The word “moral” has taken a strange journey over the last several centuries. The word did not yet exist when Plato and Aristotle composed their theories of virtue. It was only when Cicero translated Aristotle’s *Nicomachean Ethics* that he coined the term “moralis” as the Latin translation of Aristotle’s “*ēthikós*” (Online Etymology Dictionary, n.d.). It is an ironic slight to Aristotle—who favored concrete particulars in lieu of abstract forms—that the word has become increasingly abstract and all-encompassing throughout its lexical evolution, with a meaning that now approaches Plato’s “form of the good.”

We doubt that this semantic drift is a coincidence. It may instead signify a cultural evolutionary shift in people’s perceptions of moral character as increasingly generalized as people inhabit increasingly larger and more unfamiliar social networks. Here, we support this perspective with five studies. Studies 1 and 2 find that social network size correlates with the prevalence of generalized morality. Studies 1a and b explicitly tie beliefs in generalized morality to social network size with large surveys. Study 2 conceptually replicates this finding in a Hadza hunter-gatherer camp, showing that Hadza hunter-gatherers with more external exposure perceive their

campmates using more generalized morality. Studies 3 and 4 show that generalized morality can be adaptive for predicting cooperation in large and unfamiliar networks. Study 3 is an agent-based model which shows that, given plausible assumptions, generalized morality becomes increasingly valuable as social networks grow larger and less familiar. Study 4 is an experiment that shows that generalized morality is particularly valuable when people interact with unfamiliar partners in novel situations. Finally, Study 5 shows that generalized morality has risen over English-language history, such that words for moral attributes (e.g., fair, loyal, caring) have become more semantically generalizable over the last 200 years of human history.

Our Supplemental Materials summarize five additional studies that support key assumptions of our theory. Supplemental Study 1 explores the meaning of different moral attributes, including “morality.” In this study, we find that morality has a highly generalized meaning—people see “morality” as diagnostic of cooperation in virtually any situation. This study also provides a pretest that allows us to match moral attributes to the cooperation dilemmas that we use in Study 4. Supplemental Studies 2a and b show that cooperation varies across situations but not without limit. There is moderate correlation between cooperation in one situation and cooperation in other situations, which is an important assumption of our model. Supplemental Study 2a also gives us the intraclass coefficient value that we use to seed the moral profile scores in Study 4. Supplemental Study 3 provides evidence that current-day Americans intuitively view attributes like “responsible” and “compassionate” as semantically interchangeable, which is consistent with generalized morality in judgments of moral character. Supplemental Study 4 is a direct replication of Study 4. And Supplemental Study 5 is an experiment providing a complementary mechanism for why people may adopt generalized morality in large social networks. The Supplemental Materials also provide additional analyses supporting the conclusions of Studies 2 and 5.

We now turn from our central hypotheses to discussing more speculative questions about our findings. First, we consider five open questions about our theory. We then turn to discussing the implications of our findings for politics, religion, and cultural change.

Five Open Questions

Do Our Findings Apply to All Forms of Social Cognition?

Our final study showed that moral semantics have become more generalized over the past 200 years, and supplemental analyses found that the same trend did not characterize emotion semantics. Why might emotion be different? Should we expect any other forms of social cognition, like warmth and competence, to have become more generalized over history?

Given the absence of research on character complexity, there is little direct evidence to answer this question. However, some recent work has found that perceptions of personality may show the opposite trend as morality, becoming more complex in large social networks, and becoming more complex over time (Alvergne et al., 2010; Saucier et al., 2014). To explain these results, Smaldino et al. (2019) have pointed out that humans living in large complex groups occupy many social niches, providing a wider scope of behavior or habits than humans living in small-scale societies. For example, in a large industrialized society like the United States, it is possible to be

a filmmaker who hosts a book club on the weekends while also learning a language every morning. These activities simultaneously imply high levels of openness to experience (via filmmaking), extraversion (via inviting people over on the weekends for a book club), and conscientiousness (via getting up early to take language lessons). The plurality of niches makes it more likely that we will develop a plurality of dimensions to describe someone’s personality.

This research on personality shows that the relationship between social network size and generalized person perception may not characterize all aspects of social cognition. Sensitivity to niche diversity may also be an important moderator of this relationship. Morality has low sensitivity to niche diversity; the meaning of morality is similar in a university classroom and a mechanic’s garage. At the other extreme, other forms of social cognition such as competence have higher sensitivity to niche diversity: Competence in a university classroom will seldom translate to competence in a mechanic’s garage, and this sensitivity implies that generalized competence may not be a very useful heuristic for predicting competence across contexts in a socially complex society.

One future direction in the study of character complexity could examine how perceptions of character vary across history and culture for each of the “big two”—operationalized in various research programs as trustworthiness/warmth/communion versus dominance/competence/agency (Abele et al., 2021; Fiske, 2018; Oosterhof & Todorov, 2008). Recent research has found that attributes related to warmth or trustworthiness are more insensitive to niche diversity than attributes that signal dominance or competence (Eisenbruch & Krasnow, 2022). This work was focused on explaining why warmth matters more than competence in person perception, but it also has implications for asymmetric perceptions of warmth and competence complexity as societies grow larger and more complex. It is plausible that perceptions of warmth or trustworthiness become more generalized in large complex societies because they are relatively insensitive to niche diversity, whereas perceptions of competence or dominance may become more complex and context-dependent in large societies because they are more niche-sensitive.

How Do These Findings Relate to Other Theories of Moral Psychology?

Our introduction compares and contrasts our theory with existing moral taxonomy models. However, there are many other theories of moral psychology that intersect with this research. In the last decade, the theory of dyadic morality (TDM) and the closely related affective harm account (AHA) have been the main alternatives to moral taxonomy models (Gray et al., 2012, 2022; Schein & Gray, 2018). These theories focus on the psychological substrates of moral judgment; they identify how experiencing negative affect and perceiving harm encourage moral condemnation. However, TDM and AHA do not connect moral judgment to context-specific predictions about cooperation, nor do they explore how the structure of moral judgment can promote or inhibit cooperation in groups. Our account is therefore better suited as a cultural evolutionary theory of the structure of moral judgment, which complements the psychological and phenomenological focus of TDM and AHA.

Our focus on judgments of character rather than judgments of acts is also consistent with a recent push toward “person-centered morality” in moral psychology. Person-centered morality argues

that a major goal of an individual's moral psychology is to understand and predict people's behavior rather than formulate attitudes about abstract moral principles (Pizarro & Tannenbaum, 2012; Uhlmann et al., 2015). Person-centered morality also explores how moral judgment can vary across relational partners (Earp et al., 2021), and a fruitful direction for future research could explore how perceptions of character complexity vary based on relationship status and group membership. We may view the morality of a familiar partner like a family member along a vast number of different dimensions while viewing a faceless political opponent along a single dimension of immorality. This research would answer a call in moral psychology to integrate social identity context into research on morality (Hester & Gray, 2020; Schein, 2020). It would also have implications for intergroup conflicts where people are inherently more knowledgeable about members of their ingroup than members of their outgroup.

How Does Generalized Morality Differ From Internal Attributions?

Attributional style has long been a dominant paradigm in social psychology. Research has found that internal attributions are especially common in modern Western countries (Markus & Kitayama, 1991), for judgments of one's own behavior versus other people's behavior (Gilbert & Malone, 1995), and for judgments of close others' behavior rather than strangers' behavior (Taylor, 1981). One explanation of these cultural differences is that East Asian cultures practiced subsistence styles that encouraged interdependence (Talhelm et al., 2014; Thomson et al., 2018) and that early philosophical traditions in China and India emphasized a more contextualized and embedded view of the self (Vignoles et al., 2016).

Readers of this literature may wonder how perceiving generalized character differs from making internal attributions because there are similarities between these processes. Making an internal attribution, like making a judgment of generalized character, involves ignoring the context in which someone behaved. Conversely, making an external attribution, like perceiving localized morality, involves acknowledging that someone's behavior may vary meaningfully across contexts. We view differences in attributional style as one route to perceptions of localized morality. When situational pressures are salient, people may be less likely to assume that behavior was internally motivated and that this behavior reflects someone's broader character. There is evidence for this effect from Lammers et al. (2018), who found that people were less likely to make generalized inferences about single acts of prosocial or antisocial behavior after they were exposed to other forms of situational variability. Other research from Ji et al. (2023) has found that people are less likely to extrapolate from first impressions in collectivist cultures (where people are more likely to make external attributions of behavior) versus individualist cultures (where internal attributions are more common).

However, generalized morality should not only be a product of internal versus external attributions. Take a situation where a colleague's longstanding infidelity is leaked to their workplace. Even if everyone agrees that the colleague intended to be unfaithful—and did not act because of some strong situational pressure—they may disagree about what this means for the colleague's integrity in other domains. Some people may see infidelity and workplace integrity as tapping distinct aspects of

character—just like skill as a mechanical says little about skill as a programmer. Other people may take the news of infidelity as a clear signal that their colleague is untrustworthy and uncooperative in all domains of life. We encourage future research to further explore and dissect these differences between attributional style and character complexity.

Is Generalized Morality Always Functional?

In this article, we focused on how generalized morality may have culturally evolved through adaptive learning strategies. However, these are not the only pathways that could lead to the proliferation of generalized moral judgment in large groups. Generalized morality could also evolve because of more "content-driven" social learning, which is to say that it may spread because it becomes more psychologically appealing regardless of its function (Mesoudi & Whiten, 2008; Sperber, 1996). In big groups filled with more unfamiliar partners, it may be easier on the mind to think about social partners using generalized "good-bad" labels rather than considering all their virtues and vices.

In some cases, the moral domain could also become more generalized without any social learning. People may naturally switch to using generalized morality when they enter large social networks, or when they encounter strangers because it is more intuitive. Perceptions of generalized morality may also be a natural consequence of processing information about strangers because these interactions involve more abstract construals (Hess et al., 2018; Idson & Mischel, 2001). When you only have a vague impression of social partners, you may be more likely to process their behavior using abstract traits like "good" and "bad." As partners become more familiar, you may begin to ascribe them more specific moral traits, like "loyal Tar-heel" or "unwilling to share single-malt scotch," which are tied to the contexts where you become most familiar with the partner.

There may be other cues that lead people to switch to a more localized or generalized perception of morality without social learning. For example, people switch to using generalized morality when they receive consistent signals about someone's moral character, even without social learning (Lammers et al., 2018). We also show evidence in Supplemental Study 5 that people switch to using more generalized perceptions of morality when they make partner choices in larger (vs. smaller) groups.

However, it is important to bear in mind that evidence for this "switching" has exclusively come from large Western societies with vocabularies communicating generalized morality (e.g., "immoral," "good person"). Generalized morality is intuitive for people living in these groups. But in other societies, there are no words conveying generalized morality, and interacting with novel partners may be quite rare. Social learning may play a greater role in the spread of generalized morality in these societies.

To make this point clearer, we can draw an analogy to the idea of "individuation" in stereotyping and prejudice research (Taylor, 1981). Individuation describes how people can move from category-based stereotypes (e.g., Canadians are bad drivers) to individuated impressions (e.g., my Canadian friend is a good driver), as they become more motivated and attend more to a target. In large anonymous social networks, a similar process may play out for character complexity: We may begin by viewing someone's morality as generalized, and increasingly view it as localized.

But in small-scale societies where people encounter strangers less frequently, this process may look different because there are no expressions or norms to communicate generalized morality. People may view morality as localized from the start.

Studying these different mechanisms is a major future direction for this research program. Our thesis is that social learning is a key catalyst for generalized morality to evolve in societies when it becomes adaptive, and that other processes like construal level can predict when and who uses generalized morality in societies where it has already evolved. We provide some evidence for this thesis in this article, but future research could offer valuable nuances to our account.

What Does Our Model Suggest About the Evolution of Moral Psychology?

One of the clearest differences between our theory and past work is that we propose that moral pluralism arises through cultural evolution rather than biological evolution. Theories such as MFT and MAC accommodate cultural variation. For example, MFT suggests that cultural differences can explain why Indians and conservative Americans value authority and purity more than American liberals (Graham et al., 2009; Haidt et al., 1993). However, evolutionary theories of morality also assume a “first draft of the moral mind,” which not only gives humans the capacity to create moral categories but also an initial set of biologically based categories that inform intuitions about right and wrong (Haidt & Joseph, 2004). Moreover, evolutionary theories assume that humans biologically evolved these moral intuitions because of a consistent set of selection pressures that they faced in our species’ evolutionary history (Wright, 2010).

However, there is growing evidence against both of these assumptions, which we believe supports a cultural evolutionary model of moral psychology. The first assumption of universal moral intuitions was initially supported using responses of liberal and conservative American participants to the “Moral Foundations Questionnaire” (Graham et al., 2009). Cross-cultural tests of this survey questionnaire, however, revealed low measurement invariance across different countries, suggesting that the structure of morality was highly variable across cultures (Iurino & Saucier, 2020). A recent effort to develop a new cross-culturally sensitive “Moral Foundations Questionnaire V2” across 25 world nations also concluded that “the nomological network of moral foundations varied across cultural contexts” (Atari et al., 2022). We view cultural variation in the structure of the moral domain as neither a measurement problem nor an artifact of cultural values in “WEIRD” (Western, educated, industrialized, rich, and democratic) versus other societies (Henrich et al., 2010), but instead as indicative of cultural differences in the contexts and social networks in which people cooperate.

New evidence also challenges the idea that early humans faced a universal set of selection pressures that led our species to evolve genetically determined moral intuitions. The original narrative advanced by evolutionary theories of morality was that humans once lived in small kin-based tribes facing a similar set of problems: protection and care of children, detection of cheating, competition over finite resources, management of hierarchies, and pathogen avoidance (see Graham et al., 2013, for the direct list). These theories in turn assume that humans overcame these challenges by cultivating a biologically based desire for purity, loyalty, and obedience to authority.

This narrative comes from studies of extant hunter-gatherer groups, which social scientists once assumed were good proxies for human life throughout our evolutionary history (Lee & DeVore, 2017). Yet there is evidence now from archeology and ethnography that pre-Holocene human life was much more diverse than these theories assume and that many Pleistocene human groups, especially those in the late Pleistocene, did not face the kinds of social pressures that social psychologists often assume were omnipresent (Graeber & Wengrow, 2021; Singh & Glowacki, 2022).

One line of evidence suggests that hunter-gatherer groups were likely larger and more diverse than evolutionary psychology theories presume, and that loyalty may not have been a salient concern for many late Pleistocene hunter-gatherers. Instead of living within small-scale societies with tight kinship loyalties, cross-cultural evidence suggests that hunter-gatherers often resided with nonkin (Hill et al., 2011). Many groups in the late Pleistocene also oscillated between large sedentary settlements and smaller nomadic groups throughout the year based on seasonal climate variation. Recent studies have shown that foragers in Papua New Guinea, the Pacific Northwest, Alaska, and northern Australia seasonally shifted between large sedentary settlements and dispersed mobile groups depending on rainfall, fauna migration patterns, and temperature (Ames, 1994; Roscoe, 2006; White & Peterson, 1969). These findings match archeological evidence from the Neolithic and Paleolithic eras, which has found evidence of sedentary hunter-gatherer communities (Heinsohn, 2010; Marean, 2016; Snir et al., 2015).

Second, hunter-gatherer groups varied widely in their views of hierarchy and dominance and did not necessarily require obedience to authority. While it is impossible to access data about social structure and dominance beliefs from the Pleistocene, records from early colonial contact show wide variation in these beliefs across North American hunter-gatherer groups. Historical records of the Mi’kmaq and Wendat hunter-gatherers of current-day Northeastern Canada showed an explicit disavowal of dominance hierarchies and obedience to authority (de Lom d’Arce, 1905; Graeber & Wengrow, 2021), whereas other hunter-gatherers like the Calusa of current-day Florida had strong dominance hierarchies (Marquardt, 2014).

Finally, Pleistocene-era groups may not have been driven by the same kind of fairness concerns that people show in many nations today. Some studies have suggested that a sense of fairness is culturally widespread based on cross-cultural comparisons between India and the United States (Berman et al., 1985), but studies of small-scale societies have found much more mixed evidence for a universal concern for fairness. Henrich et al.’s (2004) analysis of behavior in economic games in small-scale societies, for example, concluded that “unlike University of California, Los Angeles students, Mapuche proposers rarely claimed that their offers were influenced by a sense of fairness” (p. 12). Hadza participants, who also make low offers in economic games, have strong food-sharing norms but do not appeal to fairness to justify these norms (Henrich et al., 2004; Stibbard-Hawkes et al., 2022). Complementary analyses from evolutionary anthropology have suggested that a concern about fairness may have coevolved with the rise of private property and authority systems across history (Graeber, 2012).

In sum, there may not have been a universal set of evolutionary pressures facing Pleistocene-era human groups, and a genetically evolved set of moral intuitions may not have been as functional in early human groups as we once assumed. The possible existence of universal moral concerns in hunter-gatherers is an open question

that requires more research. One avenue for this research could analyze vocabulary lists of world languages, following research on worldwide emotion concepts and personality concepts. Along these lines, the Hadza lexicon has no words for “fair,” “fairness,” “obedience,” “obey,” “just,” or “justice” (Miller et al., 2013). At this point, we consider it more likely that abstract moral categories like “loyalty” and “fairness” culturally evolve to help people organize and predict cooperation in large groups with high levels of relational mobility (Smith et al., 2018).

We grant that some evolutionary theories of morality do not make these same assumptions about early human history. For example, theories of social value orientation and welfare tradeoff ratios argue that moral intuitions evolved as a cognitive system to help humans successfully invest in partners who would increase their fitness (Delton & Robertson, 2016; McClintock & Allison, 1989). These ideas are both more plausible in our view and more compatible with our theory of moral character as a tool for partner selection.

Three Key Implications

Implications for Concept Creep

Recent research has identified “concept creep”: the growing moralizations of previously innocuous behavior (Haslam, 2016). Over the last several decades, a growing number of behaviors have been absorbed into the moral domain of large Western societies and are now viewed as diagnostic of moral good or evil.

Concept creep has coincided with exponential increases to the size of social networks via urbanization and the advent of the internet (Li, 2020), and there is good reason to think that these trends are linked. Expanding the moral domain to include events that are tenuously related to cooperation means that one can make cooperative inferences about a novel social partner with a growing set of information, in the same way, that a generalized moral concept allows one to make a prediction about someone’s likelihood of cooperating in one context using information about their previous cooperation in (tenuously linked) other contexts. This mechanism is closely related to the notion of “prevalence-induced concept change,” wherein the diminishing prevalence of overtly harmful behavior (e.g., murder) encourages people to expand the scope of immorality to include more borderline behavior (Levari et al., 2018).

Implications for the Evolution of Religion

Our theory may also explain patterns in the evolution of religious beliefs over history. One of the more enduring puzzles in the study of religion is the rise of moralizing high gods such as the Christian and Jewish gods. Historical analyses suggest that high gods with universal moral codes have spread in the last 12,000 years, coinciding with the Neolithic Revolution when many human groups transitioned to more sedentary agriculture-based communities rather than nomadic hunter–gatherer groups (Norenzayan et al., 2016; J. Watts et al., 2015). One reason for this co-occurrence may be because moralizing high gods helped regulate cooperation in these communities, encouraging people to cooperate by fostering the expectation that defection would be met with divine punishment (Johnson, 2016; Norenzayan & Shariff, 2008).

Another plausible mechanism, however, is that people developed more generalized conceptions of morality as societies grew larger,

and they projected these beliefs onto the content of gods’ minds (Purzycki et al., 2022). Some evidence supports this projection account. For example, people frequently project their own demographic and personality traits onto their perceptions of gods (Epley et al., 2009; Jackson et al., 2018; Purzycki, 2013), and project punitive characteristics onto perceptions of gods when they seek to punish norm violators (Caluori et al., 2020). This evidence makes it plausible that, once communities adopted a belief in generalized morality, their gods became arbiters of these moral beliefs rather than focusing on more context-specific domains of cooperation (e.g., individual gods policing food taboos, sacred places, or sexuality).

Implications for Political Polarization

Political polarization is rising in many countries, with a sharp rise over the last several decades of American history (Finkel et al., 2020; Iyengar et al., 2019). In political psychology, affective polarization has emerged as a particularly acute problem because many Americans express rising hostility and mistrust of opposing political parties in addition to disagreeing about policy (Iyengar et al., 2019). To explain the recent trend toward affective polarization, different models have pointed to increasing political segregation (Motyl et al., 2014), the rise of politically polarized media (Martin & Yurukoglu, 2017), and the growing alignment of social identities (e.g., White Evangelical Christian) and political party membership (e.g., Republican; Mason, 2018). Predictions from our theory suggest that the rise of online social networks may also be a factor in the rise of affective polarization (Brady et al., 2023).

The internet is unique because it rapidly expands social network size, involves a high density of interactions with strangers, and rewards negatively valenced information (Brady et al., 2021, 2023). The first two of these conditions make it likely that people will use a generalized view of morality to interact with people on social media, whereas the third makes it likely that people will view social media interaction partners negatively. These three conditions create a perfect recipe for affective polarization because adopting a generalized moral judgment will encourage people to make wide-ranging negative inferences about political opponents online using minimal information about these individuals. A potential relationship between generalized moral judgment and affective polarization shows how the function of generalized morality can backfire among polarized groups. In these groups, outgroup identification becomes a minimal cue of noncooperation, which may become emphasized for people who otherwise have little information about someone’s prosociality.

Limitations

We summarize the key limitations of our research in Table 6. In this table, we recapitulate several of the limitations that we describe in the discussion section of individual studies. We also describe factors that offset these limitations, and how future research could further address them.

Statement of Limitations

Each of our studies has limitations. For example, in Study 1, we show that generalized morality is associated with larger social

Table 6
Table of Limitations

Study	Limitation description	Possible offsets
1, 2, 5	Three of our studies use correlational data. This means that we cannot claim on the basis of these studies that increasing social network size will increase generalized morality.	These limitations are partly offset by Studies 3 and 4, which causally show that generalized morality is use useful than localized morality for predicting cooperation in large networks. Future research could detect pseudocausality by using time-series models to test how time-lagged rises in social network size of cities or towns predict rises in generalized morality in language (which could be measured through language in local newspapers).
1	The self-report measure of generalized morality in Study 1 could also be measuring how important participants find morality or social motivation.	This limitation is partly offset by our control variables in Study 1, such as religiosity, which allow us to partly control for individual differences in moral importance. It is also partly offset by the behavioral paradigm in Study 2. Finally, we replicate our analyses using just Items 1–3 in the measure that captures moral beliefs to partly offset a possible confound with social motivation. However, we encourage future research to create self-report measures of generalized morality that hold moral importance constant. Supplemental Study 3 might be a starting point for this research.
2	Our external exposure measure in the Hadza hunter–gatherer study is a proxy for social network size.	This limitation is offset by Study 1, which provides a more direct measure of social network size. We encourage future research that replicates our findings in hunter–gatherer communities using a more comprehensive social network size measure.
3	Agent-based models (ABMs) are computer simulations and do not provide empirical data supporting our claim.	ABMs can nevertheless provide a valuable tool for formalizing a causal theory, and we offset the lack of empirical data in Study 3 through our empirical studies. Studies 2 and 4 are designed to directly emulate our ABM paradigm.
3	We focus on social network size in our theorizing and agent-based model, but anonymity can also arise through other factors, such as relational mobility.	We encourage further research that tests which relational mobility can encourage generalized morality to spread across social networks, even when these networks remain the same size.
5	Our historical NLP analysis does not measure social network size directly.	This limitation is partly offset through Study 1, in which we directly measure social network size. Moreover, our focus on historical change across all English language—rather than a specific community—provides us a large sample size of text (over 155 billion words), which gives us reliable estimates of generalized morality throughout history. We encourage future research that replicates our findings using newspapers, which may correspond to distinct communities with available measures of social network size over history.
5	Our historical NLP analysis only measures changes in the English language.	Focusing on the English language exclusively can be a barrier to generalized insights into human nature. This limitation is partly offset by our study of non-English speakers in Studies 1 and 2. However, we highly recommend that other research replicates our Study 5 findings with non-English corpora.

Note. NLP = natural language processing.

networks. However, we measure generalized morality through a self-report measure that may also measure participants' subjective importance of morality. In Studies 2 and 5, we measure generalized morality more directly through behavioral and linguistic paradigms, but we must rely on proxies for social network size such as exposure outside of hunter–gatherer's local region (Study 2) and the passage of time (Study 5). Finally, many of our studies are correlational, which prevent strong causal inference. Our studies are designed to have complementary limitations (e.g., Studies 1 and 2 have higher external validity, whereas Studies 3 and 4 have higher internal validity; Study 1 measures generalized morality through self-report, whereas Study 2 measures it through a behavioral paradigm), but each study's limitation is nevertheless important to recognize. We write about how future research could address these limitations and extend our findings in the general discussion.

Conclusion

Many centuries ago, Plato and Aristotle debated whether virtue is a generalized and abstract form or whether it is a complex interaction of person and context. Here we suggest that both perspectives are apt depending on whom you ask, where you ask,

and when you ask. According to our model, humans may have culturally evolved increasingly generalized conceptions of moral character, and many people today may use generalized morality to infer cooperation in partner selection dilemmas. Understanding this variation in the structure of morality may help us chart morality's evolution throughout human history and could be the key for predicting the future of moral psychology in a more complex and globalized world.

References

- Abele, A. E., Ellemers, N., Fiske, S. T., Koch, A., & Yzerbyt, V. (2021). Navigating the social world: Toward an integrated framework for evaluating self, individuals, and groups. *Psychological Review*, *128*(2), 290–314. <https://doi.org/10.1037/rev0000262>
- Alvergne, A., Jokela, M., & Lummaa, V. (2010). Personality and reproductive success in a high-fertility human population. *Proceedings of the National Academy of Sciences of the United States of America*, *107*(26), 11745–11750. <https://doi.org/10.1073/pnas.1001752107>
- Ames, K. M. (1994). The Northwest Coast: Complex hunter–gatherers, ecology, and social evolution. *Annual Review of Anthropology*, *23*(1), 209–229. <https://doi.org/10.1146/annurev.an.23.100194.001233>

- Apicella, C. L. (2018). High levels of rule-bending in a minimally religious and largely egalitarian forager population. *Religion, Brain & Behavior*, 8(2), 133–148. <https://doi.org/10.1080/2153599X.2016.1267034>
- Apicella, C. L., Azevedo, E. M., Christakis, N. A., & Fowler, J. H. (2014). Evolutionary origins of the endowment effect: Evidence from hunter-gatherers. *The American Economic Review*, 104(6), 1793–1805. <https://doi.org/10.1257/aer.104.6.1793>
- Apicella, C. L., Marlowe, F. W., Fowler, J. H., & Christakis, N. A. (2012). Social networks and cooperation in hunter-gatherers. *Nature*, 481(7382), 497–501. <https://doi.org/10.1038/nature10736>
- Apicella, C. L., & Silk, J. B. (2019). The evolution of human cooperation. *Current Biology*, 29(11), R447–R450. <https://doi.org/10.1016/j.cub.2019.03.036>
- Atari, M., Haidt, J., Graham, J., Koleva, S., Stevens, S. T., & Dehghani, M. (2022). *Morality beyond the WEIRD: How the nomological network of morality varies across cultures*. PsyArXiv. <https://doi.org/10.31234/osf.io/q6c9r>
- Axelrod, R., & Hamilton, W. D. (1981). The evolution of cooperation. *Science*, 211(4489), 1390–1396. <https://doi.org/10.1126/science.7466396>
- Barclay, P., & Willer, R. (2007). Partner choice creates competitive altruism in humans. *Proceedings of the Royal Society B: Biological Sciences*, 274(1610), 749–753. <https://doi.org/10.1098/rspb.2006.0209>
- Barrett, H. C., Bolyanatz, A., Crittenden, A. N., Fessler, D. M., Fitzpatrick, S., Gurven, M., Henrich, J., Kanovsky, M., Kushnick, G., Pisor, A., Scelza, B. A., Stich, S., von Rueden, C., Zhao, W., & Laurence, S. (2016). Small-scale societies exhibit fundamental variation in the role of intentions in moral judgment. *Proceedings of the National Academy of Sciences of the United States of America*, 113(17), 4688–4693. <https://doi.org/10.1073/pnas.1522070113>
- Berman, J. J., Murphy-Berman, V., & Singh, P. (1985). Cross-cultural similarities and differences in perceptions of fairness. *Journal of Cross-Cultural Psychology*, 16(1), 55–67. <https://doi.org/10.1177/0022002185016001005>
- Bird, R. B., Scelza, B., Bird, D. W., & Smith, E. A. (2012). The hierarchy of virtue: Mutualism, altruism and signaling in Martu women's cooperative hunting. *Evolution and Human Behavior*, 33(1), 64–78. <https://doi.org/10.1016/j.evolhumbehav.2011.05.007>
- Blanc, G. (2020). *Modernization before industrialization: Cultural roots of the demographic transition in France*. Social Science Research Network. <https://doi.org/10.2139/ssrn.3702670>
- Bluemke, M., Jong, J., Grevenstein, D., Mikloušić, I., & Halberstadt, J. (2016). Measuring cross-cultural supernatural beliefs with self-and peer-reports. *PLOS ONE*, 11(10), Article e0164291. <https://doi.org/10.1371/journal.pone.0164291>
- Boyd, R., & Richerson, P. J. (1988). *Culture and the evolutionary process*. University of Chicago Press.
- Brady, W. J., Jackson, J. C., Lindström, B., & Crockett, M. J. (2023). *Algorithm-mediated social learning in online social networks*. <https://osf.io/yw5ah>
- Brady, W. J., McLoughlin, K., Doan, T. N., & Crockett, M. J. (2021). How social learning amplifies moral outrage expression in online social networks. *Science Advances*, 7(33), Article eabe5641. <https://doi.org/10.1126/sciadv.abe5641>
- Caluori, N., Jackson, J. C., Gray, K., & Gelfand, M. (2020). Conflict changes how people view God. *Psychological Science*, 31(3), 280–292. <https://doi.org/10.1177/0956797619895286>
- Cavalli-Sforza, L. L., & Feldman, M. W. (1981). *Cultural transmission and evolution: A quantitative approach*. Princeton University Press.
- Charlesworth, T. E. S., Caliskan, A., & Banaji, M. R. (2022). Historical representations of social groups across 200 years of word embeddings from Google Books. *Proceedings of the National Academy of Sciences of the United States of America*, 119(28), Article e2121798119. <https://doi.org/10.1073/pnas.2121798119>
- Charlesworth, T. E. S., Yang, V., Mann, T. C., Kurdi, B., & Banaji, M. R. (2021). Gender stereotypes in natural language: Word embeddings show robust consistency across child and adult language corpora of more than 65 million words. *Psychological Science*, 32(2), 218–240. <https://doi.org/10.1177/0956797620963619>
- Cheney, D. L., & Seyfarth, R. M. (1988). Assessment of meaning and the detection of unreliable signals by vervet monkeys. *Animal Behaviour*, 36(2), 477–486. [https://doi.org/10.1016/S0003-3472\(88\)80018-6](https://doi.org/10.1016/S0003-3472(88)80018-6)
- Chudek, M., Heller, S., Birch, S., & Henrich, J. (2012). Prestige-biased cultural learning: Bystander's differential attention to potential models influences children's learning. *Evolution and Human Behavior*, 33(1), 46–56. <https://doi.org/10.1016/j.evolhumbehav.2011.05.005>
- Crittenden, A. N., Sorrentino, J., Moonie, S. A., Peterson, M., Mabulla, A., & Ungar, P. S. (2017). Oral health in transition: The Hadza foragers of Tanzania. *PLOS ONE*, 12(3), Article e0172197. <https://doi.org/10.1371/journal.pone.0172197>
- Curry, O., Whitehouse, H., & Mullins, D. (2019). Is it good to cooperate? Testing the theory of morality-as-cooperation in 60 societies. *Current Anthropology*, 60(1), 47–69. <https://doi.org/10.1086/701478>
- Curtin, C. M., Barrett, H. C., Bolyanatz, A., Crittenden, A. N., Fessler, D. M., Fitzpatrick, S., Gurven, M., Kanovsky, M., Kushnick, G., Laurence, S., Pisor, A., Scelza, B., Stich, S., von Rueden, C., & Henrich, J. (2020). Kinship intensity and the use of mental states in moral judgment across societies. *Evolution and Human Behavior*, 41(5), 415–429. <https://doi.org/10.1016/j.evolhumbehav.2020.07.002>
- Davidson, R., Dey, A., & Smith, A. (2015). Executives' "off-the-job" behavior, corporate culture, and financial reporting risk. *Journal of Financial Economics*, 117(1), 5–28. <https://doi.org/10.1016/j.jfineco.2013.07.004>
- Dawkins, R. (2016). *The selfish gene*. Oxford University Press.
- de Lom d'Arce, L. A. (1905). *New voyages to North-America* (Vol. 1). A. C. McClurg.
- Delton, A. W., & Robertson, T. E. (2016). How the mind makes welfare tradeoffs: Evolution, computation, and emotion. *Current Opinion in Psychology*, 7, 12–16. <https://doi.org/10.1016/j.copsyc.2015.06.006>
- Dictionary.com. (2023). *Definition of morality*. Retrieved September 14, 2022, from <https://www.dictionary.com/browse/morality>
- Dunbar, R. I. (2003). The social brain: Mind, language, and society in evolutionary perspective. *Annual Review of Anthropology*, 32(1), 163–181. <https://doi.org/10.1146/annurev.anthro.32.061002.093158>
- Dunbar, R. I. (2009). The social brain hypothesis and its implications for social evolution. *Annals of Human Biology*, 36(5), 562–572. <https://doi.org/10.1080/03014460902960289>
- Dunford, C. (1977). Kin selection for ground squirrel alarm calls. *American Naturalist*, 111(980), 782–785. <https://doi.org/10.1086/283206>
- Earp, B. D., McLoughlin, K. L., Monrad, J. T., Clark, M. S., & Crockett, M. J. (2021). How social relationships shape moral wrongness judgments. *Nature Communications*, 12(1), Article 5776. <https://doi.org/10.1038/s41467-021-26067-4>
- Eisenbruch, A. B., & Krasnow, M. M. (2022). Why warmth matters more than competence: A new evolutionary approach. *Perspectives on Psychological Science*, 17(6), 1604–1623. <https://doi.org/10.1177/17456916211071087>
- Epley, N., Converse, B. A., Delbosc, A., Monteleone, G. A., & Cacioppo, J. T. (2009). Believers' estimates of God's beliefs are more egocentric than estimates of other people's beliefs. *Proceedings of the National Academy of Sciences of the United States of America*, 106(51), 21533–21538. <https://doi.org/10.1073/pnas.0908374106>
- Finkel, E. J., Bail, C. A., Cikara, M., Ditto, P. H., Iyengar, S., Klar, S., Mason, L., McGrath, M. C., Nyhan, B., Rand, D. G., Skitka, L. J., Tucker, J. A., Van Bavel, J. J., Wang, C. S., & Druckman, J. N. (2020). Political sectarianism in America. *Science*, 370(6516), 533–536. <https://doi.org/10.1126/science.abe1715>
- Fiske, S. T. (2018). Stereotype content: Warmth and competence endure. *Current Directions in Psychological Science*, 27(2), 67–73. <https://doi.org/10.1177/0963721417738825>
- Garg, N., Schiebinger, L., Jurafsky, D., & Zou, J. (2018). Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the*

- National Academy of Sciences of the United States of America*, 115(16), E3635–E3644. <https://doi.org/10.1073/pnas.1720347115>
- Gelfand, M. J., Raver, J. L., Nishii, L., Leslie, L. M., Lun, J., Lim, B. C., Duan, L., Almaliach, A., Ang, S., Arndottir, J., Aycan, Z., Boehnke, K., Boski, P., Cabecinhas, R., Chan, D., Chhokar, J., D'Amato, A., Subirats, M., Fischlmayr, I. C., ... Yamaguchi, S. (2011). Differences between tight and loose cultures: A 33-nation study. *Science*, 332(6033), 1100–1104. <https://doi.org/10.1126/science.1197754>
- Gendron, M., Hoemann, K., Crittenden, A. N., Mangola, S. M., Ruark, G. A., & Barrett, L. F. (2020). Emotion perception in hadza hunter–gatherers. *Scientific Reports*, 10(1), Article 3867. <https://doi.org/10.1038/s41598-020-60257-2>
- Gigerenzer, G., Hertwig, R., & Pachur, T. (Eds.). (2011). *Heuristics: The foundations of adaptive behavior*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199744282.001.0001>
- Gilbert, D. T., & Malone, P. S. (1995). The correspondence bias. *Psychological Bulletin*, 117(1), 21–38. <https://doi.org/10.1037/0033-2909.117.1.21>
- Graeber, D. (2012). *Debt: The first 5000 years*. Penguin.
- Graeber, D., & Wengrow, D. (2021). *The dawn of everything: A new history of humanity*. Penguin.
- Graham, J., Haidt, J., Koleva, S., Motyl, M., Iyer, R., Wojcik, S., & Ditto, P. (2013). Moral foundations theory: The pragmatic validity of moral pluralism. *Advances in Experimental Social Psychology*, 47, 55–130. <https://doi.org/10.1016/B978-0-12-4072367.00002-4>
- Graham, J., Haidt, J., & Nosek, B. A. (2009). Liberals and conservatives rely on different sets of moral foundations. *Journal of Personality and Social Psychology*, 96(5), 1029–1046. <https://doi.org/10.1037/a0015141>
- Gray, K., MacCormack, J. K., Henry, T., Banks, E., Schein, C., Armstrong-Carter, E., Abrams, S., & Muscatell, K. A. (2022). The affective harm account (AHA) of moral judgment: Reconciling cognition and affect, dyadic morality and disgust, harm and purity. *Journal of Personality and Social Psychology*, 123(6), 1199–1222. <https://doi.org/10.1037/pspa0000310>
- Gray, K., Young, L., & Waytz, A. (2012). Mind perception is the essence of morality. *Psychological Inquiry*, 23(2), 101–124. <https://doi.org/10.1080/1047840X.2012.651387>
- Grossmann, I., & Varnum, M. E. (2015). Social structure, infectious diseases, disasters, secularism, and cultural change in America. *Psychological Science*, 26(3), 311–324. <https://doi.org/10.1177/0956797614563765>
- Gupta, A. K. (2004). Origin of agriculture and domestication of plants and animals linked to early Holocene climate amelioration. *Current Science*, 87(1), 54–59. <https://www.jstor.org/stable/24107979>
- Haidt, J., & Joseph, C. (2004). Intuitive ethics: How innately prepared intuitions generate culturally variable virtues. *Daedalus*, 133(4), 55–66. <https://doi.org/10.1162/0011526042365555>
- Haidt, J., Koller, S. H., & Dias, M. G. (1993). Affect, culture, and morality, or is it wrong to eat your dog? *Journal of Personality and Social Psychology*, 65(4), 613–628. <https://doi.org/10.1037/0022-3514.65.4.613>
- Hamilton, W. L., Leskovec, J., & Jurafsky, D. (2018). *Diachronic word embeddings reveal statistical laws of semantic change*. ArXiv:1605.09096 [Cs]. <https://doi.org/10.48550/arXiv.1605.09096>
- Harrington, J. R., & Gelfand, M. J. (2014). Tightness-looseness across the 50 united states. *Proceedings of the National Academy of Sciences of the United States of America*, 111(22), 7990–7995. <https://doi.org/10.1073/pnas.1317937111>
- Harris, Z. S. (1954). Distributional structure. *Word*, 10(2–3), 146–162. <https://doi.org/10.1080/00437956.1954.11659520>
- Hartman, R., Blakey, W., & Gray, K. (2022). Deconstructing moral character judgments. *Current Opinion in Psychology*, 43, 205–212. <https://doi.org/10.1016/j.copsyc.2021.07.008>
- Haslam, N. (2016). Concept creep: Psychology's expanding concepts of harm and pathology. *Psychological Inquiry*, 27(1), 1–17. <https://doi.org/10.1080/1047840X.2016.1082418>
- Hauser, M. D., & Marler, P. (1993). Food-associated calls in rhesus macaques (*Macaca mulatta*): II. Costs and benefits of call production and suppression. *Behavioral Ecology*, 4(3), 206–212. <https://doi.org/10.1093/beheco/4.3.206>
- Heinsohn, T. E. (2010). Marsupials as introduced species: Long-term anthropogenic expansion of the marsupial frontier and its implications for zoogeographic interpretation. In S. Haberle, J. Stevenson, & M. Prebble (Eds.), *Altered ecologies: Fire, climate and human influence on terrestrial landscapes*, *Terra Australis* 32 (pp. 133–176). ANU E Press.
- Henrich, J. (2020). *The WEIRD people in the world: How the west became psychologically peculiar and particularly prosperous*. Penguin.
- Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., & Gintis, H. (2004). *Foundations of human sociality: Economic experiments and ethnographic evidence from fifteen small-scale societies*. Oxford University Press. <https://doi.org/10.1093/0199262055.001.0001>
- Henrich, J., & Gil-White, F. J. (2001). The evolution of prestige: Freely conferred deference as a mechanism for enhancing the benefits of cultural transmission. *Evolution and Human Behavior*, 22(3), 165–196. [https://doi.org/10.1016/S1090-5138\(00\)00071-4](https://doi.org/10.1016/S1090-5138(00)00071-4)
- Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, 33(2–3), 61–83. <https://doi.org/10.1017/S0140525X0999152X>
- Hess, Y. D., Carnevale, J. J., & Rosario, M. (2018). A construal level approach to understanding interpersonal processes. *Social and Personality Psychology Compass*, 12(8), Article e12409. <https://doi.org/10.1111/spc3.12409>
- Hester, N., & Gray, K. (2020). The moral psychology of raceless, genderless strangers. *Perspectives on Psychological Science*, 15(2), 216–230. <https://doi.org/10.1177/1745691619885840>
- Heyes, C. (2018). *Cognitive gadgets: The cultural evolution of thinking*. Harvard University Press.
- Heyes, C. M., & Frith, C. D. (2014). The cultural evolution of mind reading. *Science*, 344(6190), Article 1243091. <https://doi.org/10.1126/science.1243091>
- Hill, K. R., Walker, R. S., Zocicević, M., Eder, J., Headland, T., Hewlett, B., Hurtado, A. M., Marlowe, F., Wiessner, P., & Wood, B. (2011). Co-residence patterns in hunter–gatherer societies show unique human social structure. *Science*, 331(6022), 1286–1289. <https://doi.org/10.1126/science.1199071>
- Ibson, L. C., & Mischel, W. (2001). The personality of familiar and significant people: The lay perceiver as a social-cognitive theorist. *Journal of Personality and Social Psychology*, 80(4), 585–596. <https://doi.org/10.1037/0022-3514.80.4.585>
- Iurino, K., & Saucier, G. (2020). Testing measurement invariance of the Moral Foundations Questionnaire across 27 countries. *Assessment*, 27(2), 365–372. <https://doi.org/10.1177/1073191118817916>
- Iyengar, S., Lelkes, Y., Levendusky, M., Malhotra, N., & Westwood, S. J. (2019). The origins and consequences of affective polarization in the United States. *Annual Review of Political Science*, 22(1), 129–146. <https://doi.org/10.1146/annurev-polisci-051117-073034>
- Iyer, R., Koleva, S., Graham, J., Ditto, P., & Haidt, J. (2012). Understanding libertarian morality: The psychological dispositions of self-identified libertarians. *PLOS ONE*, 7(8), Article e42366. <https://doi.org/10.1371/journal.pone.0042366>
- Jackson, J. C., Caluori, N., Abrams, S., Beckman, E., Gelfand, M., & Gray, K. (2021). Tight cultures and vengeful gods: How culture shapes religious belief. *Journal of Experimental Psychology: General*, 150(10), 2057–2077. <https://doi.org/10.1037/xge0001033>
- Jackson, J. C., Hester, N., & Gray, K. (2018). The faces of God in America: Revealing religious diversity across people and politics. *PLOS ONE*, 13(6), Article e0198745. <https://doi.org/10.1371/journal.pone.0198745>
- Jackson, J. C., Lindquist, K., Drabble, R., Atkinson, Q., & Watts, J. (2023). Valence-dependent mutation in lexical evolution. *Nature Human Behaviour*, 7(2), 190–199. <https://doi.org/10.1038/s41562-022-01483-8>

- Jackson, J. C., Rand, D., Lewis, K., Norton, M. I., & Gray, K. (2017). Agent-based modeling: A guide for social psychologists. *Social Psychological & Personality Science*, 8(4), 387–395. <https://doi.org/10.1177/1948550617691100>
- Jackson, J. C., Watts, J., Henry, T. R., List, J.-M., Forkel, R., Mucha, P. J., Greenhill, S. J., Gray, R. D., & Lindquist, K. A. (2019). Emotion semantics show both cultural variation and universal structure. *Science*, 366(6472), 1517–1522. <https://doi.org/10.1126/science.aaw8160>
- Jackson, J. C., Watts, J., List, J.-M., Puryear, C., Drabble, R., & Lindquist, K. (2020). *From text to thought: How analyzing language can advance psychological science*. PsyArXiv. <https://doi.org/10.31234/osf.io/qat4r>
- Ji, L.-J., Lee, A., Zhang, Z., Li, Y., Wang, X.-Q., Torok, D., & Rosenbaum, S. (2023). Judging a book by its cover: Cultural differences in inference of the inner state based on the outward appearance. *Journal of Personality and Social Psychology*, 125(1), 82–99. <https://doi.org/10.1037/pspi0000413>
- Johnson, D. (2016). *God is watching you: How the fear of God makes us human*. Oxford University Press.
- Kendal, J., Giraldeau, L.-A., & Laland, K. (2009). The evolution of social learning rules: Payoff-biased and frequency-dependent biased transmission. *Journal of Theoretical Biology*, 260(2), 210–219. <https://doi.org/10.1016/j.jtbi.2009.05.029>
- Kovacs, K., & Conway, A. R. (2019). What is IQ? Life beyond “general intelligence”. *Current Directions in Psychological Science*, 28(2), 189–194. <https://doi.org/10.1177/0963721419827275>
- Laland, K. N., Odling-Smee, J., & Feldman, M. W. (2001). Cultural niche construction and human evolution. *Journal of Evolutionary Biology*, 14(1), 22–33. <https://doi.org/10.1046/j.1420-9101.2001.00262.x>
- Lammers, J., Gast, A., Unkelbach, C., & Galinsky, A. D. (2018). Moral character impression formation depends on the valence homogeneity of the context. *Social Psychological & Personality Science*, 9(5), 576–585. <https://doi.org/10.1177/1948550617714585>
- Landy, J. F., Piazza, J., & Goodwin, G. P. (2016). When it's bad to be friendly and smart: The desirability of sociability and competence depends on morality. *Personality and Social Psychology Bulletin*, 42(9), 1272–1290. <https://doi.org/10.1177/0146167216655984>
- Lee, R. B., & DeVore, I. (2017). *Man the hunter*. Routledge. <https://doi.org/10.4324/9780203786567>
- Levari, D. E., Gilbert, D. T., Wilson, T. D., Sievers, B., Amodio, D. M., & Wheatley, T. (2018). Prevalence-induced concept change in human judgment. *Science*, 360(6396), 1465–1467. <https://doi.org/10.1126/science.aap8731>
- Li, Q. (2020). *China's development under a differential urbanization model*. Springer.
- Macfarlan, S. J., Remiker, M., & Quinlan, R. (2012). Competitive altruism explains labor exchange variation in a Dominican community. *Current Anthropology*, 53(1), 118–124. <https://doi.org/10.1086/663700>
- Marean, C. W. (2016). The transition to foraging for dense and predictable resources and its impact on the evolution of modern humans. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1698), Article 20150239. <https://doi.org/10.1098/rstb.2015.0239>
- Markus, H. R., & Kitayama, S. (1991). Culture and the self: Implications for cognition, emotion, and motivation. *Psychological Review*, 98(2), 224–253. <https://doi.org/10.1037/0033-295X.98.2.224>
- Marlowe, F. (2010). *The Hadza: Hunter-gatherers of Tanzania* (Vol. 3). University of California Press.
- Marquardt, W. H. (2014). Tracking the Calusa: A retrospective. *Southeastern Archaeology*, 33(1), 1–24. <https://doi.org/10.1179/sea.2014.33.1.001>
- Martin, G. J., & Yurukoglu, A. (2017). Bias in cable news: Persuasion and polarization. *The American Economic Review*, 107(9), 2565–2599. <https://doi.org/10.1257/aer.20160812>
- Mason, L. (2018). *Uncivil agreement: How politics became our identity*. University of Chicago Press. <https://doi.org/10.7208/chicago/9780226524689.001.0001>
- McClintock, C. G., & Allison, S. T. (1989). Social value orientation and helping behavior. *Journal of Applied Social Psychology*, 19(4), 353–362. <https://doi.org/10.1111/j.1559-1816.1989.tb00060.x>
- Mesoudi, A., & Whiten, A. (2008). The multiple roles of cultural transmission experiments in understanding human cultural evolution. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 363(1509), 3489–3501. <https://doi.org/10.1098/rstb.2008.0129>
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G., & Dean, J. (2013). *Distributed representations of words and phrases and their compositionality* [Conference session]. NIPS'13: Proceedings of the 26th International Conference on Neural Information Processing Systems, Red Hook, New York, United States.
- Miller, K., Anyawire, M., Bala, G., & Sands, B. (2013). *A Hadza lexicon*. Mang'ola.
- Mischel, W. (2013). *Personality and assessment*. Psychology Press. <https://doi.org/10.4324/9780203763643>
- Morin, O., & Acerbi, A. (2017). Birth of the cool: A two-centuries decline in emotional expression in Anglophone fiction. *Cognition and Emotion*, 31(8), 1663–1675. <https://doi.org/10.1080/02699931.2016.1260528>
- Motyl, M., Iyer, R., Oishi, S., Trawalter, S., & Nosek, B. A. (2014). How ideological migration geographically segregates groups. *Journal of Experimental Social Psychology*, 51, 1–14. <https://doi.org/10.1016/j.jesp.2013.10.010>
- Murdock, G. P., & Provost, C. (1973). Measurement of cultural complexity. *Ethnology*, 12(4), 379–392. <https://doi.org/10.2307/3773367>
- Muthukrishna, M., Henrich, J., & Slingerland, E. (2021). Psychology as a historical science. *Annual Review of Psychology*, 72, 717–749. <https://doi.org/10.1146/annurev-psych-082820-111436>
- Newson, L., Postmes, T., Lea, S. E., & Webley, P. (2005). Why are modern families small? Toward an evolutionary and cultural explanation for the demographic transition. *Personality and Social Psychology Review*, 9(4), 360–375. https://doi.org/10.1207/s15327957pspr0904_5
- Norenzayan, A., & Shariff, A. F. (2008). The origin and evolution of religious prosociality. *Science*, 322(5898), 58–62. <https://doi.org/10.1126/science.1158757>
- Norenzayan, A., Shariff, A. F., Gervais, W. M., Willard, A. K., McNamara, R. A., Slingerland, E., & Henrich, J. (2016). The cultural evolution of prosocial religions. *Behavioral and Brain Sciences*, 39, Article e1. <https://doi.org/10.1017/S0140525X14001356>
- Nowak, M. A. (2006). Five rules for the evolution of cooperation. *Science*, 314(5805), 1560–1563. <https://doi.org/10.1126/science.1133755>
- Online Etymology Dictionary. (n.d.). *Moral*. Retrieved January 6, 2022, from <https://www.etymonline.com/word/moral>
- Oosterhof, N. N., & Todorov, A. (2008). The functional basis of face evaluation. *Proceedings of the National Academy of Sciences of the United States of America*, 105(32), 11087–11092. <https://doi.org/10.1073/pnas.0805664105>
- Pechenick, E. A., Danforth, C. M., & Dodds, P. S. (2015). Characterizing the Google Books corpus: Strong limits to inferences of socio-cultural and linguistic evolution. *PLOS ONE*, 10(10), Article e0137041. <https://doi.org/10.1371/journal.pone.0137041>
- Pizarro, D. A., & Tannenbaum, D. (2012). Bringing character back: How the motivation to evaluate character influences judgments of moral blame. In M. Mikulincer & P. R. Shaver (Eds.), *The social psychology of morality: Exploring the causes of good and evil* (pp. 91–108). American Psychological Association. <https://doi.org/10.1037/13091-005>
- Pollom, T. R., Cross, C. L., Herlosky, K. N., Ford, E., & Crittenden, A. N. (2021). Effects of a mixed-substance diet on the growth of Hadza children. *American Journal of Human Biology*, 33(1), Article e23455. <https://doi.org/10.1002/ajhb.23455>
- Purzycki, B. G. (2013). The minds of gods: A comparative study of supernatural agency. *Cognition*, 129(1), 163–179. <https://doi.org/10.1016/j.cognition.2013.06.010>

- Purzycki, B. G., Pisor, A. C., Apicella, C., Atkinson, Q., Cohen, E., Henrich, J., McElreath, R., McNamara, R. A., Norenzayan, A., Willard, A. K., & Xygalatas, D. (2018). The cognitive and cultural foundations of moral behavior. *Evolution and Human Behavior*, 39(5), 490–501. <https://doi.org/10.1016/j.evolhumbehav.2018.04.004>
- Purzycki, B. G., Willard, A. K., Klocová, E. K., Apicella, C., Atkinson, Q., Bolyanatz, A., Cohen, E., Handley, C., Henrich, J., Lang, M., Lesorogol, C., Mathew, S., McNamara, R. A., Moya, C., Norenzayan, A., Placek, C., Soler, M., Vardy, T., Weigel, J., ... Ross, C. T. (2022). The moralization bias of gods' minds: A cross-cultural test. *Religion, Brain & Behavior*, 12(1–2), 38–60. <https://doi.org/10.1080/2153599X.2021.2006291>
- Rai, T. S., & Fiske, A. P. (2011). Moral psychology is relationship regulation: Moral motives for unity, hierarchy, equality, and proportionality. *Psychological Review*, 118(1), 57–75. <https://doi.org/10.1037/a0021867>
- Richerson, P. J., & Boyd, R. (2008). *Not by genes alone: How culture transformed human evolution*. University of Chicago Press.
- Roscoe, P. (2006). Fish, game, and the foundations of complexity in forager society: The evidence from New Guinea. *Cross-Cultural Research*, 40(1), 29–46. <https://doi.org/10.1177/1069397105282432>
- Rozin, P., Lowery, L., Imada, S., & Haidt, J. (1999). The CAD triad hypothesis: A mapping between three moral emotions (contempt, anger, disgust) and three moral codes (community, autonomy, divinity). *Journal of Personality and Social Psychology*, 76(4), 574–586. <https://doi.org/10.1037/0022-3514.76.4.574>
- Saucier, G., Thalmayer, A. G., Payne, D. L., Carlson, R., Sanogo, L., Ole-Kotikash, L., Church, A. T., Katigbak, M. S., Somer, O., Szarota, P., Szirmák, Z., & Zhou, X. (2014). A basic bivariate structure of personality attributes evident across nine languages. *Journal of Personality*, 82(1), 1–14. <https://doi.org/10.1111/jopy.12028>
- Schein, C. (2020). The importance of context in moral judgments. *Perspectives on Psychological Science*, 15(2), 207–215. <https://doi.org/10.1177/1745691620904083>
- Schein, C., & Gray, K. (2018). The theory of dyadic morality: Reinventing moral judgment by redefining harm. *Personality and Social Psychology Review*, 22(1), 32–70. <https://doi.org/10.1177/1088868317698288>
- Schulz, J. F., Bahrami-Rad, D., Beauchamp, J. P., & Henrich, J. (2019). The Church, intensive kinship, and global psychological variation. *Science*, 366(6466), Article eaau5141. <https://doi.org/10.1126/science.aau5141>
- Schwartz, S. H. (2012). An overview of the Schwartz theory of basic values. *Online Readings in Psychology and Culture*, 2(1), 2307–0919. <https://doi.org/10.9707/2307-0919.1116>
- Shaver, P., Schwartz, J., Kirson, D., & O'Connor, C. (1987). Emotion knowledge: Further exploration of a prototype approach. *Journal of Personality and Social Psychology*, 52(6), 1061–1086. <https://doi.org/10.1037/0022-3514.52.6.1061>
- Shweder, R. A., LeVine, R. A., & Economiste, R. A. L. (1984). *Culture theory: Essays on mind, self and emotion*. Cambridge University Press.
- Simon, H. A., & Newell, A. (1958). Heuristic problem solving: The next advance in operations research. *Operations Research*, 6(1), 1–10. <https://doi.org/10.1287/opre.6.1.1>
- Singh, M., & Glowacki, L. (2022). Human social organization during the Late Pleistocene: Beyond the nomadic-egalitarian model. *Evolution and Human Behavior*, 43(5), 418–431. <https://doi.org/10.1016/j.evolhumbehav.2022.07.003>
- Smaldino, P. E. (2019). Social identity and cooperation in cultural evolution. *Behavioral Processes*, 161, 108–116. <https://doi.org/10.1016/j.beproc.2017.11.015>
- Smaldino, P. E., Lukaszewski, A., von Rueden, C., & Gurven, M. (2019). Niche diversity can explain cross-cultural differences in personality structure. *Nature Human Behaviour*, 3(12), 1276–1283. <https://doi.org/10.1038/s41562-019-0730-3>
- Smith, K. M., & Apicella, C. L. (2020a). Hadza hunter-gatherers disagree on perceptions of moral character. *Social Psychological & Personality Science*, 11(5), 616–625. <https://doi.org/10.1177/1948550619865051>
- Smith, K. M., & Apicella, C. L. (2020b). Partner choice in human evolution: The role of cooperation, foraging ability, and culture in Hadza campmate preferences. *Evolution and Human Behavior*, 41(5), 354–366. <https://doi.org/10.1016/j.evolhumbehav.2020.07.009>
- Smith, K. M., Larroucau, T., Mabulla, I. A., & Apicella, C. L. (2018). Hunter-gatherers maintain assortativity in cooperation despite high levels of residential change and mixing. *Current Biology*, 28(19), 3152–3157.e4. <https://doi.org/10.1016/j.cub.2018.07.064>
- Smith, K. M., Mabulla, I. A., & Apicella, C. L. (2022). Hadza hunter-gatherers with greater exposure to other cultures share more with generous campmates. *Biology Letters*, 18(7), Article 20220157. <https://doi.org/10.1098/rsbl.2022.0157>
- Snir, A., Nadel, D., Groman-Yaroslavski, I., Melamed, Y., Sternberg, M., Bar-Yosef, O., & Weiss, E. (2015). The origin of cultivation and proto-weeds, long before Neolithic farming. *PLOS ONE*, 10(7), Article e0131422. <https://doi.org/10.1371/journal.pone.0131422>
- Spearman, C. (1904). "General intelligence" objectively determined and measured. *The American Journal of Psychology*, 15(2), 201–293. <https://doi.org/10.2307/1412107>
- Sperber, D. (1996). *Explaining culture: A naturalistic approach*. Blackwell.
- Stibbard-Hawkes, D. N. E., Smith, K., & Apicella, C. L. (2022). Why hunt? Why gather? Why share? Hadza assessments of foraging and food-sharing motive. *Evolution and Human Behavior*, 43(3), 257–272. <https://doi.org/10.1016/j.evolhumbehav.2022.03.001>
- Sylwester, K., & Roberts, G. (2010). Cooperators benefit through reputation-based partner choice in economic games. *Biology Letters*, 6(5), 659–662. <https://doi.org/10.1098/rsbl.2010.0209>
- Talhelm, T., Zhang, X., Oishi, S., Shimin, C., Duan, D., Lan, X., & Kitayama, S. (2014). Large-scale psychological differences within China explained by rice versus wheat agriculture. *Science*, 344(6184), 603–608. <https://doi.org/10.1126/science.1246850>
- Taylor, S. E. (1981). A categorization approach to stereotyping. In D. L. Hamilton (Ed.), *Cognitive processes in stereotyping and intergroup behavior* (pp. 32). Lawrence Erlbaum.
- Thomson, R., Yuki, M., Talhelm, T., Schug, J., Kito, M., Ayanian, A. H., Becker, J. C., Becker, M., Chiu, C. Y., Choi, H.-S., Ferreira, C. M., Fülöp, M., Gul, P., Houghton-Illera, A. M., Joasoo, M., Jong, J., Kavanagh, C. M., Khutkyy, D., Manzi, C., ... Visserman, M. L. (2018). Relational mobility predicts social behaviors in 39 countries and is tied to historical farming and threat. *Proceedings of the National Academy of Sciences of the United States of America*, 115(29), 7521–7526. <https://doi.org/10.1073/pnas.1713191115>
- Turchin, P., Currie, T. E., Whitehouse, H., François, P., Feeney, K., Mullins, D., Hoyer, D., Collins, C., Grohmann, S., Savage, P., Mendel-Gleason, G., Turner, E., Dupeyron, A., Cioni, E., Reddish, J., Levine, J., Jordan, G., Brandl, E., Williams, A., ... Spencer, C. (2018). Quantitative historical analysis uncovers a single dimension of complexity that structures global variation in human social organization. *Proceedings of the National Academy of Sciences of the United States of America*, 115(2), E144–E151. <https://doi.org/10.1073/pnas.1708800115>
- Turchin, P., Whitehouse, H., Gavrillets, S., Hoyer, D., François, P., Bennett, J. S., Feeney, K. C., Peregrine, P., Feinman, G., Korotayev, A., Kradin, N., Levine, J., Reddish, J., Cioni, E., Wacziarg, R., Mendel-Gleason, G., & Benam, M. (2022). Disentangling the evolutionary drivers of social complexity: A comprehensive test of hypotheses. *Science Advances*, 8(25), Article eabn3517. <https://doi.org/10.1126/sciadv.abn3517>
- Turney, P. D., & Pantel, P. (2010). From frequency to meaning: Vector space models of semantics. *Journal of Artificial Intelligence Research*, 37, 141–188. <https://doi.org/10.1613/jair.2934>

- Uhlmann, E. L., Pizarro, D. A., & Diermeier, D. (2015). A person-centered approach to moral judgment. *Perspectives on Psychological Science*, *10*(1), 72–81. <https://doi.org/10.1177/1745691614556679>
- Vignoles, V. L., Owe, E., Becker, M., Smith, P. B., Easterbrook, M. J., Brown, R., González, R., Didier, N., Carrasco, D., Cadena, M. P., Lay, S., Schwartz, S. J., Des Rosiers, S. E., Villamar, J. A., Gavreliuc, A., Zinkeng, M., Kreuzbauer, R., Baguma, P., Martin, M., ... Bond, M. H. (2016). Beyond the ‘east-west’ dichotomy: Global variation in cultural models of selfhood. *Journal of Experimental Psychology: General*, *145*(8), 966–1000. <https://doi.org/10.1037/xge0000175>
- Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of ‘small-world’ networks. *Nature*, *393*(6684), 440–442. <https://doi.org/10.1038/30918>
- Watts, J., Greenhill, S. J., Atkinson, Q. D., Currie, T. E., Bulbulia, J., & Gray, R. D. (2015). Broad supernatural punishment but not moralizing high gods precede the evolution of political complexity in Austronesia. *Proceedings of the Royal Society B: Biological Sciences*, *282*(1804), Article 20142556. <https://doi.org/10.1098/rspb.2014.2556>
- White, C., & Peterson, N. (1969). Ethnographic interpretations of the prehistory of western Arnhem Land. *Southwestern Journal of Anthropology*, *25*(1), 45–67. <https://doi.org/10.1086/soutjanth.25.1.3629467>
- Workman, C. I., Smith, K. M., Apicella, C. L., & Chatterjee, A. (2022). Evidence against the “anomalous-is-bad” stereotype in Hadza hunter gatherers. *Scientific Reports*, *12*(1), Article 8693. <https://doi.org/10.1038/s41598-022-12440-w>
- Wright, R. (2010). *The moral animal: Why we are, the way we are: The new science of evolutionary psychology*. Vintage.
- Yarkoni, T. (2022). The generalizability crisis. *Behavioral and Brain Sciences*, *45*, Article e1. <https://doi.org/10.1017/s0140525x20001685>

Received March 22, 2023
 Revision received July 17, 2023
 Accepted July 25, 2023 ■