

The Moral Foundations of Individual-Level Foreign Policy Preferences

David Lindsey* and David Lake†

January 12, 2014

Abstract

Individuals differ in their foreign policy preferences. Some support higher defense spending and more aggressive stances toward other countries, others oppose the use of force except in the most extreme circumstances. Some prefer working through multilateral institutions, others believe it is necessary to maintain full freedom of action on all foreign policy issues. Social psychology posits that individuals hold varying moral foundations, which they group in five dimensions of care/harm, fairness/cheating, loyalty/betrayal, authority/submission, and sanctity/degradation. Liberals tend to emphasize care/harm and fairness/cheating disproportionately, while conservatives appear to weight each foundation more equally. Drawing on this work, we first replicate the association between the moral foundations and political identity and domestic policy preferences. We then extend this research to foreign policy issues, finding significant relationships between moral foundations and attitudes on a variety of policy questions. We also conduct two survey experiments that prime respondents with different moral foundations and find no significant effect on attitudes. We postulate that moral concerns are already salient in the minds of respondents when they answer questions about foreign policy.

*Ph.D. student, Department of Political Science, University of California, San Diego, La Jolla, CA 92093-0521. Email: dlindsey@ucsd.edu

†Gerri-Ann and Gary E. Jacobs Professor of Social Sciences and Distinguished Professor of Political Science, University of California, San Diego, La Jolla, CA 92093-0521. Email: dlake@ucsd.edu

Why do individuals differ in their foreign policy preferences? Even casual observation suggests that people, even within the same country, hold widely varying attitudes toward the use of force, the United Nations, globalization, and human rights. In this paper, we find that foreign policy attitudes are rooted, in part, in differentially-held moral foundations that are themselves believed to be relatively stable and enduring at the individual level.

Haidt (2012) and collaborators have found that individuals hold five moral foundations in varying degrees: care/harm, fairness/cheating, loyalty/betrayal, authority/submission, and sanctity/degradation, which we explain in more detail below. Liberals in the United States tend to emphasize care/harm and fairness/cheating, while conservatives weight each of the five foundations more equally. These foundations have been found to predict political positions on domestic issues (Koleva et al., 2012). In surveys of US respondents, we extend these foundations to foreign policy attitudes and find similar associations.

We do not claim that these moral foundations are the sole determinant of foreign policy attitudes. Opinions held by individuals have many causes. Nonetheless, these moral foundations are strongly related to attitudes on a large range of foreign policy issues, even controlling for self-described ideology and demographic traits. The relationships between specific moral foundations and particular foreign policy issues will likely surprise few readers. If one accepts the foundations, their consequences for foreign policy preferences are relatively straightforward. It is hardly counter-intuitive that individuals concerned with care/harm will want to protect weaker nations from foreign aggression or that those scoring high on sanctity/degradation will agree that their culture needs protection from foreign influence. What is surprising, however, is that just a few moral foundations can predict attitudes across such a large number of not only domestic but also foreign policy dimensions, including issues of both security and political economy.

The so-called Cold War consensus on foreign policy attitudes held that public opinion was capricious, inchoate, and even dangerous (Almond, 1950; Lippmann, 1955; Campbell et al., 1960; Converse, 1964). This view intersected with and supported theories of international relations that assumed politics stopped “at the water’s edge” and relatively autonomous statesmen (the gendered noun is intentional) were free to pursue what they regarded as the “national interest.” This consensus was shattered by the Vietnam War, which demonstrated beyond doubt that the public clearly held differing opinions over foreign policy, expressed their disagreements forcefully, and ultimately constrained policy (Verba et al., 1967; Mueller, 1973; Holsti, 1996). We now know that public opinion on foreign policy is a lot like opinion on any political topic. It is relatively coherent and, at least in the aggregate, relatively stable (Page and Shapiro, 1992).¹ In most cases, it is also “pretty prudent” or sensible given the issues of the day (Jentleson, 1992; Nincic, 1992). Reflecting this structure, and focusing mostly on security issues, scholarly consensus posits that attitudes vary along two core dimensions: militant internationalism (MI), defined by views on the use of force, and cooperative internationalism (CI), a cluster of views supporting international institutions and world order. These two dimensions, in turn, combine to create four approaches to foreign policy: hard-liners (support MI, oppose CI), accommodationists (oppose MI, support CI), internationalists (support both), and isolationists (oppose both) (Holsti, 1996). Although the questions that comprise these dimensions have changed over time, the four types continue to describe broad approaches to US foreign policy even after the Cold War. Although most of the time these foreign policy attitudes are latent, opinion can be activated by crises, especially war casualties (Mueller, 1973; Gartner, 1997; Gelpi, Feaver and Reifler, 2006; Berinsky and Druckman, 2007), and shaped by elite opinion (Zaller, 1992), the media (Powlick and Katz, 1998; Baum, 2003; Baum and Groeling,

¹See also Russett and Deluca (1981).

2009), and other cues. Once activated, moreover, the public can constrain the choices of national leaders – and leaders may even shape policy in anticipation of public opinion becoming activated in various instances.² In short, we now believe foreign policy attitudes are both meaningful and that they matter (Aldrich et al., 2006; Baum and Groeling, 2008).

Although we can identify enduring types of individuals with clusters of foreign policy attitudes, we still lack strong theories or explanations of where these schools “come from” or why individuals hold specific attitudes. The types are induced from responses to a range of questions on foreign policy and, even though they support the idea of coherent belief structures, they cannot by themselves be used to explain the underlying attitudes. Instead, scholars have attempted to explain individual attitudes by age cohort, gender, and other demographic traits or by ideology and partisanship, but whether the latter are causes or effects remains unclear (Holsti, 1996). The closest approach in foreign policy studies to ours is the hierarchy of attitudes model of Hurwitz and Peffley (1987), which derives foreign policy attitudes from underlying principles of ethnocentrism and the morality of warfare. Related studies examine the link between foreign policy attitudes and personality traits like aggression and accommodation (Herrmann, Tetlock and Visser, 1999; Herrmann, Tetlock and Diascro, 2001), distrust (in general, and between Americans and others) (Bartels, 1994; Brewer and Steenbergen, 2002; Binning, 2007), or preferences over other-regarding and positive inducements (Nincic and Ramos, 2010).³ Our moral foundations are more clearly exogenous to foreign policy attitudes and more general, since the same foundations have been shown to co-vary with a range of domestic policy attitudes as well. We also show that these moral foundations are related to both attitudes toward security policy and foreign economic policy.

²On rally effects, see Baker and Oneal (2001) or Chapman and Reiter (2004).

³On personality traits and political attitudes in general, see Gerber et al. (2011).

In international political economy, scholars have taken a substantially different tack to explaining attitudes than in the general foreign policy literature. Unlike security issues, on which we might expect – in theory – people to have relatively homogenous views, economic policies have more clearly identifiable distributional implications. Although the most useful unit of analysis – factors, sectors, or firms – remains debated, economic theory predicts that economic openness in goods or factor markets will create both relative winners and losers, and that these actors will seek to obtain policies favorable to their interests in the political arena (see Frieden and Rogowski 1996). Even as there is evidence that such predictions hold at the aggregate or group level, manifested in the ability of theories to explain actual patterns of economic openness and the surrounding political struggles, there is only modest evidence that the predictions hold at the level of individuals (Scheve and Slaughter, 2001; Kaltenhaler, Gelleny and Ceccoli, 2004; Rho and Tomz, 2012).⁴ It is increasingly clear that individuals do not define their policy preferences strictly in terms of their positions as producers within the international division of labor. Rather, they appear to have more complex policy preferences derived from their roles as consumers (Baker, 2005; Naoi and Kume, 2011), sociotropic goals like full employment or equality (Scheve and Slaughter, 2006; Mansfield and Mutz, 2009; Hainmueller, Hiscox and Margalit, 2011), religiosity (Scheve and Stasavage, 2006), gender (Goldstein, Margalit and Rivers, 2007), identity (Rankin, 2001), or ethnocentrism (Sabet, 2012). In some ways, in moving beyond economic determinants, the political economy literature is beginning to converge with the literature on foreign policy attitudes more generally. Even here, however, a large portion of the variance in individual attitudes remains unexplained.

The goal of this paper is to improve our understanding of individual-level attitudes toward a range of foreign policy issues by building from a common core of moral foundations. Part I

⁴But see Fordham (2008) and Fordham and Kleinberg (2012).

lays out moral foundation theory (MFT) and addresses several common objections. Although we cannot reject these criticisms, we argue that they do not undermine our results. Part II reports on an online survey conducted in Fall 2012, correlating the moral foundations with a battery of foreign policy questions drawn from other frequently used surveys. We do not innovate on either the moral foundation questionnaire (MFQ) or the foreign policy questions themselves, but rather use these tools precisely for their comparability with other studies. Part III presents results on two sets of survey experiments conducted in January-March 2013. In these experiments, we prime respondents on the care/harm or loyalty/betrayal foundation and then repeat the battery of foreign policy questions. Although our respondents “took” or absorbed the primes, the priming had no systematic effect on foreign policy attitudes, including for those questions that were highly related to care/harm and loyalty/betrayal in the previous survey. This suggests that the moral foundations are relatively stable and likely present in the minds of respondents when answering questions even without our manipulation.

1 The Moral Foundations

Haidt and his co-authors have identified five moral foundations (Haidt and Joseph, 2004; Haidt and Graham, 2007; Haidt and Joseph, 2007). These foundations were originally derived inductively from a survey of cultures and human practices as the “best candidates for universal cognitive modules upon which cultures construct moral matrices” (Haidt 2012, 124).⁵ The foundations are understood to be intuitive, nearly automatic responses by individuals to situations that are robust to

⁵Haidt acknowledges his debt to and the overlap of these foundations with those identified by Shweder (1991) and Fiske (see Rai and Fiske 2011). We do not assess the universality of these foundations as our sample is limited to the U.S.

rationality and justification, although initial reactions based on a foundation may be overridden by conscious thought and reasoning.⁶ When confronted with an example of incest between consenting adults, for instance, nearly everyone reacts with disgust (violation of sanctity/degradation), even though many cannot explain why; most are unpersuaded by arguments that no one was harmed in the encounter, but some may reason their way to begrudging acceptance.⁷ From this, Haidt (2012, 131) argues that moral foundations are “organized in advance of experience.”

The five foundations (in their positive and negative valences) are, briefly:⁸

- **Care/harm:** characterized by the emotion of compassion and the virtues of caring and kindness toward others.
- **Fairness/cheating:** concerns equality and proportionality in social relations, violations of which provoke anger (by victims) and guilt (by perpetrators); associated with virtues of justice and trustworthiness.
- **Loyalty/betrayal:** manifested in feeling of attachment and pride in one’s “in group” and fear of any “out group,” even though groups may be arbitrary associations (such as sports teams); characterized by patriotism and self-sacrifice toward group members, and rage and anger toward those who “defect” from the group.
- **Authority/submission:** distinguished by the emotions of respect and fear towards the authority as an institution or those in authority as individuals and the virtues of obedience and deference; unlike the other foundations, which appear bounded by neutrality at one end of the continuum, some individuals may have an active aversion to authority, in general, or to some authorities, in particular.
- **Sanctity/degradation:** disgust towards certain objects or acts that are perceived as “unclean” or taboo, offset by the virtues of temperance, chastity, piety, or cleanliness.

⁶See also Kahneman (2011) and Tetlock (2006).

⁷A common technique used by Haidt and his associates is to challenge initial reactions to violations of moral foundations with reasoned counter-arguments. Most subjects experience “moral dumbfounding” in which they cannot counter the counter-arguments but still believe the violation is morally wrong. Haidt (2012, 25, 38-40).

⁸The labels for each have varied over time; we use the labels from Haidt (2012). In the book, Haidt poses a sixth foundation of liberty/oppression. Although he is possibly tapping into another dimension here, this sixth dimension overlaps considerably with both fairness/cheating and authority/submission, muddling the framework and, in our view, highlighting the first problem identified by critics of moral foundations theory (see below). The moral foundation questionnaire does not yet contain questions (and thus scores) on liberty/oppression, so we do not explore this dimension here.

While the particular meaning of each foundation as well as “triggers” that evoke the associated emotions vary by culture, time, and context, Haidt claims that all individuals respond at least in part to all five foundations. Haidt measures the moral foundations through responses to a battery of questions on the Moral Foundations Questionnaire (MFQ), available online and reproduced in the appendix to this paper.⁹ For our purposes, the important point is that none of the questions on the MFQ are “political,” though some of the authority questions are phrased in general terms about government and law, and none address specific policy issues.

Figure 1 displays the correlations among the 30 questions on the MFQ by our respondents in the primary survey (described below). The questions are labeled down the main diagonal with the letter corresponding to the foundation they measure (“H” for care/harm, “F” for fairness/cheating, “L” for loyalty/betrayal, “A” for authority/submission, and “S” for sanctity/degradation). Boxes shaded in blue indicate that the Spearman’s rank correlation coefficient between two questions is positive; red indicates a negative correlation. Darker shading indicates values further from zero. Thus, dark shading down the main diagonal suggests that questions for the same foundation are similar (convergent validity), while light shading far from the diagonal suggests that questions measuring different foundations are not similar (discriminant validity). We see this general pattern below. The figure also makes clear that the five foundations form two natural clusters (large blue sections in the upper left and bottom right). These two clusters are known as the “individualizing” (care/harm and fairness/cheating) and “binding” (loyalty/betrayal, authority/submission, and sanctity/degradation) foundations.

[Figure 1 About Here]

⁹Available at: <http://www.yourmorals.org/explore.php> (accessed Nov. 4, 2013).

For Political Scientists, one of the most interesting patterns to emerge from the MFQ is that self-identified liberals score higher on the care/harm and fairness/cheating foundations, and lower on the others, while self-described conservatives score more evenly across all five foundations.¹⁰ Our convenience sample, described in the next section, replicates this core finding (compare Figure 2 to Figure 8.1 in Haidt 2012, 158). The two clusters of foundations in Figure 1 are essentially reproduced in this graph, with care/harm and fairness/cheating taking very similar values on average as ideology moves from left to right.¹¹

[Figure 2 About Here]

Importantly, moral foundations are also closely associated with a range of policy issues (Koleva et al., 2012). Even controlling for demographic attributes, ideology, religious attendance, and sociopolitical beliefs, moral foundations are significant predictors of attitudes toward abortion, same-sex marriage, torture, teaching creationism, gun control, flag burning, stem cell research, and other public policy concerns, including some international issues like global warming, defense spending, and terrorism. It is precisely this predictive power that we wish to extend to foreign policy preferences, not only explaining foreign policy attitudes more fully but connecting individual views on international issues to domestic issues as well.

Moral foundations theory (MFT) has been subject to a number of criticisms, most pointedly by Suhler and Churchland (2011).¹² First, the five foundations are, if not arbitrary, neither exhaustive nor distinct from one another. As Suhler and Churchland (2011, p. 2106) write, “other basic moral values exhibited by humans across various cultures have as much – or as little – call to be included

¹⁰See also Lakoff (2002; 2009), Jost, Nosek and Gosling (2008), Carney (2008), and Weber and Federico (2013)

¹¹We note that the correlation among the foundations at the *individual* level is somewhat smaller than one might infer from Figure 2. The correlation between care/harm and fairness/cheating at the individual level, for example, is 0.56.

¹²For a response see Haidt and Joseph (2011).

as do Haidt's favored five," with two strong contenders being industry and modesty. Conversely, the five foundations may all be variants of a more fundamental care/harm principle. As Suhler and Churchland (2011, p. 2107) again note, the loyalty/betrayal and sanctity/degradation foundations "may, for all we can be sure, merely be extensions of harm concerns to entities other than individual persons..." There is much to this criticism. Theoretically, one can plausibly link each of the other foundations to the care/harm principle. Empirically, as indicated in Figure 1, there is some overlap between questions on the MFQ and related foundations, especially in the two main clusters. Yet, though there may be strong reason to doubt that the foundations are autonomous modules that can be easily separated one from the others, empirically they do appear to be at least somewhat distinct (Graham et al., 2011).

Recognizing the force of this critique, our justification for using MFT and the MFQ rests on practicality. Of the possible ways of categorizing and measuring moral foundations, MFT is the most developed at the moment, and the MFQ is both publicly available and widely used by others, including in the study of domestic policy preferences. Equally, though the foundations may overlap, they do correlate differently with various domestic and, as we shall see, foreign policy issues in ways that are broadly intuitive. Is MFT in its current incarnation *the* theory of moral foundations? We doubt it. But it does provide a solid basis for examining the possible determinants of foreign policy preferences.

Second, Haidt (2012, Chapter 7) claims that the foundations are innate, emerging as adaptive responses to environmental challenges faced by our primate ancestors and especially early humans living in social groups. Suhler and Churchland (2011) argue, however, that there is no neurobiological or physiological basis in the brain for these traits and, thus, they are unlikely to be innate or genetic in origin. We find it plausible that humans are predisposed to moral foundations, but

remain steadfastly agnostic whether morality is a product of “nature” or “nurture.” Even if the human brain evolved for social skills, a position that Churchland (2011) herself broadly accepts, whether the moral foundations we use are the product of our genes or culture remains very much an open question. We find the extreme form of innateness to be unjustified by existing evidence. While a growing body of work documents a genetic basis for political attitudes, including those specifically related to foreign policy (Benjamin et al., 2012; Cranmer and Dawes, 2012; Hatemi and McDermott, 2012), there is not yet similar systematic empirical work linking genetic traits to the moral foundations.

In the end, we do not need to resolve the innateness question for our purposes. All that is necessary for us here is that moral foundations are formed prior to political awareness, or the development of an interest in and knowledge of politics within one’s community. Although there is no definite dividing line, it appears that political awareness typically begins to arise only in late childhood or early adulthood.¹³ This seems plausible. Yet, even young children have strong moral senses about caring/harm, fairness/cheating, and so on, long before they can generalize these standards to the political community in which they live and, more so, to other political communities with which they might indirectly interact as captured in our foreign policy questions. This is not to suggest that children who are more caring or concerned with fairness at a young age will carry those concerns into adulthood. Although one’s moral foundations seem to be relatively stable, this has not been tested in panel studies and it seems likely that one’s moral stance evolves over time. Nonetheless, it seems reasonable that an individual’s moral foundations are “pre-political” or broadly formed prior to engaging specific questions of politics and public policy. Although

¹³This is a neglected topic. See Hess and Torney (2009), originally published in 1967, and Furnham and Gunter (1983).

it is possible, of course, that policy preferences may affect one's moral foundations, we think the feedback is likely minimal. We proceed in the remainder of this paper to assume that moral foundations are essentially exogenous to foreign policy preferences.

2 Moral Foundations of Foreign Policy Attitudes

2.1 Survey

To examine the role played by moral foundations in foreign policy attitudes, we conducted a survey of 1,561 American adults recruited on the Amazon Mechanical Turk (MT) platform. MT provides access to a large pool of users who complete short tasks for small payments. While MT does not provide access to a random or nationally representative sample, the available pool is considerably more diverse than other convenience samples, such as college students (Paolacci, Chandler and Ipeirotis, 2010). The non-random nature of the sample does threaten the generalizability of our results, and our initial sample was more female, more liberal, and less religious than the American public (see Appendix, Table A.1). Nonetheless, we have no reason to believe that the sampling procedure threatens the internal validity of our conclusions for two reasons. First, we recruited subjects to answer a survey about their "attitudes and opinions," without any mention of either politics or morality, so selection based on those factors is unlikely. Second, we see no reason to suspect that MT users systematically differ from the general public in terms of the relationship between the moral foundations and their foreign policy attitudes, and we find that the correlations among dependent variables and demographics in our survey are similar to those found elsewhere (see Appendix, Figure A.1). This parallels the finding of Berinsky, Huber and Lenz (2012), who

replicate a number of established experimental effects using samples drawn from MT.

A second concern in using MT involves data quality. Because MT users are paid upon completion of tasks, they have incentives to complete those tasks as quickly as possible, which might lead to low quality data. Buhrmester, Kwang and Gosling (2011) find that the reliability of psychometric instruments given via MT is as good as or better than the reliability of instruments administered in person, irrespective of the compensation level. Nonetheless, we took steps to ensure the quality of our data by embedding two “catch” questions with obvious answers into our instrument and discarding questionnaires submitted by respondents who failed these questions. This led to the elimination of 368 responses (roughly 24% of our total).¹⁴

Our survey instrument contained four sets of questions: a group of demographic controls, the MFQ30 used to measure the moral foundations, a set of foreign policy questions drawn from previous surveys, and a set of questions about domestic politics used to replicate earlier work. The demographic questions were always presented first, but the order of the other questions was randomized across respondents. We discuss each of these sets of questions in turn.

For our demographic controls, we collected the same variables used by Haidt and his collaborators: age, gender, political ideology, interest in politics, education, party identification, household income, and religiosity. Because we include a number of questions about foreign economic policy, we also attempted to capture whether or not the primary wage earner in a respondent’s household works in a tradable sector of the economy. Our measure of occupation is crude - we collect data on employment in ten broad groupings, classifying agriculture, mining, and manufacturing as tradable sectors - but we have no reason to believe that employment in particular industrial sectors

¹⁴The first of these asks respondents to agree or disagree with the statement: “It is better to do good than to do bad”; respondents who disagreed were discarded. The second question is embedded in the MFQ and asks if “Whether or not someone was good at math” is relevant to their moral judgments. Respondents who found this relevant were discarded.

correlates with individual attitudes about morality, so it is unlikely that the lack of fine-grained employment data biases our results.

Our next group of questions is the standard MFQ, which consists of six questions designed to measure each of the five moral foundations, for a total of 30 questions. These questions break into two groups - the judgments subscale (which asks respondents to express agreement or disagreement with normative statements) and the relevance subscale (which asks respondents how important various factors are to their opinions about morality). These questions were originally generated for face validity and then refined on the basis of internal consistency and breadth of coverage. The developers of the MFQ have demonstrated that it has high internal consistency, high test-retest validity, and that five factors fit the data it generates well in confirmatory factor analysis. MFQ responses also have high predictive power for a variety of related attitudes and judgments (Graham et al., 2011). Scores for each foundation are calculated from the six relevant questions as a simple additive index and normalized to range from 0 to 5. Because there are six possible responses to each question, the resulting measure is effectively continuous.

The foreign policy questions on our instrument were drawn from three recent surveys conducted on nationally representative samples: the 2010 Chicago Council on Global Affairs survey, the 2011 Transatlantic Trends survey conducted by the German Marshall Fund of the United States, and the Spring 2010 Pew Global Attitudes Project survey. We selected questions from these surveys in order to maximize coverage of different topic areas in international relations and chose questions that featured substantial differences of opinion among US respondents. The selection of questions from existing instruments allows for greater comparability of our results to other research. Finally, our instrument included a series of questions about domestic politics to replicate an earlier set of findings about the relationship between the moral foundations and political attitudes

(Koleva et. al. 2012). The full survey included 93 questions and was designed to take between fifteen and twenty minutes to complete.

2.2 Survey Results

Results for all foreign policy questions are detailed in the appendix, and summarized in Table 3. In Tables 1 and 2, we present full results for two sample questions for purposes of illustration. Because the responses to each of our foreign policy questions come in ordered categories, we report the results from an ordered probit specification in each case. For each question, we examine three models: one with only each respondent's scores on the five moral foundations, a second controlling for that respondent's self-described political ideology, and a third controlling for the full set of variables listed above.

[Table 1 About Here]

We selected one highly “conservative” question for illustration, which asks respondents whether they believe that “maintaining superior military power worldwide” should be a “very important,” “somewhat important,” or “not important” goal of US foreign policy. Table 1 shows that care/harm is strongly and significantly associated with opposition to maintaining US power; this finding changes little with the inclusion of controls. Conversely, all three of the “binding” foundations – loyalty/betrayal, authority/submission, and sanctity/degradation – are significantly associated with support for maintaining US power. Even after ideology is included as a control, these foundations carry the same sign and remain significant (model 2), showing clearly that they capture attitudes above and beyond ideology. The same is true with our full battery of controls (model 3).

[Table 2 About Here]

We also selected one highly “liberal” question, which asks respondents to agree or disagree with the statement “When dealing with international problems, the US should be more willing to make decisions within the United Nations even if this means that the United States will sometimes have to go along with a policy that is not its first choice.” For this question, Table 2 shows that the care/harm, fairness/cheating, and sanctity/degradation foundations are all significant. Care/harm and fairness/cheating increase support for using the UN, while sanctity/degradation reduces it. Although there is some change in the point estimates (particularly for sanctity/degradation) after we add control variables, the overall structure of the results on the foundations are again largely unchanged by the inclusion of the controls.

Table 3 summarizes the core findings from all foreign policy questions on our survey. Within the table, the questions are split into three groups, based on the available answer categories, and then ranked on the basis of the unconditional correlation between ideology and the responses. The five columns of the table display the five moral foundations. Because ordered probit coefficients are difficult to interpret directly, each cell displays the change in the predicted probability of a specified answer for a one point change (on the 0 to 5 scale) in the specified foundation. The specifications displayed in the table include measures of all of the foundations and all demographic control variables (model 3 in Tables 1 and 2). The significance levels displayed in the table (green means positive and significant; red means negative and significant) refer to the underlying ordered probit coefficients rather than the predicted effects. The full ordered probit results and alternative specifications for each model are displayed in the Appendix (Table A.3).

[Table 3 About Here]

The correlation between the moral foundations and opinions on a wide variety of foreign policy issues is quite striking. For every question but one, responses correlate with at least one (and generally more than one) of the foundations at 95% confidence or greater. Within each category of question, listed in order of their unconditional correlation with political ideology, positive and significant results are largely found along the main diagonal, as expected, and negative and significant results are found along the secondary diagonal. The results display a theoretically meaningful pattern - the “liberal” foundations (care/harm and fairness/cheating) tend to be positively and significantly associated with the more “liberal” questions and negatively associated with the more “conservative” questions; while the “conservative” foundations (loyalty/betrayal, authority/submission, and sanctity/degradation) tend to be positively and significantly associated with the “conservative” questions. It is important to note that this association holds even *after* controlling for respondent ideology. This forms a striking visual display that the moral foundations explain a substantial portion of the variance in foreign policy attitudes over and above that which can be attributed to political ideology. It is also important to observe, however, that foundations from the same cluster sometimes point in opposite directions (as with authority/submission and sanctity/degradation on our question about growing international business ties), while on others, foundations from the different clusters point in the same direction (as with fairness/cheating and authority/submission on our question about strengthening the United Nations.) This result underscores the fact that the foundations are distinct from a unidimensional conception of ideology.

Despite the convergence of foreign policy and foreign economic policy studies noted in the Introduction, questions on globalization, trade and immigration appear to be somewhat distinct. Questions relating to foreign economic policy are, interestingly, typically less ideological and cluster in the middle of the agree/disagree section of Table 3. Nonetheless, the moral foundations are

still strongly related to foreign economic policy attitudes, though the economic items rarely respond to the loyalty/betrayal and fairness/cheating foundations, despite the prominence of related constructs in other explanations. We also see that purity concerns are frequently important to the economic policy questions, correlating with opposition to free trade and immigration, likely reflecting a fear of the “foreign”. Although our sector of employment variable is never significant (see Appendix), we do find a number of relationships with the other controls that are broadly consistent with existing findings. Yet, the moral foundations appear to highlight a rather different set of concerns about globalization rooted in authority/submission and sanctity/degradation.

The significant correlations between moral foundations and foreign policy attitudes suggest that foreign policy preferences are rooted, at least in part, in moral positions and likely to originate from the same sources as attitudes toward domestic policy issues. Nonetheless, moral foundations only account for a portion of the overall variance in foreign policy attitudes. To gauge the relative importance of the foundations, we estimate the predictive accuracy of three models: a null (intercept-only) model, a model with only our control variables, and a full model with the controls and foundations. To avoid the problem of over-fitting, we estimate the true error rate of these models through “leave-one-out” cross-validation (i.e., for each observation, we fit a model to all other observations then use this model to predict the “left out” observation). The addition of the foundations improves predictive accuracy for 23 of our 26 questions. Aggregating across all of the questions, the intercept-only model correctly predicts 38.7% of responses; the model with only our controls correctly predicts 40.6%, and the complete model with both controls and foundations predicts 42.3% correctly. Obviously, there is still considerable residual variation for all of these models, but we note that the improvement in predictive performance from including the foundations is as large as the improvement provided by the controls over the null model, despite

the fact that there are twice as many control variables, showing that the foundations are a relatively important correlate of foreign policy attitudes.

2.3 Clustering Attitudes

To further explore the moral foundations and the relationship of our work here to other studies of foreign policy attitudes, we fit a second statistical model in which we estimate a spatial model of our 26 foreign policy questions. For simplicity of estimation and exposition, in these models we dichotomize all of the responses, and look at agreement with a statement (or favoring a use of force/finding a foreign policy goal important). We also adopt a simpler two-dimensional model, where the dimensions are given by the individualizing and binding foundations, although the results of the five dimensional model are broadly similar.

For our estimation, we adapt the Bayesian IDEAL procedure, commonly used for legislative voting data (Jackman, 2001; Clinton, Jackman and Rivers, 2004). Specifically, we assume that each respondent i has an ideal point x_i in the two-dimensional moral space, while for each question, j , agreement and disagreement can also be positioned in this space at ζ_j and ψ_j respectively, where respondents utility for each response decreases in squared distance between these points and their own ideal points as well as a disturbance term. This yields the reduced form:

$$y_{ij}^* = U_i(\zeta_j) - U_i(\psi_j) = \beta_j' x_i - \alpha_j + \epsilon_{ij}$$

$$y_{ij} = 1 \Leftrightarrow y_{ij}^* > 0$$

where $\beta_j = 2 * (\zeta_j - \phi_j)$ and $\alpha_j = \zeta_j' \zeta_j - \psi_j' \psi_j$. We adopt the common assumption that $\epsilon \sim N(0, 1)$. Unlike in the typical ideal point setting, we are interested primarily in the item

parameters (α and β), rather than the respondent ideal points. Assuming these respondent ideal points to be known with *certainty* from the MFQ would allow us to estimate a probit and recover these parameters; however, we adopt a more flexible model, which accounts for the measurement error inherent in the MFQ. Rather than incorporate the MFQ scores as known data, we model respondent ideal points as parameters, using the MFQ scores as informative priors on the two latent dimensions. Specifically, our prior is that the respondent ideal points are normally distributed around the MFQ scores with a standard deviation of 0.33. We impose uninformative uniform priors on the α and β parameters and estimate the model in Stan 2.1.0 (Stan Development Team, 2013). We use four chains of length 5,000 with 2,000 discarded warmup draws, and assess convergence using the Gelman-Rubin diagnostic.

We transform the resulting estimates to calculate the “cutting line” separating agree and disagree responses, as shown in Figure 3, for our question about whether or not the United States should make decisions within the United Nations. To allow inter-question comparisons, we then express the β parameters for each question as a position vector for coordinates (β_1, β_2) where β_1 corresponds to the individualizing foundations and β_2 corresponds to the binding foundations. The direction of this vector is the same as the angle between the cutting line and the x-axis; consequently, it expresses the relative extent to which the two dimensions discriminate on a given question. Because dimensions are scaled equivalently, values close to zero degrees indicate that the agree and disagree responses are separated almost exclusively along the binding dimensions, while values close to positive or negative 90 degrees indicate agree and disagree responses are separated almost exclusively along the individualizing dimension, and values close to positive or negative 45 degrees indicate equal discrimination on the two dimensions. The magnitude of this vector indicates the signal-to-noise ratio, showing how strongly responses change as a function of

the foundations relative to the size of the error term.

[Figure 3 About Here]

We plot the end points of these vectors in Figure 4 on a polar coordinate system. The angular coordinates in the figure represent twice the angle of the vector; these are labelled around the circumference of the circle. Thus, questions located on the same radial line in Figure 4 have parallel cutting lines in the two-dimensional moral space; while questions located on opposite radial lines in the figure have orthogonal cutting lines in the moral space. Distances from the origin in the figure represent the magnitude of the vectors, so that a larger distance from the origin indicates a larger effect for the moral foundations on individual responses. Consequently, the distance between a pair of points in Figure 4 represents how similarly responses to the two questions change as a result of changes in the moral foundations; similar angular position indicates that the change is in the same direction, while similar distances from the origin indicate that change is of the same magnitude.

[Figure 4 About Here]

Figure 4 reproduces some common sense clusters. We find, for example, that questions 17 and 18 (our two questions about immigration) are practically on top of one another as are questions 7, 12, and 13, all of which concern foreign aid. To facilitate further interpretation, we use k-means clustering to identify three clusters in the figure based *solely* on the angular positions and color these clusters; the clusters detected correspond nearly exactly to the three most populated quadrants of the figure. We note that the plot reflects only the properties of the cutting line, not of “agreement”; if two questions were perfectly negatively correlated, they would be placed at the

same point in Figure 4. We do this so that “flipped” questions will remain proximate (e.g., had we asked respondents whether maintaining superior military power should *not* be an important goal of US foreign policy, this question would fall in exactly the same position as question 4 in the figure).

The first cluster, shown in green, consists of questions that discriminate more on the binding foundations than the individualizing foundations (reaching equality for question 9), and for which the two dimensions move respondents in opposite directions (i.e., increases in the binding foundations make respondents more likely to agree, while increases in the individualizing foundations make respondents less likely to agree or vice versa). In this cluster, we find a combination of issues related to economic and cultural protectionism (questions 16-19) and military interventionism, such as whether or not to prioritize maintaining superior US military power (4) and whether or not the US should undertake military intervention to ensure the oil supply (22). These questions appear loosely related to Holsti’s “hard-liner” types, in that they involve self-interested uses of force; however, rather than being similar to a set of non-cooperative attitudes on international order, we find that these militaristic attitudes are related to a set of non-cooperative attitudes on international economics.

In the red cluster, like the green cluster, the foundations move in opposite directions; however, unlike the red cluster, questions in the green cluster discriminate primarily on the individualizing dimension. In this cluster, we find issues related to foreign aid (7,12, and 13), the United Nations (5 and 8), globalization (15), and humanitarian intervention (21 and 23). This cluster appears to map fairly neatly onto Holsti’s liberal internationalist type as it draws a relationship between militarist internationalism (e.g., intervention to stop a genocide) and cooperative internationalism (e.g., support for strengthening the UN).

In the black cluster, unlike the green and red, the individualizing and binding foundations

move respondents in the same direction (i.e, either increases in both the binding and individualizing foundations make a respondent more likely to agree with the statement or increases in both make a respondent less likely to agree). Interestingly (though this may be a quirk of our questions), all of the questions in this cluster fall near the -45 degree line, indicating that the individualizing and binding foundations discriminate about equally for the questions in the cluster. The questions involved appear to reflect what might be termed “neoconservative” values, including attitudes towards democracy promotion (2), strong US leadership in world affairs (20), and the hypothetical use of force to prevent an Islamist takeover of Pakistan (24).

While the general pattern of our clusters is similar to that found by others through inductive methods, our two dimensions have clear theoretical content (as opposed to the atheoretical dimensions derived from such techniques as factor analysis). Our clusters differ from many others, however, in that we do not find any unique moral position corresponding to isolationism; instead, isolationism and unilateralism resemble one another, while differing from multilateralism. Nonetheless, the moral foundations appear to explain a substantial proportion of the coherence of foreign policy attitudes, perhaps providing microfoundations for other approaches. To assess the relative contribution of the moral foundations to the latent structure of foreign policy attitudes, we refit the model described above and add a third dimension with mean zero and standard deviation one, allowing the model to find the best additional dimension inductively.¹⁵ We then compare the performance of this model to our model, as well as a null model. In the dichotomous data, the null model correctly predicts 66.2% of all responses, while the latent moral model predicts 72.2% correctly, and the addition of a third dimension improves prediction to 77.1% correct. Thus, the moral

¹⁵While the first two dimensions are identified by the informative priors, this third dimension is not globally identified by our assumptions. It is locally identified (Rivers, 2003; Clinton, Jackman and Rivers, 2004), so as we have no basis for constraining parameters, we simply proceed with the locally identified model, selecting a mode and taking care to ensure that our sampler does not cross over.

model seems to have captured a substantial proportion of the underlying coherence of attitudes as compared to a higher dimensional alternative.

3 Priming Experiments

Thus far, we have examined general effects of the moral foundations on relatively generic attitudes, but the effect of political rhetoric and public media on the link between moral foundations and foreign policy attitudes is potentially important, particularly for active policy disputes. The frames used by politicians and the media often appeal to morality. If this moral rhetoric is able to substantially shift opinions, then, while the moral foundations might still supply a partial explanation of variation in viewpoints, their effects are likely to be highly variable and contingent on specific elite-level rhetorical strategies. On the other hand, Moral Foundations Theory suggests that the foundations act primarily through an immediate, intuitive effect when respondents are presented with an issue, so their effects should remain essentially constant even when respondents are primed for moral concerns, either by political rhetoric or an experimental treatment. Here, we find strong support for the latter expectation.

3.1 Priming Surveys

To test the effects of priming and rhetoric, we conducted two survey experiments in January and March of 2013, in which we primed respondents with one of the moral foundations, then posed the same foreign policy questions used in our initial survey. In the first experiment, we primed respondents for the care/harm foundation; in the second, we primed the loyalty/betrayal foundation. We again recruited American adults on MT yielding 535 respondents for the first experiment and 583

for the second. These respondents answered a reduced set of demographic questions (ideology, education, and gender), the same foreign policy questions as respondents in the first round, and two catch questions, as well as completing the priming protocol.

In designing any priming experiment, a core decision is how “heavy-handed” to make the treatment. We chose a relatively subtle treatment, described in detail below. We use a subtle prime because we believe it more accurately reflects the theoretical construct of political rhetoric of national leaders in a mature democracy. If one politician evokes sanctity to appeal to one possible constituency, his opponent is likely to reframe the issue as one of harm or fairness to appeal to an alternative constituency. Average citizens then get bombarded by conflicting frames, the net effect of which is unclear. We anticipate that a more subtle prime will more realistically capture the net effect of competing rhetoric. The use of a more “subtle” prime also helps us overcome inferential threats that plague heavier priming. Several recent papers have failed to replicate prominent priming studies and attributed existing results to various forms of experimenter bias or hypothesis guessing (Doyen et al., 2012; Skelton, Loveland and Yeagley, 1996; Zemack-Rugar, Bettman and Fitzsimons, 2007). Using a “strong” prime introduces a much higher risk that respondents will guess (correctly or incorrectly) the purpose of experiment and adjust their behavior accordingly, while our more subtle primes mitigate this threat.

For our primes, we used the scrambled sentence protocol introduced by Srull and Wyer (1979), which has been adopted by numerous other studies in diverse contexts (Bargh, Chen and Burrows, 1996; Epley and Gilovich, 1999; Shariff and Norenzayan, 2007). We present respondents with 12 sets of 5 words, each of which can be unscrambled into a grammatical, four-word sentence when leaving out one word. We ask respondents to unscramble these sentences, which are identical between the treatment and control conditions with the exception of a single “target” word, which

changes in 9 of the 12 sentences. In the control condition, these target words are neutral with respect to morality (tall, encountered, punctuality, music, goal, quickly, cat, and expensive). In the treatment condition, the “target” word changes to a foundation-related word taken from the Moral Foundations Dictionary,¹⁶ a set of words related to each of the moral foundations developed by Graham, et. al. (2009). We selected four virtue primes for the foundation being primed (caring, compassion, safety, and shelter for harm; loyal, unity, community, and family for loyalty) as well as four vice primes (endangered, cruelty, violently, and ruined for harm; excluded, spy, treacherously, and deserted for loyalty). The sentences are presented in pairs (a virtue sentence followed by a vice sentence) separated by buffer sentences (neutral sentences identical in the two conditions). As an example, a care/harm prime asks respondents to unscramble the words “Locate Woman Caring Was The” into the sentence “The woman was caring.” The neutral or control prime for this sentence unscrambles as “The woman was tall.” In the loyalty/betrayal condition, this becomes “The woman was loyal.”

The twelfth and final word group in the protocol unscrambles in two valid ways, only one of which uses the target word.¹⁷ By comparing the proportion of treated respondents to the proportion of control respondents who used the “target” word in their sentences, we are able to determine whether or not the priming succeeded in mentally activating the moral constructs we sought to evoke (Higgins and King, 1981). We find that a substantially (and significantly) higher proportion of respondents used the target word in the treatment conditions, indicating that the priming was

¹⁶Available online at <http://www.moralfoundations.org/downloads/moral%20foundations%20dictionary.dic> (Accessed March 1, 2013)

¹⁷This sentence could be unscrambled as “The teacher was qualified” in both treatment and control conditions. In the harm condition, respondents could also form the sentence “The teacher was sympathetic”; in the loyalty/betrayal condition, respondents could also form the sentence “The teacher was devoted.”; while in the control condition, respondents could also form the sentence “The teacher was young” (the words “sympathetic”, “devoted”, and “young” appeared in the same position in the jumbles).

successful (see appendix).

3.2 Results of Priming Experiments

Overall, our experiments did not identify any significant or systematic effects of the primes on foreign policy attitudes. We present results for the same two illustrative questions as in section two. For each question, the first column contains all respondents, the second reports only liberal respondents, and the third reports only conservative respondents. Table 4 displays the results for the care/harm priming, while Table 5 displays the results for the loyalty/betrayal priming.

[Table 4 About Here]

Examining Table 4 shows the muted effects of care/harm priming, with the prime (the variable “treated”, which takes on a value of one when respondents received the treatment prime and zero otherwise) taking on only marginal ($p < .1$) significance in two of the models of the UN question. Ideology takes on the same sign and roughly the same magnitude for both questions as it did in the original survey (Tables 1 and 2), which suggests that our samples are comparable. Priming for the care/harm foundation, however, appears to have made respondents, if anything, less supportive of the UN and more supportive of US military power. This is the opposite of the expectation generated by our first wave, in which higher care/harm scores were correlated with more support for the UN and less support for US military power. Breaking apart the sample into liberals and conservatives shows some potential heterogeneity (particularly on the UN question), but we still see no meaningful treatment effects.

[Table 5 About Here]

As in Table 4, Table 5 replicates our original finding on ideology for both questions. We see, again, no meaningful priming effects at the conventional level. Within the sample as a whole, we find marginally significant ($p < .1$) effects for the loyalty/betrayal priming, and when limiting our sample to conservatives, the coefficient on assignment to treatment becomes significant at the conventional (i.e., 95%) level; however, the priming made our respondents *less* supportive of US power, despite the positive correlation between loyalty/betrayal scores and support for US power in the first wave. As discussed below, we attach little meaning to this finding, as it was one of the few significant coefficients out of a large number tested.

?? presents the results of the care/harm priming for all foreign policy questions, while Table 7 presents the results for the loyalty/betrayal priming. In both tables, the results displayed represent the change in the predicted probability of a given answer category for treated respondents (as compared to control respondents) after controlling for ideology, gender, and education. As in Table 4, the first column presents results for the full sample, the second column presents results only for liberal respondents, and the third column presents results for only conservative respondents. The final column presents a “treatment on the treated” estimate by coding only respondents in the treatment condition who included the target word in the test priming sentence as treated.

[?? About Here]

?? reveals that the care/harm priming experiment did not generate strong effects. Assignment to treatment has a statistically significant effect at 95% confidence in only four models (out of 104 estimated), which is almost exactly the expected rate of false positives; consequently, we do not attribute any meaning to the particular coefficients found to be significant. In addition to the lack of statistical significance, no clear substantive patterns emerge in the data. If priming moved

attitudes meaningfully, we would expect to see the sign on our treatment variable match the sign on the care/harm foundation from the original survey (Table 3); that is, if higher care/harm scores are associated with increased support for a given US policy goal, we would expect to see the care/harm priming also increase support for that goal. We find, however, that the sign on the treatment dummy matches the sign from the original survey only 14 out of 26 times. If we limit attention to questions where care/harm was found to be significant in the initial survey, the sign matches only 6 out of 16 times. In short, we find no effect.

[Table 7 About Here]

As with the care/harm priming, we find null results on the loyalty/betrayal priming displayed in Table 7. Once again, the coefficient on our dummy for treatment is significant in four out of 104 models, which suggests that the associations we find are likely the product of chance. Similarly, the sign for treatment matches the sign on the loyalty/betrayal foundation from Table 3 only 16 out of 26 times. If we restrict attention to cases where the coefficients were significant in the first round, the sign matches 5 out of 11 times. Once again, this leads us to conclude that the effects detected are simple noise.

While it is possible that the prime was too weak to cause significant changes in answers to the foreign policy questions, the test described above does show that the priming successfully activated the relevant concepts.¹⁸ Consequently, we interpret the results as an indication that moral concerns are already salient in the minds of respondents when they answer questions about foreign policy. Because individuals already view foreign policy issues through a moral lens, priming has little or no effect. Thus, we suspect, but do not show here, that political rhetoric invoking moral

¹⁸Similarly, it is possible that our sample size was insufficient to pick up the effects. We address this possibility in the appendix.

concerns is unlikely to substantially shift opinions on foreign policy issues.

4 Conclusion

The results of our primary survey show that moral foundations are statistically significant and often substantively meaningful predictors of a diverse body of foreign policy attitudes. Although the results linking particular moral foundations to specific policy issues are largely intuitive, the findings as a whole clearly indicate that foreign policy attitudes are associated with moral worldviews above and beyond what is captured in political ideology. Rather than any single correlation, this link between moral foundations and foreign policy attitudes overall is our most important conclusion. This extends the findings of Koleva et al. (2012) that the moral foundations significantly predict variation in opinion for a diverse group of domestic policy issues, while challenging the claims of Hurwitz and Peffley (1987) that domestic and foreign policy are best explained separately. Rather, moral foundations appear to underlie a range of public policy preferences.

Our findings also demonstrate considerable similarity between the domains of security policy and foreign economic policy. The same moral foundations prove to be statistically significant predictors of both sets of attitudes. The same foundations are also significant predictors of both relatively generic attitudes (e.g., whether the US should protect weaker nations against aggression) and more specific policy questions (e.g., whether the US should use force in response to a Chinese invasion of Taiwan). Furthermore, the effect of the foundations demonstrates a certain nuance; we see, for example, that the care/harm foundation is positively correlated with support for certain uses of force and negatively correlated with others, which suggests it cannot be reduced to a simple militarism dimension.

Our experiments support Haidt's contention that the moral foundations are intuitive and automatically activated when considering political and other issues. Priming has no additional effect on attitudes, implying that our survey respondents are already considering moral implications when filling out the questionnaire. Of course, intuition may be overridden by further assessment, unlikely in the context of our online survey. Actual foreign policy decision makers or even voters at the end of a long campaign in which foreign policy issues are debated at length may reach more considered views that are more or less "moral" than our results suggest, but, for many individuals who do not evaluate foreign policy issues in depth, intuition may be all that matters.

Although the moral foundations prove to be a novel and significant predictor of foreign policy attitudes, they are by no means the only determinant of these attitudes. We generally find that several our demographic control variables also have significant predictive power. Moral foundations are not the sole and likely not even the most important determinant of foreign policy attitudes, but they are important sources of opinion on foreign policy issues. This implies not only, as in past studies, that foreign policy attitudes are well structured, but that these attitudes are rooted in general moral worldviews that analysts would do well to take seriously.

References

- Aldrich, John, Christopher Gelpi, Peter Feaver, Jason Reifler and Kristin Thompson Sharp. 2006. "Foreign Policy and the Electoral Connection." *Annual Review of Political Science* 9:477–502.
- Almond, Gabriel A. 1950. *The American People and Foreign Policy*. New York: Harcourt Brace.
- Baker, Andy. 2005. "Who Wants to Globalize? Consumer Tastes and Labor Markets in a Theory of Trade Policy Beliefs." *American Journal of Political Science* 49(4):924–938.
- Baker, William D. and John R. Oneal. 2001. "Patriotism or Opinion Leadership? The Nature and Origins of the 'Rally Round The Flag' Effect." *Journal of Conflict Resolution* 45(5):661–687.

- Bargh, John A., Mark Chen and Lara Burrows. 1996. "Automaticity of Social Behavior: Direct Effects of Trait Construct and Stereotype Activation on Action." *Journal of Personality and Social Psychology* 71(2):230–244.
- Bartels, Larry M. 1994. "The American Public's Defense Spending Preferences in the Post-Cold War Era." *The Public Opinion Quarterly* 58(4):479–508.
- Baum, Matthew A. 2003. *Soft News Goes to War: Public Opinion and American Foreign Policy in the New Media Age*. Princeton, NJ: Princeton University Press.
- Baum, Matthew A. and Tim J. Groeling. 2008. "The Relationships Between Mass Media, Public Opinion, and Foreign Policy: Towards a Theoretical Synthesis." *Annual Review of Political Science* 11:39–65.
- Baum, Matthew A. and Tim J. Groeling. 2009. *War Stories: The Causes and Consequences of Public Views of War*. Princeton, NJ: Princeton University Press.
- Benjamin, Daniel J., David Cesarini, Matthijs J. H. M. van der Loos, Christopher T. Dawes, Philipp D. Koellinger, Patrik K. E. Magnusson, Christopher F. Chabris, Dalton Conley, David Laibson, Magnus Johannesson and Peter M. Visscher. 2012. "The genetic architecture of economic and political preferences." *Proceedings of the National Academy of Sciences* 109(21):8026–8031.
URL: <http://www.pnas.org/content/109/21/8026.abstract>
- Berinsky, Adam J., Gregory A. Huber and Gabriel S. Lenz. 2012. "Evaluating Online Labor Markets for Experimental Research: Amazon.com's Mechanical Turk." *Political Analysis* 20(3).
- Berinsky, Adam J. and James N. Druckman. 2007. "Public Opinion Research and Support for the Iraq War." *Public Opinion Quarterly* 71(1):126–141.
- Binning, Kevin R. 2007. "'Its Us Against the World': How Distrust in Americans versus People-In-General Shapes Competitive Foreign Policy Preferences." *Political Psychology* 28(6):777–799.
- Brewer, Paul R. and Marco R. Steenbergen. 2002. "All Against All: How Beliefs about Human Nature Shape Foreign Policy Opinions." *Political Psychology* 23(1):39–58.
- Buhrmester, Michael, Tracy Kwang and Samuel D. Gosling. 2011. "Amazon's Mechanical Turk: A New Source of Inexpensive, Yet High-Quality, Data?" *Perspectives on Psychological Science* 6(1):3–5.
- Campbell, Angus, Philip E. Converse, Warren E. Miller and Donald E. Stokes. 1960. *The American Voter*. New York: Wiley.
- Carney, Dana R. 2008. "The Secret Lives of Liberals and Conservatives: Personality Profiles, Interaction Styles, and the Things They Leave Behind." *Political Psychology* 29(6):807–840.
- Chapman, Terrence L. and Dan Reiter. 2004. "The United Nations Security Council and the Rally 'Round the Flag Effect." *Journal of Conflict Resolution* 48(6):886–909.

- Clinton, Joshua, Simon Jackman and Douglas Rivers. 2004. "The Statistical Analysis of Roll Call Data." *American Political Science Review* 98(2):355–370.
- Converse, Philip E. 1964. The Nature of Belief Systems in Mass Publics. In *Ideology and Its Discontents*, ed. D.E. Apter. New York: Free Press pp. 206–211.
- Cranmer, Skyler and Christopher Dawes. 2012. "The Heritability of Foreign Policy Preferences." *Twin Research and Human Genetics* 15(1):52–59.
- Doyen, Stéphane, Olivier Klein, Cora-Lise Pichon and Axel Cleeremans. 2012. "Behavioral Priming: It's All in the Mind, but Whose Mind?" *PLoS ONE* 7(1):e29081.
- Epley, Nicholas and Thomas Gilovich. 1999. "Just Going Along: Nonconscious Priming and Conformity to Social Pressure." *Journal of Experimental Social Psychology* 35(6):578 – 589.
- Fordham, Benjamin O. 2008. "Economic Interests and Public Support for American Global Activism." *International Organization* 62(1):163–182.
- Fordham, Benjamin O. and Katja B. Kleinberg. 2012. "How Can Economic Interests Influence Support for Free Trade?" *International Organization* 66(2):311–328.
- Frieden, Jeffrey A. and Ronald Rogowski. 1996. The Impact of the International Economy on National Policies: An Analytical Overview. In *Internationalization and Domestic Politics*, ed. Robert O. Keohne and Helen V. Milner. Cambridge University Press pp. 25–47.
- Furnham, Adrain and Barrie Gunter. 1983. "Political Knowledge and Awareness in Adolescents." *Journal of Adolescence* 6(4):373–385.
- Gartner, Scott S. 1997. "All Politics Are Local: Local Loses and Individual Attitudes toward the Vietnam War." *Journal of Conflict Resolution* 41(5):699–694.
- Gelpi, Christopher, Peter D. Feaver and Jason Reifler. 2006. "Success Matters: Casualty Sensitivity and the War in Iraq." *International Security* 30(3):pp. 7–46.
- Gerber, Alan S., Gregory Huber, David Doherty and Conor Dowling. 2011. "The Big Five Personality Traits in the Political Arena." *Annual Review of Political Science* 14:265–287.
- Goldstein, Judith, Yotam Margalit and Doug Rivers. 2007. "Producer, Consumer, Family Member: The Relationship Between Trade Attitudes and Family Status." Palo Alto, CA, Stanford University.
- Graham, Jesse, Brian Nosek, Jonathan Haidt, Ravi Iyer, Spassena Koleva and Peter Ditto. 2011. "Mapping the Moral Domain." *Journal of Personality and Social Psychology* 101(2):366–385.
- Haidt, Jonathan. 2012. *The Righteous Mind: Why Good People are Divided by Politics and Religion*. New York: Pantheon.
- Haidt, Jonathan and Craig Joseph. 2004. "Intuitive Ethics: How Innately Prepared Intuitions Generate Culturally Variable Virtues." *Daedalus* 133(4):55–66.

- Haidt, Jonathan and Craig Joseph. 2007. "When Morality Opposes Justice: Conservatives Have Moral Intuitions that Liberals May Not Recognize." *Social Justice Research* 20:98–116.
- Haidt, Jonathan and Craig Joseph. 2011. "How Moral Foundations Theory Succeeded in Building on Sand: A Response to Suhler and Churchland." *Journal of Cognitive Neuroscience* 23(9):2117–2122.
- Haidt, Jonathan and Jesse Graham. 2007. "When Morality Opposes Justice: Conservatives Have Moral Intuitions that Liberals may not Recognize." *Social Justice Research* 20:98–116.
- Hainmueller, Jens, Michael Hiscox and Yotam Margalit. 2011. "Do Concerns About Labor Market Competition Shape Attitudes Toward Immigration? New Evidence." Annual Meeting of the International Political Economy Society. Charlottesville, VA.
- Hatemi, Peter K. and Rose McDermott. 2012. "The genetics of politics: discovery, challenges, and progress." *Trends in Genetics* 28(10):525 – 533.
- Herrmann, Richard K., Philip E. Tetlock and Matthew N. DiScro. 2001. "How Americans Think About Trade: Reconciling Conflicts Among Money, Power, and Principles." *International Studies Quarterly* 45(2):191–218.
- Herrmann, Richard K., Philip E. Tetlock and Penny S. Visser. 1999. "Mass Public Decisions to Go to War: A Cognitive-Interactionist Framework." *American Political Science Review* 93(3):553–573.
- Hess, Robert D. and Judith V. Torney. 2009. *The Development of Political Attitudes in Children*. Piscataway, NJ: Transaction Books.
- Higgins, E.T. and G. King. 1981. Accessibility of social constructs: Information processing consequences of individual and contextual variability. In *Personality, Cognition, and Social Interaction*. Hillsdale, NJ: Erlbaum pp. 69–121.
- Holsti, Ole R. 1996. *Public Opinion and American Foreign Policy*. Ann Arbor: University of Michigan Press.
- Hurwitz, Jon and Mark Peffley. 1987. "How Are Foreign Policy Attitudes Structured? A Hierarchical Model." *American Political Science Review* 81(4):1099–1120.
- Jackman, Simon. 2001. "Multidimensional Analysis of Roll Call Data via Bayesian Simulation: Identification, Estimation, Inference, and Model Checking." *Political Analysis* 9(3):227–241.
URL: <http://pan.oxfordjournals.org/content/9/3/227.abstract>
- Jentleson, Bruce W. 1992. "The Pretty Prudent Public: Post Post-Vietnam American Opinion on the Use of Military Force." *International Studies Quarterly* 36(1):pp. 49–73.
- Jost, John T., Brian Nosek and Samuel Gosling. 2008. "Ideology: Its Resurgence in Social, Personality, and Political Psychology." *Perspectives on Psychological Science* 3(2):126–136.
- Kahneman, Daniel. 2011. *Thinking Fast and Slow*. New York: Farrar, Straus, and Giroux.

- Kaltenhaler, Karl C., Ronald D. Gellens and Stephen J. Ceccoli. 2004. "Explaining Citizen Support for Trade Liberalization." *International Studies Quarterly* 48:829–851.
- Koleva, Spassena, Jesse Graham, Ravi Iyer and Peter Ditto. 2012. "Tracing the threads: How five moral concerns (especially Purity) help explain culture war attitudes." *Journal of Research in Personality* 46:184–194.
- Lakoff, George. 2002. *Moral Politics: How Liberals and Conservatives Think*. Chicago: University of Chicago Press.
- Lakoff, George. 2009. *The Political Mind: A Cognitive Scientist's Guide to Your Brain and Its Politics*. New York: Penguin.
- Lippmann, Walter. 1955. *Essays in the Public Philosophy*. Boston: Little Brown.
- Mansfield, Edward D. and Diana C. Mutz. 2009. "Support for Free Trade: Self-Interest, Sociotropic Politics, and Out-Group Anxiety." *International Organization* 63(3):425–457.
- Mueller, John E. 1973. *Presidents and Public Opinion*. New York: Wiley.
- Naoi, Megumi and Ikuo Kume. 2011. "Explaining Mass Support for Agricultural Protectionism: Evidence from a Survey Experiment During the Global Recession." *International Organization* 65(04):771–795.
- Nincic, Miroslav. 1992. "A Sensible Public: New Perspectives on Popular Opinions and Foreign Policy." *Journal of Conflict Resolution* 36(4):772–789.
- Nincic, Miroslav and Jennifer Ramos. 2010. "Ideological Structure and Foreign Policy Preferences." *Journal of Political Ideologies* 15(2):119–141.
- Page, Benjamin I. and Robert Y. Shapiro. 1992. *The Rational Public: Fifty Years of Trends in Americans' Foreign Policy Preferences*. Chicago: University of Chicago Press.
- Paolacci, Gabriele, Jesse Chandler and Panagiotis Ipeirotis. 2010. "Running experiments on amazon mechanical turk." *Judgment and Decision Making* 5(5):411–419.
- Powlick, Philip and Andrew Katz. 1998. "Defining the American Public Opinion/Foreign Policy Nexus." *International Studies Review* 42:29–61.
- Rai, T.S. and A.P. Fiske. 2011. "Moral Psychology Is Relationship Regulation: Moral Motives for Unity, Hierarchy, Equality, and Proportionality." *Psychological Science* 118:57–75.
- Rankin, David M. 2001. "Identities, Interests, and Imports." *Political Behavior* 23(4):351–376.
- Rho, Sungmin and Michael Tomz. 2012. "Industry, Self-Interest, and Individual Preferences over Trade Policy." Annual Meeting of the International Political Economy Society. University of Virginia, Charlottesville, VA.
- Rivers, Douglas. 2003. "Identification of Multidimensional Spatial Voting Models." Stanford University.

- Russett, Bruce and Donald R. Deluca. 1981. "Don't Tread on Me': Public Opinion and Foreign Policy in the Eighties." *Political Science Quarterly* 96(3):381–399.
- Sabet, Shahrzad. 2012. "Feelings First: Non-Material Factors as Moderators of Economic Self-Interest Effects on Trade Preferences." Annual Meeting of the International Political Economy Society. Charlottesville, VA.
- Scheve, Kenneth and David Stasavage. 2006. "Religion and Preferences for Social Insurance." *Quarterly Journal of Political Science* 1(3):255–286.
- Scheve, Kenneth F. and Matthew J. Slaughter. 2001. "What determines individual trade-policy preferences?" *Journal of International Economics* 54(2):267 – 292.
- Scheve, Kenneth and Matthew Slaughter. 2006. Public Opinion, International Economic Integration, and the Welfare State. In *Globalization and Egalitarian Redistribution*, ed. S. Bowles, P. Bardhan and M. Wallerstein. Princeton University Press.
- Shariff, Azim F. and Ara Norenzayan. 2007. "God Is Watching You: Priming God Concepts Increases Prosocial Behavior in an Anonymous Economic Game." *Psychological Science* 18(9):803–809.
- Shweder, R.A. 1991. *Thinking Through Cultures: Expeditions in Cultural Psychology*. Cambridge, MA: Harvard University Press.
- Skelton, J. A., Joan E. Loveland and Jennifer L. Yeagley. 1996. "Recalling symptom episodes affects reports of immediately-experienced symptoms: Inducing symptom suggestibility." *Psychology & Health* 11(2):183–201.
- Strull, T.K. and R.S. Wyer. 1979. "The role of category accessibility in the interpretation of information about persons: Some determinants and implications." *Journal of Personality and Social Psychology* 37:1660–1667.
- Stan Development Team. 2013. *Stan Modeling Language Users Guide and Reference Manual, Version 2.0*.
URL: <http://mc-stan.org/>
- Suhler, Christopher and Patricia Churchland. 2011. "Can Innate, Modular 'Foundations' Explain Morality? Challenges for Haidt's Moral Foundations Theory." *Journal of Cognitive Neuroscience* 23(9):2103–2116.
- Tetlock, Philip E. 2006. *Expert Political Judgment: How Good Is It? How Can We Know?* Princeton, NJ: Princeton University Press.
- Verba, Sidney, Richard A. Brody, Edwin B. Parker, Norman H. Nie, Nelson W. Polsby, Paul Ekman and Gordon S. Black. 1967. "Public Opinion and the War in Vietnam." *The American Political Science Review* 61(2):317–333.
- Weber, Christopher R. and Christopher M. Federico. 2013. "Moral Foundations and Heterogeneity in Ideological Preferences." *Political Psychology* 34(1):107–126.

Zaller, John. 1992. *The Nature and Origins of Mass Opinion*. New York: Cambridge University Press.

Zemack-Rugar, Yael, James R. Bettman and Gavan J. Fitzsimons. 2007. "The Effects of Nonconsciously Priming Emotion Concepts on Behavior." *Journal of Personality and Social Psychology* 93(6):927–939.

Figure 1: Correlation Among the MFQ Questions (Spearman's ρ)

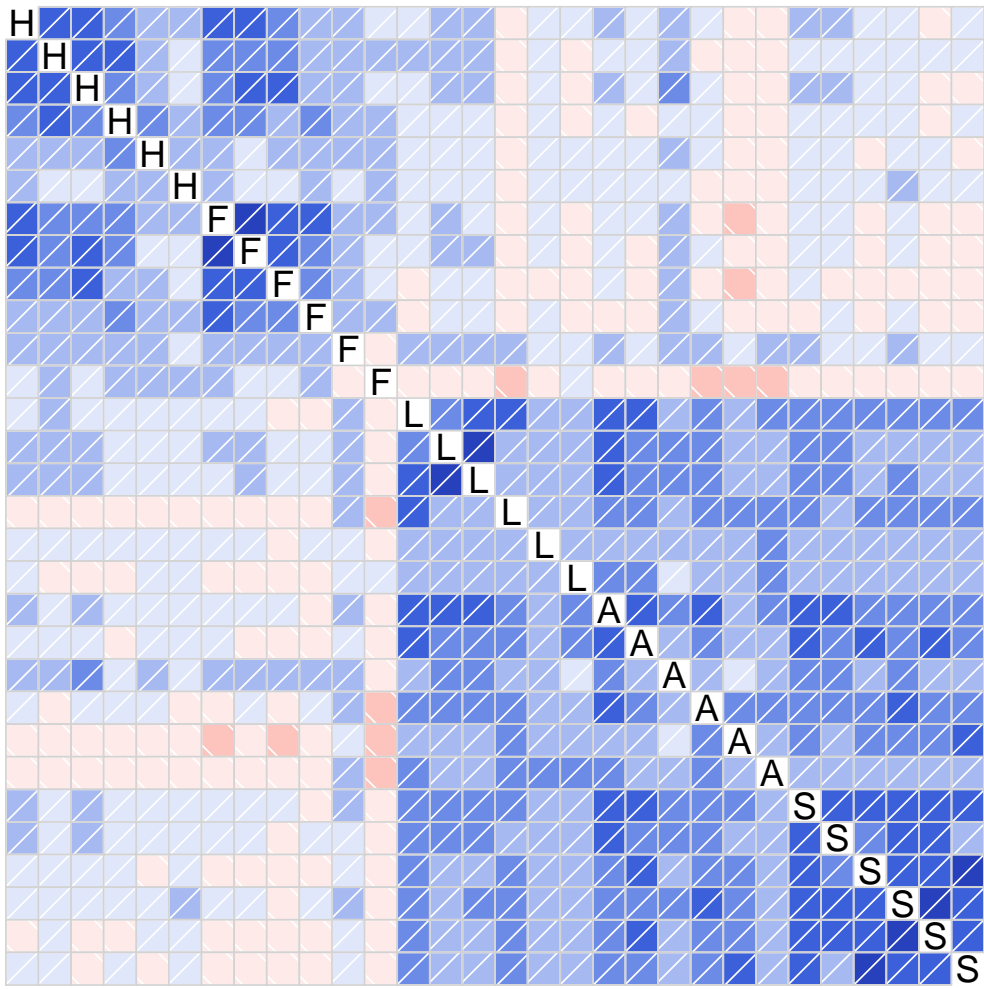


Figure 2: Ideology and Average Foundation Scores

Ideology and the Moral Foundations




Figure 3: Example Cutting Line for Use UN Question

Cutting Line for US Should Make More Decisions within the UN




Figure 4: Results of Two Dimensional Estimation and Clustering




Table 1: Ordered Probit Models of US Power Question

The statements below present possible goals for US foreign policy. Please indicate whether you think each goal should be very important, somewhat important, or not important to making US policy: Maintaining superior military power worldwide.

	Foundations Only	With Ideology	Full Controls
Harm	0.36 (0.06)***	0.32 (0.06)***	0.36 (0.06)***
Fairness	0.09 (0.06)	0.05 (0.06)	0.07 (0.07)
Ingroup	-0.34 (0.06)***	-0.32 (0.06)***	-0.32 (0.06)***
Authority	-0.27 (0.06)***	-0.22 (0.06)***	-0.23 (0.07)***
Purity	-0.18 (0.04)***	-0.12 (0.04)**	-0.12 (0.05)*
Ideology		-0.11 (0.03)***	-0.06 (0.04) ⁺
Gender			-0.03 (0.08)
Age			-0.02 (0.00)***
Party ID			-0.09 (0.05) ⁺
Political Interest			-0.16 (0.06)**
Education			0.04 (0.04)
Income			-0.06 (0.03) ⁺
Religiosity			-0.01 (0.03)
Tradable Sector			0.10 (0.14)
	Cut Points		
Very Important/Somewhat Important	-1.13 (0.21)	-1.51 (0.25)	-2.38 (0.34)
Somewhat Important/Not Important	0.23 (0.21)	-0.12 (0.24)	-0.91 (0.34)
Num. Obs.	1140	1113	1019
Log Likelihood	-1048.89	-1016.18	-902.13
Log Likelihood without Foundations	-1229.01	-1106.57	-979.78
Likelihood Ratio	360.24***	180.77***	155.31***

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, + $p < 0.1$

Distribution of Responses: Very Important–300; Somewhat Important–489; Not Important–376.

Table 2: Ordered Probit Models of Use UN Question

Please indicate whether you agree or disagree with the following statements: When dealing with international problems, the U.S. should be more willing to make decisions within the United Nations even if this means that the United States will sometimes have to go along with a policy that is not its first choice.

	Foundations Only	With Ideology	Full Controls
Harm	-0.17 (0.05) ^{***}	-0.09 (0.05) ⁺	-0.11 (0.06) ⁺
Fairness	-0.40 (0.06) ^{***}	-0.29 (0.06) ^{***}	-0.30 (0.06) ^{***}
Ingroup	0.03 (0.05)	0.01 (0.05)	0.03 (0.06)
Authority	-0.04 (0.06)	-0.09 (0.06)	-0.10 (0.06)
Purity	0.26 (0.04) ^{***}	0.14 (0.04) ^{***}	0.13 (0.04) ^{**}
Ideology		0.24 (0.03) ^{***}	0.22 (0.03) ^{***}
Gender			0.03 (0.07)
Age			0.01 (0.00) ^{***}
Party ID			0.03 (0.05)
Political Interest			0.08 (0.06)
Education			-0.05 (0.03)
Income			0.00 (0.03)
Religiosity			-0.01 (0.03)
Tradable Sector			-0.18 (0.13)
	Cut Points		
Agree Strongly/Agree Somewhat	-2.45 (0.20)	-1.48 (0.23)	-1.06 (0.32)
Agree Somewhat/Neither Agree nor Disagree	-1.27 (0.20)	-0.25 (0.23)	0.21 (0.31)
Neither Agree nor Disagree/Disagree Somewhat	-0.60 (0.20)	0.45 (0.23)	0.91 (0.32)
Disagree Somewhat/Disagree Strongly	-0.02 (0.20)	1.09 (0.23)	1.57 (0.32)
Num. Obs.	1119	1096	1005
Log Likelihood	-1567.35	-1487.17	-1343.43
Log Likelihood without Foundations	-1670.47	-1513.90	-1369.81
Likelihood Ratio	206.25 ^{***}	53.45 ^{***}	52.74 ^{***}

^{***} $p < 0.001$, ^{**} $p < 0.01$, ^{*} $p < 0.05$, ⁺ $p < 0.1$

Distribution of Responses: Agree Strongly–217; Agree Somewhat–434; Neither Agree nor Disagree–245; Disagree Somewhat–140; Disagree Strongly–107.

Table 3: Survey Results

Dependent Variable	Correlation with Ideology	Harm	Fairness	Ingroup	Authority	Purity
Goals of US Foreign Policy: Change in Probability of “Very Important”						
Strengthening the UN	0.278	2.5%	8.7%	3.6%	4.7%	0.5%
Protecting Human Rights	0.254	11.3%	6.4%	1.9%	-6.1%	1.2%
Protecting Weaker Nations Against Foreign Aggression	0.099	5.1%	1.4%	2.5%	-1.3%	2.4%
Promoting Democracy	-0.046	0.4%	1.6%	4.9%	-0.4%	1.2%
Stopping the Flow of Illegal Drugs into US	-0.260	-1.8%	2.3%	4.3%	7.4%	11.7%
Maintaining Superior US Military Power	-0.378	-12.6%	-2.1%	9.1%	6.2%	3.3%
Agree/Disagree Questions: Change in Probability of “Strongly Agree”						
US Should Make More Decisions within the UN	0.428	2.1%	4.8%	-0.6%	2.2%	-2.8%
Developed Countries Have Moral Obligation to Reduce Hunger	0.272	7.7%	5.4%	-0.4%	-3.0%	0.9%
Econ Power More important than Military power	0.264	3.7%	2.7%	-0.7%	-1.8%	-1.6%
US Should Provide Food and Medicine to Needy	0.254	6.1%	3.8%	-0.5%	-3.9%	2.2%
US Should Provide Development Aid	0.244	4.2%	3.6%	-0.2%	-2.5%	1.4%
Growing Trade and Business Ties With Other Countries are Good	0.102	1.9%	2.5%	1.0%	4.8%	-4.0%
US Should Provide Aid to Increase its Influence	-0.082	-0.4%	0.6%	1.9%	1.1%	0.3%
US Should Limit Import of Foreign Products	-0.095	-0.9%	0.6%	-0.4%	2.2%	3.1%
US Should Exert Strong Leadership in World	-0.140	0.1%	-0.4%	4.6%	4.3%	2.0%
US Should Limit Immigration By High-skill Workers	-0.218	-2.0%	0.5%	0.6%	1.7%	1.4%
US Should Limit Immigration By Low-skill Workers	-0.270	-2.4%	-0.3%	0.4%	3.2%	1.1%
War Sometimes Necessary to Maintain Order in World	-0.278	-10.2%	1.0%	4.3%	4.5%	1.7%
War Sometimes Necessary to Obtain Justice	-0.279	-10.8%	2.8%	3.7%	4.2%	0.5%

Our Culture Needs Protection Against Foreign Influence	-0.387	-2.1%	-1.8%	1.6%	1.6%	2.2%
Would you favor or oppose the use of US troops... (Probability of “Strongly Favor”)						
To Deal with Humanitarian Crises	0.179	6.6%	3.1%	1.5%	1.4%	0.8%
To Stop a Genocide	0.109	6.0%	7.3%	1.5%	-0.9%	1.9%
In Response to North Korean Invasion of SK	-0.055	0.4%	-1.1%	1.2%	0.4%	0.1%
In Response to Chinese Invasion of Taiwan	-0.106	-0.9%	0.0%	1.9%	-1.1%	0.4%
To Prevent Islamist Takeover of Pakistan	-0.110	-0.6%	0.4%	2.0%	0.6%	1.1%
To Ensure the Oil Supply	-0.312	-2.7%	-1.3%	3.7%	2.9%	1.3%

The numbers in the table represent the change in the predicted probability of giving the designated response for a one unit change in each foundation (from 2 to 3) with all other variables held constant at their means. Bolded, colored text indicates that the underlying ordered probit coefficient is statistically significant at $p < .05$. Green text indicates a positive effect, while red text indicates a negative effect.

All results come from an ordered probit specification with controls for: Ideology, Gender, Age, Party ID, Political Interest, Education, Income, Religiosity, and Tradable Sector.

Correlation with ideology is Spearman’s rank correlation coefficient for the uncontrolled, bivariate relationship between each dependent variable and respondent ideology.

Table 4: Care/Harm Priming Results for UN and US Power Questions

	Use UN (All)	Use UN (Lib)	Use UN (Cons)	Power (All)	Power (Lib)	Power (Cons)
Treated	0.17 (0.10)	0.01 (0.13)	0.37 (0.21)	-0.11 (0.11)	-0.08 (0.14)	-0.32 (0.23)
Gender	0.17 (0.10)	0.17 (0.14)	0.10 (0.21)	0.10 (0.11)	-0.03 (0.14)	0.23 (0.23)
Ideology	0.09** (0.03)	0.09 (0.09)	-0.22 (0.15)	-0.25*** (0.03)	-0.42*** (0.10)	-0.28 (0.16)
Education	-0.04 (0.04)	-0.04 (0.06)	-0.03 (0.08)	0.06 (0.05)	0.14* (0.07)	-0.09 (0.09)
	Cut Points					
Strongly/Somewhat Agree	-0.53* (0.22)	-0.64 (0.33)	-1.97* (0.87)			
Somewhat Agree/Neutral	0.60** (0.21)	0.61 (0.33)	-1.02 (0.86)			
Neutral/Somewhat Disagree	1.64*** (0.22)	1.68*** (0.34)	-0.17 (0.86)			
Somewhat/Strongly Disagree	2.24*** (0.24)	2.40*** (0.38)	0.57 (0.87)			
Very/Somewhat Important				-1.22*** (0.23)	-1.32*** (0.35)	-2.10* (0.92)
Somewhat/Not Important				0.01 (0.23)	-0.05 (0.35)	-0.78 (0.91)
AIC	1229.45	679.56	313.17	945.40	541.03	211.29
BIC	1262.53	708.32	334.32	970.22	562.53	227.27
Log Likelihood	-606.72	-331.78	-148.58	-466.70	-264.52	-99.65
Deviance	1213.45	663.56	297.17	933.40	529.03	199.29
Num. obs.	462	269	104	463	266	106

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, · $p < 0.1$

Table 5: Loyalty/Betrayal Priming Results for UN and US Power Questions

	Use UN (All)	Use UN (Lib)	Use UN (Cons)	Power (All)	Power (Lib)	Power (Cons)
Treated	0.03 (0.10)	-0.01 (0.12)	0.19 (0.20)	0.17 ⁺ (0.10)	0.10 (0.13)	0.49* (0.22)
Gender	0.21* (0.10)	0.24 ⁺ (0.13)	0.07 (0.20)	0.00 (0.10)	0.00 (0.13)	0.13 (0.22)
Ideology	0.11*** (0.03)	0.11 (0.08)	0.12 (0.13)	-0.18*** (0.03)	-0.13 (0.09)	-0.35* (0.14)
Education	-0.10* (0.04)	-0.05 (0.06)	-0.24** (0.08)	0.03 (0.05)	-0.05 (0.06)	0.21* (0.09)
	Cut Points					
Strongly/Somewhat Agree	-0.60** (0.21)	-0.49 (0.31)	-0.85 (0.82)			
Somewhat Agree/Neutral	0.44* (0.21)	0.66* (0.31)	-0.02 (0.81)			
Neutral/Somewhat Disagree	1.34*** (0.21)	1.65*** (0.32)	0.81 (0.82)			
Somewhat/Strongly Disagree	2.08*** (0.23)	2.31*** (0.34)	1.71* (0.83)			
Very/Somewhat Important				-1.10*** (0.22)	-1.33*** (0.33)	-1.14 (0.88)
Somewhat/Not Important				0.05 (0.22)	-0.11 (0.32)	-0.09 (0.88)
AIC	1374.23	786.16	351.30	1056.24	640.63	241.90
BIC	1407.97	815.87	373.33	1081.56	662.96	258.42
Log Likelihood	-679.12	-385.08	-167.65	-522.12	-314.32	-114.95
Deviance	1358.23	770.16	335.30	1044.24	628.63	229.90
Num. obs.	501	303	116	502	305	116

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, + $p < 0.1$

Table 6: Harm Priming Results

Dependent Variable	Full Sample	Liberals	Conservatives	ATT
Goals of US Foreign Policy: Change in Probability of “Very Important”				
Strengthening the UN	2.3%	-0.9%	-3.5%	0.8%
Protecting Human Rights	10.1%	12.7%	4.6%	10.6%
Protecting Weaker Nations Against Foreign Aggression	2.5%	4.6%	-3.8%	3.0%
Promoting Democracy	0.2%	-6.2%	1.3%	1.9%
Stopping the Flow of Illegal Drugs into US	-4.6%	-8.8%	-11.1%	-0.2%
Maintaining Superior US Military Power	3.6%	2.8%	9.2%	0.2%
Agree/Disagree Questions: Change in Probability of “Strongly Agree”				
US Should Make More Decisions within the UN	-4.9% ⁺	-0.2%	-5.9% ⁺	-0.1%
Developed Countries Have Moral Obligation to Reduce Hunger	3.4%	0.7%	3.1%	2.5%
Econ Power More important than Military power	-2.2%	-2.3%	-12.1%	-2.6%
US Should Provide Food and Medicine to Needy	0.6%	1.5%	0.7%	1.9%
US Should Provide Development Aid	-0.8%	-1.1%	-0.0%	-2.9%
Growing Trade and Business Ties With Other Countries are Good	1.6%	6.7%	-9.1%	2.1%
US Should Provide Aid to Increase its Influence	-0.2%	-1.0%	-2.1%	-1.1%
US Should Limit Import of Foreign Products	-1.5%	-2.0%	-0.1%	-0.8%
US Should Exert Strong Leadership in World	-1.6%	-2.7%	-1.0%	-4.0%
US Should Limit Immigration By High-skill Workers	0.5%	-3.4%	4.8%	0.6%
US Should Limit Immigration By Low-skill Workers	-1.5%	-2.6%	-0.8%	-1.6%
War Sometimes Necessary to Maintain Order in World	0.8%	0.7%	-2.5%	-0.2%
War Sometimes Necessary to Obtain Justice	-0.1%	-0.8%	-0.4%	-0.9%

Our Culture Needs Protection Against Foreign Influence	-0.4%	-1.7%	0.5%	-0.5%
Would you favor or oppose the use of US troops... (Probability of “Strongly Favor”)				
To Deal with Humanitarian Crises	-1.6%	-4.1%	1.3%	-2.2%
To Stop a Genocide	-1.1%	-4.3%	-4.5%	-0.8%
In Response to North Korean Invasion of SK	2.4%	-4.1%	5.0%	-3.4%
In Response to Chinese Invasion of Taiwan	0.2%	-0.7%	-1.1%	-2.3%
To Prevent Islamist Takeover of Pakistan	-0.0%	-1.0%	-0.5%	-2.0%
To Ensure the Oil Supply	1.9%	0.8%	3.9%	1.8%

The numbers in the table represent the change in the predicted probability of giving the designated response for members of the treatment group. Bolded, colored text indicates that the underlying ordered probit coefficient is statistically significant at $p < .05$. Green text indicates a positive effect, while red text indicates a negative effect. The symbol + designates significance at $p < .10$

The ATT column codes only respondents who used the target word in the treatment condition as treated. All results come from an ordered probit specification with controls for: ideology, gender, and education.

Table 7: Ingroup Priming Results

Dependent Variable	Full Sample	Liberals	Conservatives	ATT
Goals of US Foreign Policy: Change in Probability of “Very Important”				
Strengthening the UN	0.0%	2.8%	-13.7%	-5.7%
Protecting Human Rights	4.7%	6.5%	-4.0%	2.2%
Protecting Weaker Nations Against Foreign Aggression	3.9%	5.4%	-2.5%	1.6%
Promoting Democracy	2.9%	3.4%	-1.5%	-0.8%
Stopping the Flow of Illegal Drugs into US	2.2%	5.1%	0.9%	4.3%
Maintaining Superior US Military Power	-5.2% ⁺	-2.8%	-11.3%	-5.6%
Agree/Disagree Questions: Change in Probability of “Strongly Agree”				
US Should Make More Decisions within the UN	-1.0%	0.3%	-6.4%	-1.1%
Developed Countries Have Moral Obligation to Reduce Hunger	-3.6%	-4.5%	-9.6%	-5.9% ⁺
Econ Power More important than Military power	-0.1%	1.4%	-15.3% ⁺	-1.3%
US Should Provide Food and Medicine to Needy	-1.3%	-3.1%	1.5%	-4.0%
US Should Provide Development Aid	-0.1%	-1.9%	1.8%	-5.7%
Growing Trade and Business Ties With Other Countries are Good	0.3%	-2.6%	-2.3%	-6.9% ⁺
US Should Provide Aid to Increase its Influence	0.4%	-0.1%	1.5%	-2.5%
US Should Limit Import of Foreign Products	-1.7%	-0.7%	-8.0% ⁺	1.1%
US Should Exert Strong Leadership in World	-4.5% ⁺	-4.8% ⁺	-3.8%	-4.7% ⁺
US Should Limit Immigration By High-skill Workers	-0.6%	-1.5%	1.2%	1.6%
US Should Limit Immigration By Low-skill Workers	-1.5%	0.4%	-6.6%	0.5%
War Sometimes Necessary to Maintain Order in World	-3.8%	-5.5%	1.7%	-2.7%
War Sometimes Necessary to Obtain Justice	-4.0% ⁺	-3.5%	-5.6%	-2.7%

Our Culture Needs Protection Against Foreign Influence	-2.5%	-3.4%	-2.6%	-2.4%
Would you favor or oppose the use of US troops... (Probability of “Strongly Favor”)				
To Deal with Humanitarian Crises	1.3%	0.9%	3.1%	1.2%
To Stop a Genocide	-1.6%	-3.5%	-5.9%	-1.3%
In Response to North Korean Invasion of SK	-0.2%	-0.6%	-0.8%	-0.7%
In Response to Chinese Invasion of Taiwan	1.2%	3.2% ⁺	-0.5%	2.3%
To Prevent Islamist Takeover of Pakistan	0.1%	1.0%	-1.2%	0.3%
To Ensure the Oil Supply	-0.0%	-1.0%	-0.8%	-1.1%

The numbers in the table represent the change in the predicted probability of giving the designated response for members of the treatment group. Bolded, colored text indicates that the underlying ordered probit coefficient is statistically significant at $p < .05$. Green text indicates a positive effect, while red text indicates a negative effect. The symbol ⁺⁺ designates significance at $p < .10$ and the symbol ⁺ designates significance at $p < .20$

The ATT column codes only respondents who used the target word in the treatment condition as treated. All results come from an ordered probit specification with controls for: ideology, gender, and education.