

# How Americans Judge: A Topology of Moral Communities



JAMES CHU  AND SEUNGWON LEE 

*Despite research on differences in moral logics across demographic categories, the overall community structure in how Americans share standards of judgment, and hence a fundamental basis for categorical inequalities, remains unclear. To identify communities of shared moral logics, we inductively code judgments in interview transcripts from a probability sample of Americans. We then identify clusters in a network induced by similarities in how Americans judge. We find that competence and prosociality emerge as primary logics by which Americans judge others positively. Gender is the strongest predictor for which moral logics Americans deploy in daily life. Finally, different communities emerge in judgments of institutions, or in negative judgments, suggesting that Americans deploy various moral logics depending on context, which suggests possibilities for bridging categorical divides.*

**Keywords:** moral logics, judgment, praise and denigration, networks

How people judge helps define group boundaries. People who express judgment using shared moral standards are often perceived as belonging to the same group (Lamont 2000; Lamont and Molnár 2002). In this sense, shared standards by which people express judgment toward others, what we call moral logics, contribute to and maintain boundaries and categories (Massey 2007; Sayer 2005; Valentino and Vaisey 2022). In turn, these categorical distinctions can justify unequal distributions of social rewards (Tilly 1998; Tomaskoveic-Devey and Avent-Holt 2019) or shape who organizes together to challenge existing distributions of social rewards (Lipset and Marks 2001).

Motivated by the salience of patterns of shared moral judgment in producing categories, a rich, interdisciplinary line of research has documented how demographic groups in the United States are characterized by distinctive moral logics (Lamont 2000; Thornton, Ocasio, and Lousbury 2012; Finkel et al. 2020; Graham, Haidt, and Nosek 2009; Feinberg and Willer 2019; Fiske 2018; Fiske et al. 2002; Koch et al. 2016). These studies reveal that demographic groups demonstrate distinct ways of expressing judgment (or being judged), but in terms of understanding shared systems of moral standards, they proceed analytically in reverse. Rather than beginning by mapping

**James Chu** is assistant professor of sociology at Columbia University, United States. **Seungwon Lee** is a PhD candidate at Columbia University, United States.

© 2024 Russell Sage Foundation. Chu, James, and Seungwon Lee. 2024. "How Americans Judge: A Topology of Moral Communities." *RSF: The Russell Sage Foundation Journal of the Social Sciences* 10(5): 141–64. <https://doi.org/10.7758/RSF.2024.10.5.06>. Both authors contributed equally to this manuscript. Direct correspondence to: James Chu, at [jyc2163@columbia.edu](mailto:jyc2163@columbia.edu), 606 West 122nd St, New York, NY 10027, United States.

Open Access Policy: *RSF: The Russell Sage Foundation Journal of the Social Sciences* is an open access journal. This article is published under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 Unported License.

moral logics by how Americans cluster together in how they express judgment, and identifying whether these clusters cohere by race, gender, class, or politics, these studies begin with groups and ask whether distinctive moral logics divide them. Even though the assumption of salient groups is analytically appropriate for other purposes, it means that the overall community structure in how Americans share standards of judgment—and hence, a fundamental basis for categorical boundaries—remains obscure. Here we dispense with the presumption of existing groups, and instead inductively identify moral logics by shared patterns in how people judge others while narrating their lives. Rather than asking, for instance, whether the working class has a distinctive logic of judging or being judged, we ask: to what extent can we identify distinctive communities, within which people express judgment using the same moral standards? How well is membership predicted by race, gender, class, or politics?

We answer these questions by coding interview transcripts from the American Voices Project, an unprecedented set of thousands of narrative interviews among a probability sample of Americans. Each interview followed the same semi-structured protocol about respondents' lives, which facilitates comparison. We code each instance of judgment for who is being judged and by what standard. Then, drawing on tools of network analysis, we identify communities of Americans who share similar standards for praising or denigrating various others. We call this network a topology of everyday judgment because it reveals how Americans are located in terms of similarities in how they express judgment in their daily lives. Finally, we discern the extent to which membership in these communities is predicted by categories such as race, gender, poverty, and political identity.

Our study contributes proof of concept for a method that relies on a diverse sample of narratives to induce a topology of moral communities. This method addresses limitations registered in prior research on moral judgment. Although ethnographic research on moral judgment has revealed important insights, it has also been criticized for having selective samples (Bellah 2007, 190). By contrast,

survey methods rely on diverse samples but lack the naturalism afforded by ethnographic research. Survey participants are typically invited to judge hypothetical situations, or to brainstorm and rate hypothetical individuals with varying social characteristics, such as immigrants, or the unemployed. This procedure can reveal differences in how individuals judge the same hypothetical situations, but whether individuals are similar in latent judgments of hypothetical situations is less relevant to the construction of perceived group boundaries than similarities in how they express judgment toward actual people and situations in their lives. Our study addresses both limitations through inductive coding of life narratives from a probability sample of Americans. From here, we can observe similarities in how individuals express judgment and investigate whether and how Americans cluster by distinctive moral logics.

Aside from this methodological contribution, the theoretical contributions of our findings are twofold. First, we find that Americans indeed demonstrate distinct logics in how they judge, but discrete clusters can only be observed when decomposed by who they are judging and by which valence. Rather than a single topology of how Americans judge, multiple topologies are in play. Consistent with pragmatic conceptualizations of culture (Swidler 1986), we suggest that if moral logics are an important driver of group boundaries, then the fact that individuals draw on various sets of moral logics, depending on context, offers new possibilities for bridging categorical divides. Our second finding is that gender is the strongest predictor of the various moral logics Americans deploy in daily life. Our work shows that differences in how we judge cohere closely with gender differences, which is consistent with the enduring inequalities observed through gender segregation in occupations, the division of labor within families, or how parenthood is evaluated (Schwartz and Rubel 2005; Ridgeway 2011; Goldin 2014). To be sure, our null results do not mean other kinds of demographic differences are irrelevant in all contexts. Yet the socialization processes and life experiences that drive differences in how Americans judge appear most distinct across the gender divide.

### WHAT IS A MORAL LOGIC, AND HOW CAN ONE BE OBSERVED?

Following conventions of research on boundaries, valuation, and field theory (Bourdieu and Passeron 1977; Bourdieu 1984; Lamont and Molnár 2002), we define judgment as classifications of alters as good or right, versus bad or wrong. Inherent in this definition is moral worth: to praise or valorize others is to venerate them for being good or right, and to denigrate others is to classify them as bad or wrong (Thai 2022). When expressed standards for what constitutes moral worth are coordinated and shared with others, we call this a moral logic (Lamont 2000). Critically, this definition concerns how people share in their expression of judgment, which is what we claim can be captured through inductive coding of narrative interviews. Where our definition focuses on moral logics as a shared vocabulary for expressing judgment, some scholars have a more expansive definition of moral logic, which encompasses intrapsychic and latent schemata that motivate moral decision-making (Vaisey 2009; D'Andrade 1995). Because individual schemata are not easily captured through semi-structured narrative interviews, our focus here is strictly about the structure of social relations, rather than intrapsychic features of individuals.

When understood as systems of shared vocabularies for expressing judgment, moral logics help to constitute symbolic boundaries that can crystallize inequalities by establishing group boundaries, as well as which group is more deserving of social rewards (Lamont and Molnár 2002; Bourdieu 2000; Tilly 1998). Those who express judgment using the same standards are perceived as belonging to the same group. For instance, the temperance movement was constituted when the erstwhile middle class—largely rural and Protestant, receding in numeric superiority—shared and coordinated in a moral logic of temperance. This logic is evinced in shared *philippics* against urbanites as immoral and inebriated (Gusfield 1986). Other examples of shared moral judgment as a way of boundary-making abound (Lamont 2000; Gast and Okamoto 2016). A key reason to care about moral logics, then, is that they help to produce categories and structural inequalities.

To identify moral logics, cultural sociologists often rely on archival or qualitative accounts. For instance, Robert Bellah and colleagues (2008) rely on interviews to identify regularities in how Americans evaluate others and themselves. A core argument is that the common moral vocabulary or first language of Americans is self-reliance and individualism (Bellah et al. 2008, 20). Jacqui Frost and Penny Edgell (2022) analyze data from forty-eight focus groups to identify what they call a logic of care. Michèle Lamont's (2000) ethnographic study of working-class men suggests that white participants share a moral logic of self-discipline, whereas black participants share a moral logic of care for others. More broadly, ethnographic studies in educational contexts reveal class differences in how people are socialized to express various standards of moral worth, whether in the case of parenting, or experiences in high school and college (Lareau 2011; Calarco 2014; Khan 2011; Jack 2019; Rivera 2016).

Another approach—taken primarily by psychologists—relies on quantitative patterns within responses to questionnaires that ask respondents to judge particular groups of people or situations. For instance, research on stereotype content identifies key axes of judgment that stereotypically apply to distinct groups in society. Seymour Rosenberg and colleagues (1968) asked respondents to categorize sixty-four personality traits in a way that traits in a category are likely to describe the same person. By analyzing shared patterns, the authors conclude that two fundamental axes govern how people judge others. This basic approach has been extended across various research programs (Abele and Wojciszke 2007b; Fiske, Cuddy, and Glick 2007), such as research suggesting that the axes are warmth and competence (Fiske et al. 2002) or communal and agentic (Abele and Wojciszke 2007a). Recent contributions in this line of research ask survey participants to brainstorm various groups and how they differ, identifying empirical regularities from this data (Koch et al. 2016; Zou and Cheryan 2017; Lassetter, Hehman, and Neel 2021).

Both approaches have important strengths. Qualitative research offers unique depth and

recognition of context, producing clear, and often emic, understandings of various logics that govern how people judge. Surveys to elucidate stereotype content explicitly ask diverse samples of respondents for their judgments about various groups, offering greater confidence that the results are general. A core analytic challenge that underlies both approaches, however, is that they begin with categorical groups and ask what distinguishes them in terms of how they express judgment (or are judged). If the goal is to observe moral logics, then the appropriate analysis should begin by identification of shared systems of moral standards before, rather than after, identification of how these logics differ across demographic categories. For instance, black and white working-class men may have different moral logics, but this finding does not reveal whether race or class are the primary fissures that divide how people judge, or perhaps gender or politics. Even with sufficiently diverse samples, as in the case of stereotype content research, the elicited judgments pertain to hypothetical groups of people (men, or women). Men are stereotypically judged for competence, and women for warmth, but this finding does not imply that people can be classified within two coherent systems of shared moral standards.

To address this analytic challenge, we draw inspiration from network operationalizations of culture (DiMaggio 2011; Goldberg 2011; Boutyline and Vaisey 2017). Specifically, we operationalize moral logics as clusters in networks induced by similarities in how everyday judgments are expressed. In network operationalizations of culture, the goal is not to understand the attitudes or attributes of any given individual, but the structure of relations or similarities among individuals. For instance, John Levi Martin (2002) argues that culture consists of unspoken rules that tie certain beliefs together, and that by observing the network structure of how individuals overlap in their beliefs, we can detect these latent rules in action (also see Mohr 1998). The example that liberals tend to drink lattes illustrates how ideological divisions can also drive polarization in lifestyle choices (DellaPosta, Shi, and Macy 2015). Here we adopt the same rationale as

these studies: by observing clusters and other constraints in the network structure of how individuals overlap in their judgments, we see moral logics in action. To identify moral logics, the goal is not to understand how any given American judges, but rather to map the structure of similarities in how Americans express judgment in their lives. We call this a topology of everyday judgment.

Although we did not enter this research deductively, with preregistered hypotheses, prior work offers a rich set of starting intuitions that guided our analysis—both in terms of what network structures we would anticipate in a topology of everyday judgment, and the relevant demographic characteristics that are likely to demarcate people who judge in similar ways. One assumption common in research on stereotype content is symmetry, denoting that usage of a standard when judging positively implies usage of the same standard when judging negatively. At an individual level, this means that observing how someone judges positively is all that is needed to know what moral standard they use to judge negatively (or vice versa). Symmetry implies that the networks of similarities in how Americans judge will be structurally equivalent, whether people are judging positively, negatively, or when ignoring valence. That is, symmetry implies that by studying clusters in how Americans judge positively, we will know how they will cluster negatively, or how they will cluster if we ignore valences of how people judge.

A sociological intuition invoked in prior studies on moral logics, which is somewhat in tension with symmetry, is multiplicity. Pragmatic views of culture posit that, in settled times where power relations and categories are relatively static, moral logics are not shared value systems that govern action but rather a toolkit or set of resources by which individuals justify lines of action (Swidler 1986; also see DiMaggio 2002; Campbell 1998; Hechter et al. 1999; Kaufman 2004). In contrast to schemata that serve as the cognitive infrastructure for how people think or judge (Boutyline and Soter 2021; Vaisey 2009), logics as understood here are about making one's judgment legible to others: what Wright Mills (1940) calls a "vocabulary of motive" or what Robert Bellah and col-

leagues call a “moral vocabulary” (2008, 21). This linguistic analogy is useful in illustrating how a toolkit perspective implies the possibility of moral polyglots, where agents deploy moral logics in ways that are dependent on setting, just as people can speak different languages in different situations. The testable empirical trace of this theoretical claim is multiplicity, or the presence of various topologies of moral judgment, rather than one. Rather than judging all actors using the same shared standards, people may rely on different sets of moral logics for judging some actors, and not others.

In terms of the relevant demographic categories that we expect to predict membership in various moral logics, prior work sensitizes us to gender, race, class, and politics. For instance, Lamont’s work leads us to expect white men to cluster in their judgments of self-reliance and diligence, but black men to form a distinct community of judgment that places greater value for prosociality and helping others (2000). Given research on cultural capital, we might expect class to predict adoption of distinct moral logics. Given stereotypical differences in how men and women are judged, we might expect gender as a strong predictor of which moral logics people adopt. As a reflection of rising political polarization (see Finkel et al. 2020; Graham, Haidt, and Nosek 2009), and the possibility of different moral foundations, we might also anticipate divisions by political identification. To be sure, there are numerous other potential cleavages. We could investigate the role of generational or religious (such as Christians versus non-Christians) categories in predicting how Americans judge. To keep our analysis tractable, however, we focus on key axes motivated clearly in prior work, saving additional possibilities for future work.

## DATA AND METHODS

Our analysis relies on transcripts of interviews collected as part of the American Voices Project

(AVP). Participants were selected using cluster random sampling of census tracts and block groups, with oversampling of households in the bottom half of the income distribution. Additional detail about the construction of the AVP sample is described in the introduction (see Edin et al. 2024, this issue), but for our purposes, the relevant features of this data are that the interviewees reflect a diverse probability sample of American households, and that interviews were conducted by trained interviewers through a consistent protocol. Interviews from this project were transcribed and made available for analysis on a secure server.

We recruited and trained three research assistants (RAs) to code these interviews for attributions of praise and denigration in NVivo. For each judgment, we asked RAs to identify whether the judgment was praise or denigration, who is being judged, by what standards, and a keyword in the interview that summarizes the judgment. By our definition, observations that are empirical and do not contain normative judgments are beyond the scope of our coding. Additionally, we further limited coding to praise and denigration originating from or directed to the interviewee.<sup>1</sup>

To ensure that the task was tractable for RAs, the authors produced a systematic coding scheme based on our own inductive coding until saturation ( $N = 37$  interviews). This inductive coding of transcripts led us to derive thirty-six standards by which judgments were expressed—eighteen positive and eighteen negative. We later consolidated the standards that appeared fewer than five times across the coding of our RAs, with a final set of standards encompassing ten positive and twelve negative standards of judgment. General judgments wherein specific criteria could not be inferred from context were excluded from the analyses (such as “My mom is an amazing person” does not imply a clear standard of moral worth, unless the preceding section of the transcript described the mother’s generosity).<sup>2</sup> For full

1. We originally wanted to include cases where the interviewee describes another individual judging a third party, but we concluded that these are second order judgments—perceptions about how others judge—that are analytically distinct from moral logics.

2. One concern is that removing these general judgments leads us to overstate the distinctiveness of how Americans judge. Hence we also conducted the same analyses with the generic judgments included as a robust-

transparency, we also include our original coding scheme in the online appendix.<sup>3</sup>

Similarly, we inductively identify alters being judged. Americans may not praise politicians using the same standards they use to praise their children. Hence, our original analytic strategy was to investigate different standards used to judge specific alters. Because of the limited frequency of the occurrence of certain alters, however, we do not report on results of such an analysis here. Instead, in the present manuscript, we have aggregated alters into two alter categories that are analytically distinct: individuals and institutional actors, who are being judged as representatives of institutions, such as the police, politicians, educators, health-care workers, and the like. Table A.1 presents a list of all alters coded in these two categories.

Because of funding limitations, we could not hand-code all 1,613 interview transcripts in the AVP corpus. Given resource constraints and absent *ex ante* hypotheses that would inform power analyses, we opted to randomly sample and code as many transcripts as possible, resulting in an analytic sample of 355 transcripts. A comparison of the demographic characteristics of the 355 randomly sampled interviewees, versus the broader sample of 1,613 interviewees, is reported in appendix C, which shows no statistically significant differences when multiple hypothesis corrections are included. The selected transcripts were hand coded by the authors and three research assistants, with several transcripts being coded by multiple coders to check for reliability. Details of the coding and validation process are discussed in appendix D.

### ANALYTIC STRATEGY

Our analysis proceeds as follows. We first enumerate the frequency by which interviewees

used various standards of judgments. Second, we construct a network where the nodes are interviewees, and the connections between these nodes (or edges) represent how much each interviewee pair overlaps in how they judge. To quantify overlap we calculate a Jaccard index for all pairs of interviewees, which represents the ratio of judgments with shared standards over the total number of judgments deployed across a pair of interviewees. This index ranges from zero (two respondents used completely different standards in their judgments) to one (two respondents used exactly the same standards in all their judgments).<sup>4</sup>

After producing a network of similarities in how Americans express judgment, we then apply community detection computational techniques (the Louvain algorithm) to identify clusters of shared standards, or moral logics (Traag, Walman, and van Eck 2019). This algorithm identifies clusters that maximize modularity, defined as the density of the weighted edges within clusters relative to weighted edges outside. In lay terms, the algorithm runs through various ways of grouping interviewees, each time checking to see how strong the connections are within the candidate group (that is, the degree to which people overlap in how they judge) versus outside the group. The algorithm continues through various groupings until it finds one in which similarities within groups are as strong as possible, relative to similarities with others outside the groups. We first apply this procedure to the full dataset of judgment, including both positive and negative valences and all alters being judged. We repeat this procedure for four subsets of judgments: praise about individuals, denigrations about individuals, praise about institutional actors, and denigrations about individuals.

Third, we characterize the various standards that constitute each emergent community. We

---

ness test. Whether generic positive and generic negative judgments are included, the results are substantively identical to our main findings except that there is an additional cluster characterized by generic judgments. See appendix A (<https://www.rsfjournal.org/content/10/5/141/tab-supplemental>).

3. See appendix B.

4. For example, if the judgments of interviewee A relied on standards of competence three times, safe one time, and sociable one time; and if the judgments of interviewee B relied on standards of competence four times and diligence two times, the Jaccard index of overlap between A and B is  $0.636 \approx (3 + 4) / (3 + 1 + 1 + 4 + 2)$ .

use pairwise t-tests to detect whether the standards used in each community are different from those used outside the community. To aid in substantive interpretation, we use radar plots to indicate the prevalence of each standard, indicating the frequency by which a given standard is invoked, with asterisks to indicate statistical significance. We additionally label each community according to the standards that are most prevalent.

Fourth, we investigate how well demographic categories like race, gender, class, and politics predict membership in various moral communities. Because the dependent variable is an unordered category corresponding to membership in various communities, multinomial logistic regressions are used to predict the membership. As independent variables, we include the demographic categories of interest: self-identified race (black, white, Hispanic or Latino, Others),<sup>5</sup> gender (women or Others, men), poverty (having a household income under the poverty line following the Census Bureau, adjusted for household size, versus not being under the poverty line), and political identity (Democrat, Republican, independent, no preference or missing).<sup>6</sup> As covariates, we include age (eighteen to thirty-four, thirty-five to fifty-four, fifty-five or older, missing), educational level (no bachelor's degree or missing, bachelor's or higher degree), region (Midwest, Northeast, South, West), and coder fixed effects (to condition for cross-person variation in how coders approached the interviews).

Finally, to offer an overall assessment of which demographic category is most predictive of membership, we present analyses of how these demographics contribute to model fit. To

identify which demographic categories best predict membership overall, we test how the Akaike Information Criterion (AIC) changes when demographic categories are included in, versus left out of, the multinomial logistic model predicting membership in various moral logics. In a null model, we include only coder fixed-effects. We then compare the AIC from the null model with one where a single demographic category is included. Smaller AIC values denote better model fit. Hence, when we take the difference between the null model and a model where the demographic category is included, a larger difference indicates that the demographic category contributes more to model fit.

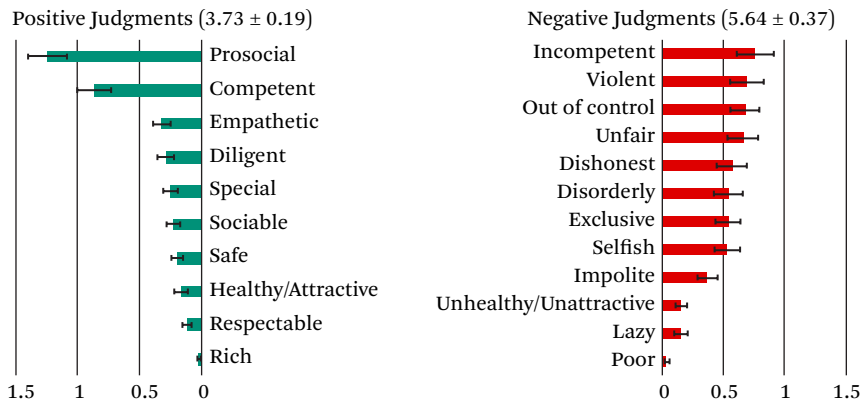
Throughout our analyses, we introduce various robustness checks as appropriate. For instance, we analyze whether our findings are sensitive to how we coded for standards. To do so, we repeat our analyses when pooling certain standards that could logically be considered subsets of other standards (for example, perhaps honesty is a subset of prosociality). Additionally, we use principal components analysis (PCA) to investigate the factor structure of the standards deployed in the judgments we coded, to check whether our coding could be aggregated further.

## RESULTS

In this section, we report the results in the following order. First, we describe the frequency of various judgments in the data. Second, we demonstrate networks of shared standards in all judgments—both positive and negative—toward all alters. Third, we divide judgments by the alter category and valence to present four

5. Although we think those who identify as Asian are important, fewer than eleven interviewees identified as Asian in our sample of  $n = 355$ . Hence, for this analysis we group them under Others, not as an indication of theoretical disinterest, but as a reflection of pragmatic data limitations.

6. In gender and educational level variables, we originally coded other gender and missing education level as separate categories. However, only a few respondents belong to these groups, which raised the concern about disclosing their identities. Therefore, we put other gender and women together, and missing educational level and no bachelor's degree together. The assumption is that respondents in these small categories are more similar with respondents in the larger categories with which they have been grouped. When we refer in our subsequent analysis and discussion to a larger group of respondents (for example, women), such references also include any members of small groups (such as other gender) that have been combined with the larger group for disclosure avoidance purposes. The results are substantially the same regardless of how these categories are coded.

**Figure 1.** Descriptive Statistics of Count of Judgments

Source: Authors' calculations.

Note: Numbers in the plot titles represent the mean  $\pm$  standard error. The optimal way to view the figures in this article is in color. We refer readers of the print edition of this article to <https://www.rsfjournal.org/content/10/5/141> to view the color versions.

networks of shared standards in judgments: praise about individuals, denigrations about individuals, praise about institutional actors, and denigrations about institutional actors. Along with each network topology, we also report the standards that different clusters in these networks deploy in their judgments. Finally, we describe the key demographic cleavages in these four networks.

### Frequency of Various Judgments

Figure 1 summarizes the frequency of praise and denigration, based on each standard. To illustrate what these standards capture, we also provide quotes and keywords associated with each in table A.2. On average, interviewees made more negative judgments (mean count = 5.64) than positive ones (mean count = 3.73). In praise, prosociality and competency were the two primary standards of judgments. Whereas praise involved two primary standards, the frequency of standards used in everyday denigration were more evenly distributed, as evidenced by the distribution of frequencies.<sup>7</sup> These descriptive results suggest that—in the aggre-

gate—Americans are more diverse in their negative judgments (denigration), relying on a variety of different standards. By contrast, Americans are more specialized in matters of praise.

### A Multiplicity of Topologies

How, if at all, do Americans cluster into communities regarding the standards they use to praise and denigrate others? As presented in figure 2, we do not detect clear communities in a network representing similarities in the standards people use to judge, including both positive and negative judgments, and both individuals and institutional actors. No clusters emerge from visual inspection, and, more important, our community detection algorithm cannot robustly identify groups. The maximum modularity of the algorithm is 0.119, about half that of other networks that will be presented in the following sections (which, on average, show a maximum modularity of 0.225). These results, of course, do not mean that no moral standards by which Americans judge others are shared. Perhaps a clearer community structure

7. We demonstrate the exact phrases or keywords for the judgments as word clouds in online appendix E. The results are consistent with the findings in figure 1. Keywords indicating prosociality (for example, support or help) or competence (for example, smart) emerge as the most frequent keywords for praise. In denigrations, keywords for being violent (for example, abusive, beat, or hit) unfair (unfair), out of control (for example, drugged, crazy), and disorderly (for example, racist) show up at the center.



**Figure 2.** Structure of Moral Communities Induced from Shared Standards of Judgments (Modularity = 0.119)



Source: Authors' calculations.

Note: Network ties are defined by shared judgments toward the same alter groups including both praise and denigration and both individual actors and institutional actors. Shading represents communities detected using the Louvain algorithm. Each node is an interviewee, and edges represent overlap in usage of standards of judgment. Thicker edges represent greater overlap. Edges across different clusters are not presented in the graph for clearer visualization of the clusters.

emerges when investigating how Americans judge individuals versus institutional actors, or by valence, suggesting that moral logics are present or salient only in specific contexts.

Figure 3 illustrates the communities identified in a network of shared standards of praise toward individuals. We include radar plots that indicate mean differences in the frequency of each standard between the members of a community and all other interviewees outside of the community. For instance, positive values indicate that the standard was more likely to be observed among members of the community than the sample average, with stars indicating statistically significant t-tests of differences in means between the community and other communities. Line width is proportional to the size of the community; the larger the community the thicker the line. Negative values indicate

that the standard was less likely to be observed.<sup>8</sup>

Shared standards of praise about individuals (figure 3) are characterized by two main communities, with two peripheral communities. These clusters are clearly distinguished, with a modularity of 0.211. The largest ( $N = 96$ , upper right) community consists of interviewees who shared standards in praising individuals for competence. Interviewees belonging to the second largest ( $N = 92$ , upper left) community shared standards in praising individuals for being prosocial. The associated radar plot implies that members of these two communities were specialists—relying rarely on other standards. Indeed, those in the prosocial community were less likely to praise individuals as being competent.<sup>9</sup> On the periphery of this pair of core communities is a third community con-

8. Full results of these t-tests are reported in online appendix F.

9. As an additional check, we use negative binomial regressions with coder fixed effects and demographic controls to test whether interviewees who praised individuals for prosociality were indeed less likely to praise for competence. We find this indeed to be the case (online appendix G).


sisting of individuals who did not judge for prosociality and competence, but instead on a more general set of standards, such as diligence, being special, and empathetic ( $N = 64$ , bottom). Our community detection algorithm also reports another small ( $N = 18$ ) cluster that does not appear to use a clear set of standards. Another eighty-five interviewees were isolates, that is, did not overlap in their standards of judgments with others, or were not recorded as praising individuals.

Is the structure of denigration a mirror image of praise? In a word, no. Clear clusters also emerge in a topology of denigration (the modularity score as maximized by our community detection is 0.238), but the topology is asymmetric from praise. Whereas a network induced by similar standards for praise is characterized by two large communities, each of which has a specialized standard (competence or prosociality), the main communities in a topology of denigration are characterized by sharing across combinations of diverse standards. Figure 4 shows that the center of the network is a large community ( $N = 60$ , bottom center) characterized primarily by judgments of

violence. (The radar plot indicates that other standards like being out of control, dishonest, or disorderly are also statistically significantly more likely to appear, as well.) The second largest community ( $N = 54$ , upper left) is primarily characterized by judgments of individuals who are out of control. Notably, people in this community eschewed judgments of violence, dishonesty, laziness, or disorderliness. Additionally, judgments of incompetence form a discrete community, but only at the periphery ( $N = 32$ , upper center), and those of selfishness are found in a larger, yet still peripheral community that includes judgments of impoliteness ( $N = 53$ , right side). Judgments about laziness do not appear as a discrete community and are infrequent. Hence we find little evidence of symmetry in the network structure of how people positively and negatively judge individuals.

We next turn to how Americans praise institutional actors, such as police, educators, religious figures, and politicians. Figure 5 shows that the network structure of shared praise toward institutional actors is again clearly defined, with a maximum modularity score of

**Figure 3.** Structure of Moral Communities Induced from Shared Standards of Praise Toward Individuals (Modularity = 0.211)

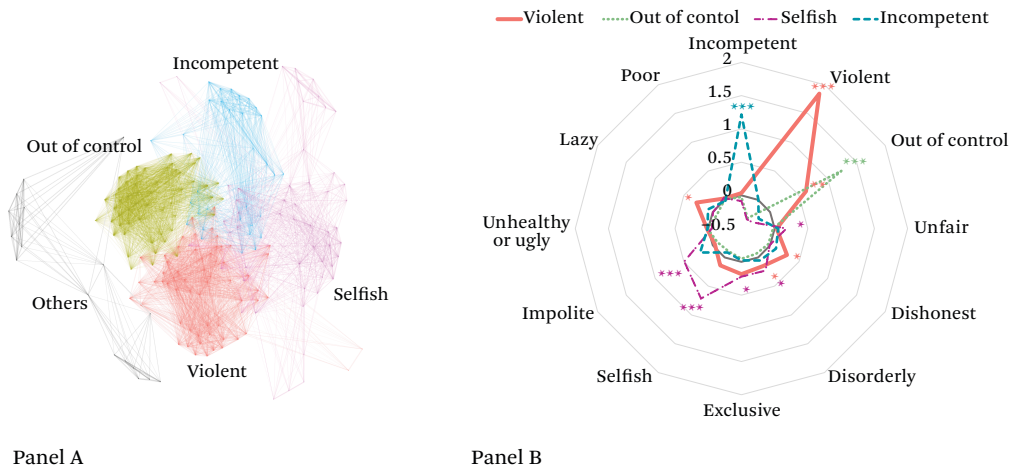


Source: Authors' calculations.

Note: Edge width in the network is weighted by the jaccard index. Width of the lines in the radar plot is weighted by the size of each cluster. Edges across different clusters are not presented in the graph for clearer visualization of the clusters. Asterisks in the radar plot represents the statistical significance of the difference between the entire sample mean and the group mean.

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

**Figure 4.** Structure of Moral Communities Induced from Shared Standards of Denigration Toward Individuals (Modularity = 0.238)



Source: Authors' calculations.

Note: Edge width in the network is weighted by the jaccard index. Width of the lines in the radar plot is weighted by the size of each cluster. Edges across different clusters are not presented in the graph for clearer visualization of the clusters. Asterisks in the radar plot represents the statistical significance of the difference between the entire sample mean and the group mean.

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

0.274. Although competence and prosociality also appear as large clusters, the network structure of how people praise institutional actors differs in two key ways. First is a cluster by which Americans specialize in judging institutional actors for empathy, or a willingness to tolerate and welcome individuals. For example, a respondent who had immigrated to the United States decades ago described their experience with medical professions: “At that hospital, both doctors and nurses, people who work on the front lines—since I had to get in through the emergency room—provided a great service, to be honest, they have my utmost respect, the profession they practice, even though I told them that I [didn’t have medical insurance], they still did everything they could, because the person’s health comes first.”

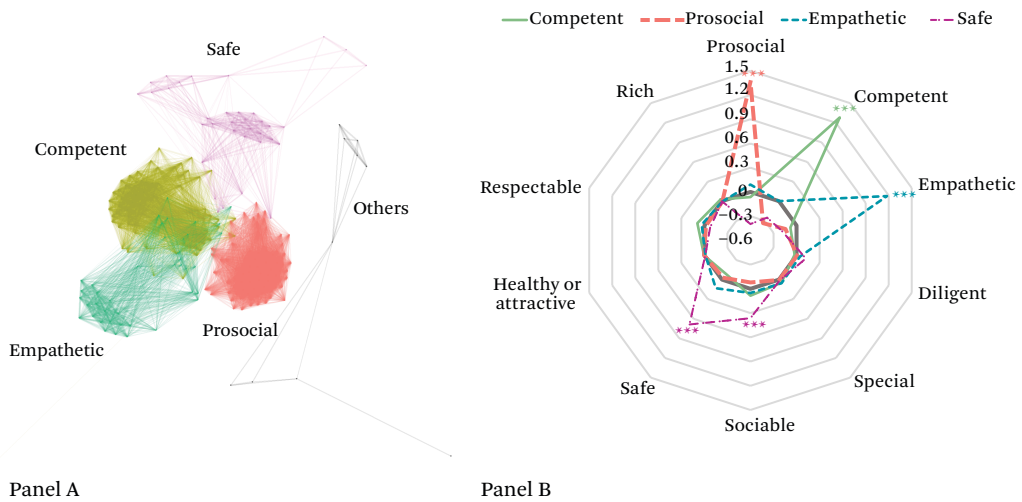
Second is a smaller, secondary cluster of Americans who judge institutional actors for being safe, as when a white male respondent answered a question about police, “But the law enforcement I feel that they’re there if I need them. I feel protected with them and they’re

like an insurance policy until you really need them that they don’t exist, but when you need them, you need them. . . . With law enforcement I respect them and I know I’m protected by them.”

The presence of these sizable, alternative, clusters means that the set of moral logics by which Americans judge institutional actors cannot be assumed to mirror how they judge individuals.

Finally, we show how Americans cluster in their denigration of institutional actors in figure 6. The modularity of this network is somewhat lower (0.175). Cross-linkages across communities are higher, reflecting perhaps the shared, public nature of critical discourse about institutions. To the extent that distinctive communities can be identified, the largest community ( $N = 77$ , bottom left) is characterized primarily by judgments of disorderliness, incompetence, and violence. The second largest community consists of judgments about dishonesty, especially about politicians ( $N = 66$ , left side). The topology of denigration is characterized by an even distribution of cluster

**Figure 5.** Structure of Moral Communities Induced from Shared Standards of Praise Toward Institutional Actors (Modularity = 0.274)



Source: Authors' calculations.

Note: Edge width in the network is weighted by the jaccard index. Width of the lines in the radar plot is weighted by the size of each cluster. Edges across different clusters are not presented in the graph for clearer visualization of the clusters. Asterisks in the radar plot represents the statistical significance of the difference between the entire sample mean and the group mean.

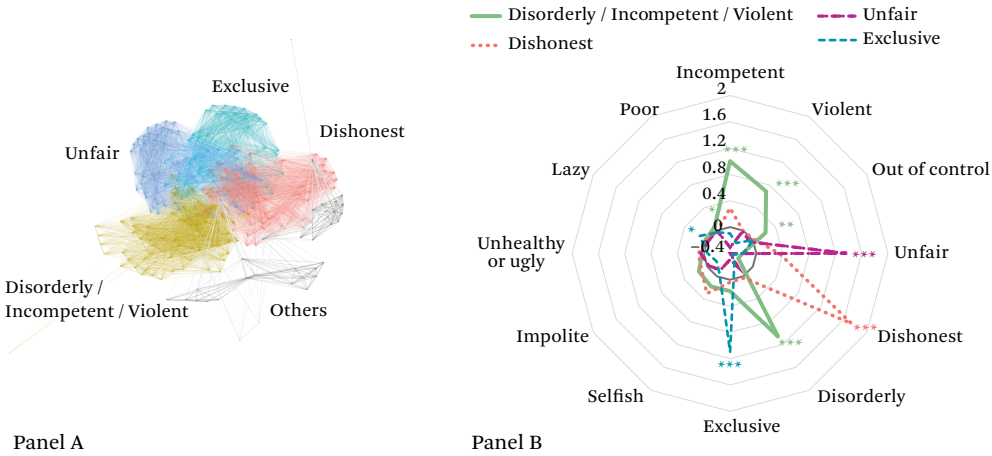
\*\*\*  $p < .001$ ; \*\*  $p < .01$ ; \*  $p < .05$

sizes, and lower modularity, relative to a topology of praise. The structural difference likely reflects the wider range of moral repertoires Americans use when they denigrate institutional actors compared to when they praise others. For instance, largely overlapping with the foregoing two main communities are two other sizable communities: the unfair community ( $N = 55$ , upper right) and exclusive community ( $N = 49$ , upper side). Indeed, there are also two peripheral and smaller communities with fewer than twenty members. The presence of these communities again underscores asymmetry between praise and denigration: judgments of uncaring or unsupportive (selfish) institutions appear in one of the smallest communities ( $N = 14$ ), whereas judgments of supportive and caring institutions appears in the largest, central community in matters of praise.

The general insight that emerges from investigation of these various topologies is that praise and denigration are not symmetric, and

the way that Americans judge institutional actors is not identical with how they judge individuals. Rather than being mirror images, whereby judgments occur along axes that have positive and negative valence, the differences in topology imply that the moral logics Americans deploy differs between praise and denigration. As a robustness test, we address concerns that this lack of symmetry is an artifact of our coding. In online appendix H, table 2–5, we show how our results are robust even if we re-code standards using a less granular scheme, as when we group diligence and competence, group lazy and incompetence, group empathetic, respectful, and prosocial, and group exclusive, impolite, and selfish. Additionally, in online appendix I, we show that our results are robust to further aggregation of how we coded for judgments. To do so, we conducted PCA to reduce the dimensionality of the twenty-two primary standards (ten positive and twelve negative) that we defined. The finding of asymmetry remains robust.

**Figure 6.** Structure of Moral Communities Induced from Shared Standards of Denigration Toward Institutional Actors (Modularity = 0.175)



Source: Authors' calculations.

Note: Edge width in the network is weighted by the jaccard index. Width of the lines in the radar plot is weighted by the size of each cluster. Edges across different clusters are not presented in the graph for clearer visualization of the clusters. Asterisks in the radar plot represents the statistical significance of the difference between the entire sample mean and the group mean.

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

**Describing Key Demographic Cleavages**

Having inductively identified moral logics through communities in a network of similarities by how Americans judge, we now turn to the second part of our research question: what demographic categories best predict membership in these various moral logics? For each topology, we briefly describe the primary demographic distinctions, which are also the ones presented in figures 3 through 6. However, for transparency, we report full results from multinomial logistic regressions predicting membership in all communities as a function of demographics in online appendix J.

When it comes to the moral logics by which Americans praise individuals, our results broadly align intuitions from prior work. For instance, membership in different moral communities (or moral logics) is predicted by gender. Women are more likely to praise individuals for their prosociality than men: the predicted probability of women is 29.7 percent, and that of men 18.7 percent. This finding implies that the stereotypes applied toward cer-

tain groups of people are reflected in their judgments: women, who are stereotypically considered more prosocial, tend also to judge other individuals by the same standard. For instance, a white woman in her seventies said the following about her husband, reminiscing on when they were still young:

He said, "Nobody should have to bury their parents alone." And he took off from work, to come with me, having just lost his mother, to come with me to bury my mom and, my step dad went to his family to be buried, and we're driving back and he still, you know, every couple of weeks he asks me to marry him but, you know, I could do worse than this guy. We really like each other, we get along well, we're good friends, and look how kind he is!

Here the respondent praised her husband for being kind and supportive, the very standards by which women are stereotypically judged for.

Democrats are about twice as likely as Re-

publicans to adopt a moral logic of prosociality when it comes to praising individuals. This result affirms findings suggesting that political polarization in America has led to segregated moral communities, each with distinct moral foundations (Chen and Rohla 2018; Finkel et al. 2020).<sup>10</sup> Race is also a significant determinant of moral logics for praise. Those who identify as Latino or Hispanic are more likely to praise individuals primarily in terms of prosociality (predicted probability 31.9 percent) than those who identify as black (20.1 percent) or as white (25.4 percent). Conversely, they are less likely to judge for competence than those identifying as white or black (Latino or Hispanic 19.5 percent, white 29.7 percent, black 29.3 percent). We conjecture this finding may be attributable to immigration experiences, where individuals who identify as Latino or Hispanic rely on family networks, as opposed to institutional ones, for support. Despite research suggesting that lower income individuals value prosociality or behave more generously (Piff et al. 2010), we did not find that low-income individuals were more likely to appear in the prosocial community.

When we investigate the other topologies, only gender remains as a consistent divide that structures membership in various moral communities. When it comes to a topology of how Americans denigrate individuals, gender remains the main predictor of moral logics: women are more likely than men to denigrate individuals for being violent (women 19.0 percent, men 8.5 percent). The modal judgments in this community tended to involve frustrations about family or partners, usually male, being abusive and violent. A black female respondent describes her ex-husband at the time when they were going through divorce, “He was stalking me. He would sit outside in the parking lot and finally . . . Oh, and he was cutting my car off. He was doing all this crazy stuff because I told him I was going to divorce him.”

Father figures are another frequently denigrated alter in this community. Another female respondent mentions how violent her father used to be: “He used to hit my mom and us too

and I say, what my dad used to do to us I won’t do that to my kids, I’m going to be the best mom for my kids, I don’t want them to suffer like I did.”

By contrast, men are more likely than women to belong to the second largest community, where “out of control” was the primary standard (men 23.3 percent, women 13.2 percent). The defining characteristic of this community is that the members denigrate other individuals for being out of control (such as addicts), but eschew other negative judgments. For example, a male respondent described his father in the following way: “My daddy raised us. When my daddy raised us, it was just really rough. My dad is an alcoholic, so we kind of grew up poor in poverty. I really don’t want to say nothing bad, but truth be told, it was b—.”

While recalling his mother, he added similar judgments about his mother and stepfather: “What can I say about this lady? She’s a mess, really. My mom and my stepdad had met. My stepdad was on drugs. . . . We didn’t have no structure or disciplines, kind of did what we wanted to do pretty much.”

In terms of how Americans praise institutional actors, we find that gender again emerges as the most salient predictor. Women are significantly more likely than men to judge institutions primarily in terms of their prosociality (women 21.7 percent, men 14.7 percent), whereas men are more likely than women to judge institutions in terms of their competence (men 23.2 percent, women 15.9 percent).

Finally, when investigating demographic dividing lines in moral logics for how Americans denigrate institutional actors, gender is again the primary category that stratifies these communities. Men are more likely than women to denigrate institutional actors for being dishonest (men 26.8 percent, women 16.0 percent) and for being disorderly, incompetent, and violent (men 26.5 percent, women 21.6 percent). Here we also find important racial distinctions. Black participants were more likely to denigrate institutions for multiple features such as disorderliness, incompetence, and, as expected from the Black Lives Matter movement, vio-

10. Although, as we will see, political polarization is not salient in predicting moral logics in negative judgments, or in judgments of institutional actors.

lence. For instance, a black male interviewee criticized police officers for being racist and violent, concluding with the following statement: “For me honestly, people have been getting beat up and killed by the police all my life, so it’s like nothing new to me, to me I’m numb to it, like I expect to be harrassed by the police, just because I’m black in America. . . . We’ve been knowing that police has been killing f—ing minorities forever.”

Other demographic fault lines, such as indicators of poverty or political identification, do not emerge as statistically significant predictors.

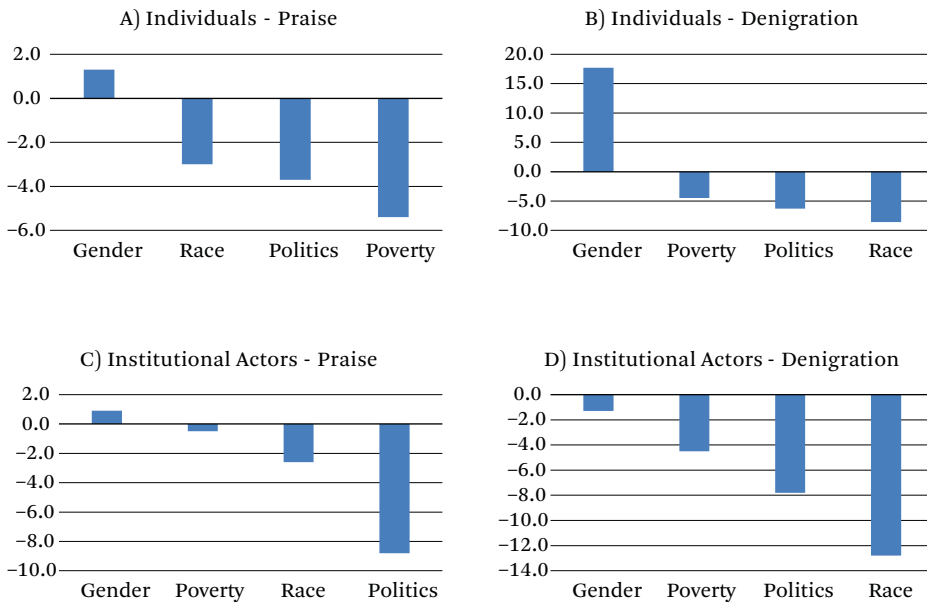
**Gender Is the Strongest Overall Predictor of Moral Logics**

To offer a more comprehensive analysis of which demographic category best predicts what moral logics people deploy, figure 7 presents the differences in model fit between a null model where only coder fixed effects are used in a multinomial logistic regression to predict

membership of the moral cluster, and an otherwise equivalent model where we add a demographic variable—either race, gender, poverty, or political identification—to the null model. A key reason to rely on AIC results, relative to our earlier descriptive results, is that these analyze how well demographic categories predict membership across all communities simultaneously, rather than investigating descriptive differences within communities one at a time (such as using paired t-tests), which also increases concerns about false positive results.

Panels A through D indicate the difference in AIC between each model. As noted, if this difference is positive, then it indicates that the demographic category is predictive of membership in various moral logics. If the difference is negative, then it means the demographic category does not improve model fit, and the gender category does not predict sorting across moral logics. We repeat this analysis for all four networks: praise about individuals, denigrations about individuals, praise about institu-

**Figure 7.** Difference Between the AIC from Null Model and the AIC from Model with Each Demographic Variable Added



Source: Authors’ calculations.

Note: Null models predict membership of clusters using only coder-fixed effects. The AICs from the null models are compared with models where each of the demographic variables is added to the coder-fixed effects.

tional actors, and denigrations about institutional actors, respectively.

Panel A demonstrates that gender is the only demographic category that improves model fit in praising individuals. When compared with our earlier descriptive results, the other categories no longer appear to improve model fit. Although the descriptive results suggest that they are salient in predicting membership in certain communities, these AIC results suggest they do not meaningfully predict membership across communities in general. Panel B again shows that gender is a predictor of membership in moral logics of how people denigrate individuals, and the other demographic categories again appear to contribute little to model fit. Panel C continues to indicate that gender improves model fit, whereas inclusion of the other categories does not. Panel D indicates, however, that all demographic categories fail to improve model fit, even if gender remains (relatively speaking) more salient than other demographic variables. We conjecture that this is due to the low modularity of communities in denigrations about institutional actors, meaning that the boundaries between communities are poorly defined in the first place, and hence demographic categories are less likely to predict membership in these communities.

## DISCUSSION AND CONCLUSION

We inductively identified instances of praise or denigration in transcripts of life narratives from a random sample of Americans, as well the standards by which these moral judgments were expressed. By detecting communities of people who have the same standards in how they express judgment, we sought to identify coherent moral logics. We did not find clear moral logics when investigating overall similarities in how people judge. Instead, we detected communities only when we analyzed praise and denigration, as well as judgments of individuals versus institutional actors, separately. We found little evidence of symmetry. The results instead offer support for claims of multiplicity, where the moral logics that people deploy depend on specific circumstances, such as the valence of judgment, and who is being judged. Notably, the fact that clusters are discernible (that is, modularity is high in the com-

munity detection algorithm) only when decomposed by valence and alter means that no overarching structure links the various maps together.

We also compare the importance of various demographic factors in predicting moral logics. We showed that gender, rather than race, class, or political identification, best predicted membership across the moral logics inductively detected through shared judgment. Other demographic categories were relevant only in specific contexts. For instance, even conditioning on income, gender, race, and other demographic characteristics, identifying as a Democrat was the strongest predictor of membership in a community that judges individuals for prosociality. Negative judgments of institutions were structured primarily by race, especially identifying as black. Relative to other racial groups, black individuals were more likely to judge institutions as disorderly, incompetent, and violent.

Our findings offer important implications for two ongoing theoretical conversations. First, consistent with social psychological research on stereotype content, we found that people cluster in how they express positive judgments toward individuals, with most relying either on prosociality or competence, which was also reflected in our factor analysis. Yet we also found that the standards by which Americans expressed judgment had a different community structure when they were expressing positive or negative judgments, and when they were judging institutional actors or individuals. One avenue for stereotype content research to explore, then, is whether different judgments emerge when priming survey respondents to think about stereotypically positive and negative judgments separately. Another avenue to explore is stereotypical judgments of institutional actors—rather than demographic groups.

Second, to qualitative studies on moral logics, our results suggest the presence of multiple topologies. People rely on different sets of moral logics, depending on circumstances like who is being judged and by what valence. Like many other researchers, our interest in identifying moral logics was motivated by its salience in structuring categories. Moral logics are im-



portant sources of categorical distinction, such as when white working-class men express moral logics of self-discipline and categorize black working-class men, who are negatively judged as lacking self-discipline, as being different (Lamont 2000). Yet a finding of multiplicity is consistent with theories that liken moral logics to toolkits that people deploy, rather than schemata that dictate attitudes (Swidler 1986). People may primarily judge for competence or prosociality, but this is situational. When it comes to judgments of institutional actors, they may draw on alternative standards like empathy; when it comes to negative judgments, people draw on moral logics that are more complex mixtures of various moral standards. Multiplicity matters because it implies the possibility that cleavages can be cut. White and black working-class men may very well share moral logics when judging institutional actors such as their employers. Additionally, we were surprised to not find a moral logic cohering around standards associated with individuality, given strong claims about the first language of Americans being individuality. We conducted additional robustness tests to investigate whether these nonresults were simply because of irregularities in coding structure. Factor analysis, however, revealed that prosociality and competence are indeed still the primary factors.

To be clear, the absence of evidence is not definitive evidence of absence. Perhaps people do not judge for diligence and individuality unless they are explicitly asked to make moral judgments about others, which American Voices Project protocol did not do. Perhaps people are more likely to attempt to convey a consistent moral logic across various situations when explicitly asked about how they judge, but when making judgments without such prompts, people are more likely to demonstrate multiplicity. We would only be able to conclude that our findings definitively challenge prior work only after these (and other) possibilities are ruled out, but the surprising lack of findings warrants further research.

Finally, where we expected all demographic categories to emerge as important dividing lines for moral logics, we found gender func-

tioned as a general dividing line and that the other categories were situational in their importance. These differences are likely attributable to enduring differences in socialization and life outcomes by gender (Risman 2004; Goldin 2014; England 2017). For instance, living with an abusive spouse conditions the judgments of many women to be about violence (indeed, domestic abuse was a recurring theme in the interview transcripts). These material differences in life experiences, as well as differences in socialization, shapes how people recount their lives and the standards by which they judge. Given prior work, we expected class to emerge as a strong predictor of moral logics. We conjecture that class might have emerged as a more general dividing line if our operationalization for class (falling below the poverty line) was more comprehensive. Because this operationalization is comparatively thin, it may explain why this demographic characteristic does not emerge as a strong predictor of moral logics. More generally, even though the magnitude of demographic differences discussed in our analyses are substantively significant (for example, women are more than twice as likely to be observed to judge individuals for violent behaviors), the assessment of demographic correlates within each community is likely to be underpowered.

Aside from helping address the above limitations in statistical precision and power, an expanded sample would enable study of finer-grained distinctions in standards people use to describe specific alters, such as differences in how pairs of alters are judged, such as between mothers versus fathers, those who are black versus white, or politicians in Democratic versus Republican parties. Additionally, we studied how these communities varied in terms of race, gender, income, and political identification—but these are hardly exhaustive. We encourage more deductive research that preregisters hypotheses about dimensions such as religious affiliation (for example, are self-identifying Christians more or less likely to judge, who do they judge, and upon what standards), age (we found that older Americans are more likely to judge, but do they differ in how?), region (are there distinctive modes of judgment unique to certain spatially segregated commu-

nities in America?), or occupation (does becoming a police officer imply spillover effects on how someone judges their family?).

These possibilities for further work underscore the methodological contribution of this work. Where qualitative research offers the appropriate level of thick description and understanding of individuals to identify moral logics, it lacks the scale of survey-based research to sample a population and identify how standards of judgment are shared more broadly. By contrast, by explicitly asking for judgments of groups, survey-based research lacks the naturalism and depth afforded by ethnographic or interview-based research. Using network representations of shared judgment through inductively coded interview transcripts of a diverse sample of Americans is a way to maintain some of the strengths of both approaches while reducing their respective weaknesses. Like other studies in this issue (Sauder, Shi, and Lynn 2024; Zilberstein et al. 2024), which have

relied on this dataset to inductively identify various ways that Americans understand concepts like luck or agency, this article draws on the omnibus nature of a narrative interview to maintain greater analytic depth at scale. Here, our results reveal the promise of large-scale qualitative data collection, such as that in the American Voices Project, in elucidating a fundamental basis of distinctions and categories in American society. If how people express judgment makes and maintains boundaries, then the fact that Americans are segregated into distinctive moral communities implies a key process that makes material inequalities durable in the long durée. Yet, as our network analysis reveals, Americans also share multiple sets of moral logics that are drawn upon in varying situations. The moral repertoires of Americans are diverse and situational. Hence, even if group boundaries are produced from perceived differences in moral logics, these boundaries are malleable.

**Table A.1.** Summary of Coding Structure for Targets of Judgment (Alters)

Alter Group	Alters	Examples
<b>Individuals</b>	Parents	mother, father, stepmother, stepfather
	Children	daughter, son, stepdaughter, stepson
	Spouse	wife, husband, ex-wife, ex-husband
	other family	grandparents, grandchildren, siblings, in-laws, other relatives
	romantic relationship	boyfriend, fiancé, girlfriend, fiancée, ex-boyfriend, ex-girlfriend
	Acquaintances people at work	friend, acquaintance coworker
<b>Institutional actors</b>	justice system	police officers, judges
	education system	teachers, educators
	religious system	religious leaders, pastors, god
	health-care system	doctors, nurses, medical system
	other social system	landlord, financial system
	political groups	Republicans, Democrats, specified political movement and ideology (such as BLM, Racism, LGBTQ)
	Neighborhood demographic groups people or society in general	neighborhood actors (such as HOA) by race, by gender, by sexuality, and so on

Source: Authors' codings.

**Table A.2.** Standards of Judgment, with Example Quotes

Variable (Common Keywords)	Mean $\pm$ S.E.	Example Quote
Number of total positive judgments	3.730 $\pm$ 0.186	
Prosocial (caring, generous, supportive)	1.245 $\pm$ 0.080	I got good people, I got good neighbors around me. We watch and we look out and support each other
Competent (skilled, talented, smart, strong)	0.868 $\pm$ 0.070	She's also very smart, she has always had a 4.0 GPA. Always been I would say brought up for classes, grade wise always has been talented in academics
Empathetic (welcoming, tolerant)	0.321 $\pm$ 0.036	The police were wonderful, extremely empathetic and they told me what to do, and they told me to get cameras and take pictures
Diligent (responsible, hard-working)	0.290 $\pm$ 0.034	Because my parents were hard working. They started their lives and careers, particularly in this country later than most do, and are definitely both by the bootstraps up people and come from humble origins.
Special (creative, inspired, fun, cool)	0.251 $\pm$ 0.030	But my mom, I look up to her a lot. Because she just has, she's always been really guided by what inspires her and what she wants to do in life. So especially since me and my sister moved out of the house like after she turned 50, she was having the best time of her life. . . . So, she's definitely an inspiration for me for sure.
Sociable (many friends, popular, well connected)	0.228 $\pm$ 0.028	He's amazing, he got a lot of friends, too many friends but got to love him, always got to have somebody around but not lately because of everything that's going on
Safe (protecting, tame, gentle)	0.197 $\pm$ 0.024	But the law enforcement I feel that they're there if I need them, I feel protected with them and they're like an insurance policy until you really need them that they don't exist, but when you need them, you need them. . . . With law enforcement I respect them and I know I'm protected by them
Healthy or physically attractive (beautiful, gorgeous, attractive)	0.166 $\pm$ 0.028	When I was working for that body shop in [PERSONAL NAME] is when I met this girl, just I think, at that point, I was already, you know, kind of feeling a little lonely and wanting something else in my life, and she was really beautiful
Respectable (polite, civil)	0.118 $\pm$ 0.019	I don't have a problem with law enforcement. I really like law enforcement. I welcome them into my life at any time, because [PERSONAL NAME] a very respectable person of law enforcement
Rich (wealthy, upscale)	0.020 $\pm$ 0.007	I've never really had a problem with the last president, President Trump. I mean, I figured since he was or is a billionaire what can it hurt.
Number of total negative judgments	5.637 $\pm$ 0.376	

*(continued)*

**Table A.2.** (continued)

Variable (Common Keywords)	Mean $\pm$ S.E.	Example Quote
Incompetent (idiot, weak, dumb)	0.758 $\pm$ 0.077	I ended up getting married to an idiot.
Violent (cruel, abusive)	0.690 $\pm$ 0.070	Hell, no, they trying to kill us, like erasing all the black people in the street, so that they know who is making things like they can't just go out and kill us, because they know we still got a system, a justice system here and that justice system still go, you still go to go through it, even if you don't right and get through it, you still got to go through it.
Out of control (unpredictable, unstable, reckless)	0.673 $\pm$ 0.060	But she has a substance problem. Not every day, but every once in a while. And she's unpredictable. You never know when that's coming. And so, I am doing everything that I have always . . . You know, supported her, always try to get her to address whatever issues she has, but she can be a problem from time to time. It is not drugs, to my knowledge anyway, it is alcohol.
Unfair (unjust, unfair, double standard)	0.656 $\pm$ 0.064	There are ways to get into the country legally, but during this time, if you don't have the resources to do that, illegally is the only way. I do understand that, too. But for a person to be afraid like, "Oh my goodness, if I get caught, then I'll be sent home," that's not fair. Let that person to get some kind of chance to explain it, "Hey, this is what's going on. I would like to get that chance to become a legal resident here, but I don't have the resources."
Dishonest (crooked, liar, shady)	0.566 $\pm$ 0.063	And then when she was [young], she got taken away by the state. The [INAUDIBLE] here is like really crooked. My daughter dropped a knife . . . and it was just a little pinprick, and hardly even bled at all, and because I didn't take her to the doctor, they said that it was neglect and they took her away . . . after that happened.
Disorderly (illegal, racist, nasty, improper)	0.538 $\pm$ 0.060	The president is such an abysmal human being. Okay. And how he would degrade President Obama and demean him say he's from Kenya and all that nonsense. How can you not become more left leaning when he's labeling these Hispanics are all rapists and murderers, come on already, but also to his horrible just horrendous history with women, and abusing women.
Exclusive (intolerant, judgmental)	0.535 $\pm$ 0.052	I remember reading something about the correlation with educational level and religion, and like it decreases -- and I can definitely see that because -- yeah, I did have a lot of problems with some of the teachings. Everybody always says, "That's not my church." I get it. [PERSONAL NAME] But overall, I can't agree with some of these practices. I just find some of them to be very exclusionary.

**Table A.2.** (continued)

Variable (Common Keywords)	Mean ± S.E.	Example Quote
Selfish (uncaring, greedy, un- supportive)	0.530 ± 0.053	You know, the doctors, all they are pill pushers, you know, they want to give you some kind of pill. You got one problem and they want to give you a pill that will cause us thirty or forty problems, so that way they can give you thirty or forty more pills for each one of that problems, and none of them fix it or make it go away. . . . They want to have you come back and do a dual-office visit that's totally meaningless, except to pay their Mercedes Benz payment.
Impolite (spoiled, crude)	0.369 ± 0.042	I will say most of the customers, they tend to have a lot of attitude you know, as employees we can't. We can't, we have to be nice. But it's just those customers, if you let them know like let's say, why is your EDT machine not working. I'm like, "Well, I don't know. It's the system. It's our system." And they start yelling at you so I'm like, "Well, I just work here." So there's a lot of customers that are like that.
Unhealthy or physically unattractive (cross, ugly)	0.155 ± 0.024	I was gross back then I guess. In my eyes, I was gross. I wore baggy clothes, apparently, I hated showering.
Lazy (deadbeat, bum)	0.152 ± 0.028	He did like all the rest of the deadbeat dads do. They'll get a job and as soon as Child Support catch up with them, they quit the job and go get another job. He was doing that kind of stuff but he was also going to school.
Poor	0.039 ± 0.010	He also, was scraping pennies together. He was a server at [PERSONAL NAME], so he didn't even have forks. He didn't even have plates; paper towels, toilet paper, he just lived this incredible life of squalor. And, I was like, "Well, I have money. So, I can pay for this. I can get your house in order." I kind of like, strung him along for a little bit, and then it got too much to deal with.

Source: Authors' codings and calculations.

## REFERENCES

- Abele, Andrea E., and Bogdan Wojciszke. 2007a. "Agency and Communion from the Perspective of Self Versus Others." *Journal of Personality and Social Psychology* 93(5): 751–63.
- . 2007b. "Communal and Agentic Content in Social Cognition: A Dual Perspective Model." In *Advances in Experimental Social Psychology*, vol. 50, edited by James M. Olson and Mark P. Zanna, 195–255. Waltham, Mass.: Academic Press.
- Bellah, Robert N. 2007. "Reading and Misreading 'Habits of the Heart.'" *Sociology of Religion* 68(2): 189–93.
- Bellah, Robert N., Stephen M. Tipton, William M. Sullivan, Ann Swidler, and Richard Madsen. 2008. *Habits of the Heart: Individualism and Commitment in American Life*. Berkeley: University of California Press.
- Bourdieu, Pierre. 1984. *Distinction: A Social Critique of the Judgement of Taste*. Cambridge, Mass.: Harvard University Press.
- . 2000. *Pacalain meditations*. Redwood City, Calif.: Stanford University Press.
- Bourdieu, Pierre, and Jean-Claude Passeron. 1977. *Reproduction in Education, Society and Culture*. Thousand Oaks, Calif.: Sage Publications.
- Boutyline, Andrei, and Laura K. Soter. 2021. "Cultural Schemas: What They Are, How to Find Them, and What to Do Once You've Caught One." *American Sociological Review*, 86(4), 728–758.
- Boutyline, Andrei, and Stephen Vaisey. 2017. "Belief Network Analysis: A Relational Approach to Understanding the Structure of Attitudes." *American Journal of Sociology* 122(5): 1371–447.
- Calarco, Jessica McCrory. 2014. "Coached for the Classroom: Parents' Cultural Transmission and Children's Reproduction of Educational Inequalities." *American Sociological Review* 79(5): 1015–37. <https://doi.org/10.1177/0003122414546931>.
- Campbell, Colin. 1998. *The Myth of Social Action*. Cambridge: Cambridge University Press.
- Chen, M. Keith, and Ryne Rohla. 2018. "The Effect of Partisanship and Political Advertising on Close Family Ties." *Science* 360(6392): 1020–24. <https://doi.org/10.1126/science.aaq1433>.
- D'Andrade, Roy G. 1995. *The development of cognitive anthropology*. Cambridge University Press.
- DellaPosta, Daniel, Yongren Shi, and Michael Macy. 2015. "Why Do Liberals Drink Lattes?" *American Journal of Sociology* 120(5): 1473–511.
- DiMaggio, Paul. 2002. "Why Cognitive (and Cultural) Sociology Needs Cognitive Psychology." In *Culture in Mind: Toward a Sociology of Culture and Cognition*, edited by Karen A. Cerula, 274–82. New York: Routledge.
- . 2011. "Cultural Networks." In *The Sage Handbook of Social Network Analysis*, edited by John Scott and Peter J. Carrington, 286–310. London: Sage Publications
- Edin, Kathryn J., Corey D. Fields, David B. Grusky, Jure Leskovec, Marybeth J. Mattingly, Kristen Olson, and Charles Varner. 2024. "Listening to the Voices of America." *RSF: The Russell Sage Foundation Journal of the Social Sciences* 10(5): 1–31. <https://doi.org/10.7758/2024.10.5.01>.
- England, Paula. 2017. *Households, Employment, and Gender: A Social, Economic, and Demographic View*. New York: Routledge.
- Feinberg, Matthew, and Robb Willer. 2019. "Moral Reframing: A Technique for Effective and Persuasive Communication Across Political Divides." *Social and Personality Psychology Compass* 13(12): e12501.
- Finkel, Eli J., Christopher A. Bail, Mina Cikara, Peter H. Ditto, Shanto Iyengar, Samara Kiar, Lilliana Mason, et al. 2020. "Political Sectarianism in America." *Science* 370(6516): 533–36.
- Fiske, Susan T. 2018. "Stereotype Content: Warmth and Competence Endure." *Current Directions in Psychological Science* 27(2) :67–73. <https://doi.org/10.1177/0963721417738825>.
- Fiske, Susan T., Amy J. C. Cuddy, and Peter Glick. 2007. "Universal Dimensions of Social Cognition: Warmth and Competence." *Trends in Cognitive Sciences* 11(2): 77–83.
- Fiske, Susan T., Amy J. C. Cuddy, Peter Glick, and Jun Xu. 2002. "A Model of (Often Mixed) Stereotype Content: Competence and Warmth Respectively Follow from Perceived Status and Competition." *Journal of Personality and Social Psychology* 82(6): 878. <https://doi.org/10.1037/0022-3514.82.6.878>.
- Frost, Jacqui, and Penny Edgell. 2022. "'I Believe in Taking Care of People': Pushing Back Against Rationalized Institutions with a Logic of Care." *Poetics* 90: 101593.
- Gast, Melanie Jones, and Dina G. Okamoto. 2016. "Moral Or Civic Ties? Deservingness and Engagement Among Undocumented Latinas in Non-Profit Organisations." *Journal of Ethnic and Migration Studies* 42(12): 2013–30.

- Goldberg, Amir. 2011. "Mapping Shared Understandings Using Relational Class Analysis: The Case of the Cultural Omnivore Reexamined." *American Journal of Sociology* 116(5): 1397-436.
- Goldin, Claudia. 2014. "A Grand Gender Convergence: Its Last Chapter." *American Economic Review* 104(4): 1091-119.
- Graham, Jesse, Jonathan Haidt, and Brian A. Nosek. 2009. "Liberals and Conservatives Rely on Different Sets of Moral Foundations." *Journal of Personality and Social Psychology* 96(5): 1029. <https://doi.org/10.1037/a0015141>.
- Gusfield, Joseph R. 1986. *Symbolic Crusade: Status Politics And The American Temperance Movement*. Champaign: University of Illinois Press.
- Hechter, Michael, James Ranger-Moore, Guillermina Jasso, and Christine Home. 1999. "Do Values Matter? An Analysis of Advance Directives for Medical Treatment." *European Sociological Review* 15(4): 405-30.
- Jack, Anthony Abraham. 2019. *The Privileged Poor*. Cambridge, Mass.: Harvard University Press.
- Kaufman, Jason. 2004. "Endogenous Explanation in the Sociology of Culture." *Annual Review of Sociology* 30(1): 335-57.
- Khan, Shamus. 2011. *Privilege: The Making of Adolescent Elite At St. Paul's School*. Princeton, N.J.: Princeton University Press.
- Koch, Alex, Roland Imhoff, Ron Dotsch, Christian Unkelbach, and Hans Alves. 2016. "The ABC of Stereotypes About Groups: Agency/Socioeconomic Success, Conservative-Progressive Beliefs, and Communion." *Journal of Personality and Social Psychology* 110(5): 675. <https://doi.org/10.1037/pspa0000046>.
- Lamont, Michèle. 2000. *The Dignity of Working Men: Morality and the Boundaries of Race, Class, and Immigration*. Cambridge, Mass.: Harvard University Press.
- Lamont, Michèle, and Virág Molnár. 2002. "The Study of Boundaries in the Social Sciences." *Annual Review of Sociology* 28(1): 167-95. <https://doi.org/10.1146/annurev.soc.28.110601.141107>.
- Lareau, Annette. 2011. *Unequal Childhood: Class, Race, and Family Life, with an Update a Decade Later*. Berkeley: University of California Press.
- Lassetter, Bethany, Eric Hehman, and Rebecca Neel. 2021. "The Relevance Appraisal Matrix: Evaluating Others' Relevance." *Journal of Personality and Social Psychology* 121(4): 842-64. <https://doi.org/10.1037/pspi0000359>.
- Lipset, Seymour M., and Gary Marks. 2001. *It Didn't Happen Here: Why Socialism Failed in the United States*. New York: W. W. Norton.
- Martin, John L. 2002. "Power, Authority, and the Constraint of Belief Systems." *American Journal of Sociology* 107(4): 861-904.
- Massey, Douglas S. 2007. *Categorically Unequal: The American Stratification System*. New York: Russell Sage Foundation.
- Mills, C. Wright. 1940. "Situated Actions and Vocabularies of Motive." *American Sociological Review* 5(6): 904-13. <https://doi.org/10.2307/2084524>.
- Mohr, John W. 1998. "Measuring Meaning Structures." *Annual Review of Sociology* 24(1): 345-70.
- Piff, Paul K., Michael W. Kraus, Stephane Cote, Bonnie Hayden Cheng, and Dacher Keltner. 2010. "Having Less, Giving More: The Influence of Social Class on Prosocial Behavior." *Journal of Personality and Social Psychology*. 99(5): 771-84. <http://doi.org/10.1037/a0020092>.
- Ridgeway, Cecilia. L. 2011. *Framed by Gender: How Gender Inequality Persists in the Modern World*. Oxford: Oxford University Press.
- Risman, Barbara. J. 2004. "Gender as a Social Structure: Theory Wrestling with Activism." *Gender and Society* 18(4): 429-59.
- Rivera, Lauren A. 2016. *Pedigree: How Elite Students Get Elite Jobs*. Princeton, N.J.: Princeton University Press.
- Rosenberg, Seymour, Carnot Nelson, and P. S. Vivekananthan. 1968. "A Multidimensional Approach to the Structure of Personality Impression." *Journal of Personality and Social Psychology* 9(4): 283-94.
- Sauder, Michael, Yongren Shi, and Freda Lynn. 2024. "Multiple Meritocracies: A Text-Based Analysis of Personal Narratives Revealing Distinct Frames of Success." *RSF: The Russell Sage Foundation Journal of the Social Sciences* 10(5): 87-117. <https://doi.org/10.7758/RSF.2024.10.5.04>.
- Sayer, Andrew. 2005. *The Moral Significance of Class*. Cambridge: Cambridge University Press.
- Schwartz, Shalom H., and Tammy Rubel. 2005. "Sex Differences in Value Priorities: Cross-Cultural and Multimethod Studies." *Journal of Personality and Social Psychology* 89(6): 1010.
- Swidler, Ann. 1986. *Talk of Love: How Culture Matters*. Chicago: University of Chicago Press.
- Thai, Mai. 2022. "Policing and Symbolic Control." *American Journal of Sociology* 127(4): 1183-220.
- Thornton, Patricia H., William Ocasio, and Michael

- Lounsbury. 2012. *The Institutional Logics Perspective: A New Approach to Culture, Structure and Process*. Oxford: Oxford University Press.
- Tilly, Charles. 1998. *Durable Inequality*. Berkeley: University of California Press.
- Tomaskovic-Devey, Donald, and Dustin Avent-Holt. 2019. *Relational Inequalities: An Organizational Approach*. New York: Oxford University Press.
- Traag, Vincent A., Ludo Waltman, and Nees Jan van Eck. 2019. "From Louvain to Leiden: Guaranteeing Well-Connected Communities." *Scientific Reports* 9(1) 5233.
- Vaisey, Stephen. 2009. "Motivation and Justification: A Dual-Process Model of Culture in Action." *American Journal of Sociology* 114(6): 1675–715.
- Valentino, Lauren, and Stephen Vaisey. 2022. "Culture and Durable Inequality." *Annual Review of Sociology* 48(1): 109–29. <https://doi.org/10.1146/annurev-soc-030320-102739>.
- Zou, Linda X., and Sapna Cheryan. 2017. "Two Axes of Subordination: A New Model of Racial Position." *Journal of Personality and Social Psychology* 112(5): 696–717. <https://doi.org/10.1037/pspa0000080>.
- Zilberstein, Shira, Elena Ayala-Hurtado, Mari Sanchez, and Derek Robey. 2024 "The Self in Action: Narrating Agentic Moments." *RSF: The Russell Sage Foundation Journal of the Social Sciences* 10(5): 118–40. <https://doi.org/10.7758/RSF.2024.10.5.05>.