

To: Participants at IPES Conference
From: Michael Tomz (tomz@stanford.edu)
Re: The Effects of International Agreements on Preferences and Beliefs
Date: November 2006

At the IPES Conference, I will present initial findings from a new project about the effects of international agreements. The project uses survey-based experiments to study how treaties (and other international agreements) change the preferences and beliefs of voters and policymakers.

The enclosed paper describes findings from one experiment, which was administered to a random sample of US voters.

My oral presentation will also cover experiments that were administered to members of the British Parliament last summer.

The project, recently funded by a five-year CAREER grant from the National Science Foundation, is still at an early stage. This is precisely the moment when advice from the IPES community could make a major difference. I look forward to your comments and suggestions!

The Effects of International Agreements on Foreign Policy Preferences

Michael Tomz
Stanford University
tomz@stanford.edu

November 2006

Preliminary. Comments welcome!

I am grateful for financial support from the National Science Foundation (CAREER grant SES-0548285). The staff at Knowledge Networks—especially Mike Dennis and Sergei Rodkin—provided invaluable assistance with this project.

Abstract: Do international agreements transform the way citizens think about foreign policy, or are such agreements mere scraps of paper that have little effect on policy preferences? I investigate this question by offering the first-ever experimental analysis of international agreements. The experiments, embedded in public opinion surveys, directly measure the effect of international agreements on foreign policy preferences and, at the same time, avoid problems of endogeneity that have stymied previous research. Citizens, I find, are far more likely to oppose policies that would violate international legal agreements than to oppose *otherwise identical* policies that would not trammel upon existing pacts. Moreover, the political effects of international agreements diffuse widely throughout the population, affecting conservatives as well as liberals and crossing other demographic divides. But citizens do not follow international legal agreements blindly. When the case for pursuing a policy is strong enough, a clear majority will endorse the policy, even though it would violate international law. Thus, international legal agreements strongly influence—but do not dictate—preferences about foreign policy.

Under what conditions do international legal agreements affect the behavior of states?

Few questions are as central to the study and practice of global politics. According to the United Nations Treaty Series, more than 50,000 international agreements are now in force. Although these agreements are “legally binding,” sovereign states have no higher power to which they can appeal when another party reneges. The tension between legal theory and international anarchy has stimulated much debate about whether international law is irrelevant, or whether it truly influences countries’ policies.

Previous research has not settled the debate. As Simmons (1998: 89) points out, “Several studies have tried to demonstrate a correlation between legal standards and state behavior, sometimes employing large databases and statistical techniques, but most are unconvincing in demonstrating causation, or even in providing an explanatory link between the actions taken and the existence of agreements or normative considerations.” This ambiguity is not necessarily the fault of political scientists or lawyers; rather, it stems from the challenging nature of the subject itself. As most contributors to the literature now recognize, problems of endogeneity make it extremely difficult to know whether, and in what ways, international agreements influence state behavior.

In this paper, I advance our understanding of international agreements by supplementing the analysis of observational data with experiments involving randomized treatment and control. A pure field experiment, in which the researcher forces some governments to sign international agreements and others to abstain, is clearly out of the question, but other interventions are both feasible and informative. I summarize the outcome of a survey-based experiment that directly measured the effect of international agreements and, at the same time, avoided problems of endogeneity that have stymied previous research.

Data from the experiment support three conclusions. First, international legal agreements transform policy preferences. Citizens are far more likely to oppose policies that would violate international law than to oppose *otherwise identical* policies that would not trammel upon existing treaties and other agreements. Second, the political effects of international law diffuse widely throughout the population, affecting conservatives as well as liberals and crossing other demographic divides. Third, citizens do not follow international law blindly. When the case for pursuing a policy is strong enough, a clear majority will endorse the policy, even though it would violate international law. In summary, international law shapes—but does not dictate—the foreign policy preferences of citizens.

In the remainder of this paper, I discuss the limits of existing evidence, present the findings from a survey-based experiment, and conclude by enumerating the strengths and weaknesses of my experimental approach and discussing avenues for future research.

1. The Limits of Existing Evidence

International relations scholars disagree about the effect of formal international agreements. Some say that such agreements have little independent effect on foreign policy (e.g. Mearsheimer 1994, 2001). In a world of sovereign states, participation in treaties and other legal pacts is entirely voluntary. Countries sign agreements they are predisposed to follow, and they remain parties while the agreement serves their interests. If the commitment ever becomes inconvenient, however, countries may withdraw unilaterally. In fact, most treaties allow parties to end their legal obligations by giving just a few months' advance notice. When a graceful exit is not possible, parties can break the agreement without fear of punishment by a world government.

Other analysts argue that formal agreements shape the way countries behave (e.g. Keohane 1984; Fortna 2003).¹ According to this view, when leaders sign an international agreement, it becomes more costly to take actions the agreement forbids and less costly to pursue policies the agreement condones. Why might the agreement change incentives? One plausible mechanism is reputation: putting an agreement in writing may increase the reputational cost of reneging. A second mechanism concerns norms, rather than interests: perhaps people feel it would be morally wrong to break promises to foreigners. Under either of these two mechanisms, international agreements would shape the way countries behave.

Due to limitations of existing data, it has been extremely difficult to resolve the debate. We simply do not know the conditions under which international agreements matter, or when the act of signing an international agreement is most likely to demonstrate credibility.

Previous research has relied entirely on the historical record to estimate the effect of international agreements. Have countries that signed environmental protection treaties polluted less than countries that did not sign? Has respect for human rights, arms control, and free trade been greater among countries that entered agreements on these topics than among countries that did not?

If agreements arose from a purely random process, the use of historical data would be unproblematic. In reality, though, countries choose whether and on what terms to enter international agreements. Thus, it is hard to know whether the historical correlation between agreements and behavior is a consequence of the agreement itself, or is due to cross-national differences in the baseline propensity to take the kinds of actions the agreement requires. Most countries honor their agreements most of the time, but this does not prove that agreements shape

¹ For an excellent review of the literature as applied to international trade and the GATT/WTO, see Busch and Reinhardt (2002).

foreign policy. According to skeptics, agreements reflect but do not change the pre-existing interests of states.

We can bring the existing debate—and the roadblock of endogeneity—into sharper relief by drawing on Rubin’s (1974) counterfactual account of causality. Suppose we are interested in Y_i , a measure of country i ’s behavior on a given issue at a particular time. The impact of an international agreement on Y_i is $\delta_i = Y_{i1} - Y_{i0}$, where Y_{i1} represents the way i would behave if party to the agreement, and Y_{i0} signifies how the same country would behave if not party to the agreement. The quantity δ_i tells what difference, if any, the agreement makes.

Unfortunately, the causal effect δ_i is unobservable. We can imagine how country i might behave in both the agreement and the no-agreement conditions, but we cannot observe both Y_{i1} and Y_{i0} for the same i at the same time. After all, no country can be in the treatment and control regimes simultaneously. Holland (1986) has called the inability to observe the same unit under both treatment and control “the fundamental problem of causal inference.”

Although the individual-level effect δ_i is beyond reach, scholars have tried to infer *the average* causal effect δ by comparing the observed Y for countries that signed the agreement with the observed Y for countries that did not. In this context, the average effect of the agreement is estimated as $\hat{\delta} = (Y_1 | A = 1) - (Y_0 | A = 0)$, where the indicator A takes a value of 1 when the country is party to the agreement and 0 otherwise.

This standard estimator is equal to the true effect δ plus two potential sources of bias. The first source of bias, “different baseline propensities,” arises when signatories and non-signatories differ in their fundamental tendency to do Y , even in the absence of an international agreement. The second source of bias, “different treatment effects,” arises when the agreement

would produce stronger (or weaker) effects on the group that actually signed than on the group that did not.

With a bit of algebra, we can show that

$$\hat{\delta} = \delta + [(Y_0 | A = 1) - (Y_0 | A = 0)] + \pi[(\delta | A = 1) - (\delta | A = 0)],$$

where π is the proportion of the sample that did not join the agreement, $\delta | A = 1$ is the average effect on those who actually signed, and $\delta | A = 0$ is the average effect the agreement would have exerted on the remaining countries, if contrary to fact they had chosen to sign. The first term in brackets gives the bias from different baseline propensities; the second gives the bias from heterogeneous treatment effects.

Expressing $\hat{\delta}$ in this way helps clarify the claims and the limitations of the existing literature. Researchers have found that countries often comply with international agreements, and in some cases the observed level of Y differs systematically between signatories and non-signatories. To the skeptic, these estimated effects are artifacts of different baselines: those who signed were more inclined to do Y in the first place.

Essentially, skeptics argue that that δ is approximately zero but our estimate $\hat{\delta} \neq 0$ because the conditional mean $(Y_0|A=1)$ exceeds the conditional mean $(Y_0|A=0)$. If we could eliminate baseline differences between the two groups, the skeptic argues, the apparent effect of the international agreement would disappear. To convince the skeptic that international agreements matter, and to obtain unbiased estimates of the causal effect more generally, it is important to remove any baseline differences. I argue below that experiments can achieve this goal by design, whereas observational studies can do so only with great difficulty.

Against the skeptics, legalists argue that a country can alter its interests and behavior by signing an international agreement. Moreover, sophisticated legalists hypothesize that

agreements affect different countries to different degrees. For example, Hathaway (2005) and Raustiala and Slaughter (2002) suggest that the relationship between international law and state behavior depends on domestic institutions such as the judiciary, the media, political parties, and interest groups. Others claim that sensitivity to agreements varies with the rule of law or the degree of democracy. These and other domestic institutions vary considerably across states. As a consequence, δ_i should differ from one i to the next.

Heterogeneity in δ_i creates a second source of bias: a correlation between signatories and susceptibility to treatment. If leaders are rational, they will weigh the anticipated effects of the agreement when deciding whether to sign. This rational behavior introduces a systematic relationship between signatory status (A_i) and the treatment effect (δ_i). Using the previous notation, $(\delta | A = 1) \neq (\delta | A = 0)$ in the presence of self-selection. This heterogeneity is, of course, interesting in its own right and an important subject of study. Unless the heterogeneity is controlled, however, $\hat{\delta}$ will be a biased estimate of the treatment effect.

To draw valid inferences from non-experimental data, we need statistical correctives that allow us to approximate the attributes of a genuine experiment. Some researchers address this problem with control variables: they model foreign policy as a function of international agreements and controls that correct for differences between signatories and non-signatories (e.g. Simmons 2000). The goal is to make signatories and non-signatories comparable after conditioning on the X 's, such that any systematic difference in behavior would reflect the causal effect of the agreement, rather than distinct baselines or different sensitivities to treatment.²

² Heckman-type selection models can be viewed variants of the same strategy: using a function of one or more control variables to address the problem of endogeneity. For applications of this approach to IMF agreements, see Przeworski and Vreeland (2000) and Vreeland (2003).

The success of this approach depends on a comprehensive set of controls, however. To solve the bias problem with control variables, the researcher must condition on all variables that correlate with the outcome and membership in the agreement (Besley and Case 2000). This can be quite a challenge, made more severe by informational asymmetries in international relations. Governments have private information, which they withhold not only from other countries but (presumably) from academic researchers, as well! When governments have pertinent but private information about their baseline interests or their sensitivity to treatment, the set of control variables is likely to be incomplete and estimates of the agreement's effect will be biased.

Some of these problems can be minimized through the use of panel data. If countries are tracked over a number of years, the insertion of fixed effects for countries or dyads can help correct for unobserved heterogeneity. Hathaway (2002), Goldstein, Rivers and Tomz (2007), Simmons (2004), and Tomz, Goldstein, and Rivers (2005) have used this approach to study the effects of international agreements on human rights and international trade. For the strategy to succeed completely, though, the determinants of state policy must be additive and time-invariant. If the decision to enter an agreement and the choice of Y depend on common variables that change over time, omitting those variables will lead to biased estimates of δ , even in the presence of fixed effects (Besley and Case 2000).

There is a second statistical option. Instead of using controls that are correlated with both the treatment and the outcome, one could seek an instrument that affects the outcome only indirectly, via the treatment variable. With a technique such as two-stage least squares regression, the instrument can be used to obtain consistent estimates of the treatment effect. This approach makes sense in theory but has serious problems in practice: it has been nearly

impossible to find valid instruments: ones that correlate strongly with the presence or absence of an international agreement but have no independent bearing on foreign policy.

In summary, problems of endogeneity make it difficult to infer the effects of international agreements. Several studies have tried to address this problem via control variables, but many questions about causality remain. The next section explains how experiments can isolate the causal effects of international agreements.

2. An Experiment-Based Approach

The core idea of this paper is to supplement observational studies by embedding experiments in interviews with citizens and elites. Some interviewees hear about a hypothetical or historical foreign policy situation in which leaders have signed a legally binding international agreement. Others consider exactly the same situation, *sans* any international agreement. By comparing the views of participants in the treatment condition (international agreement) versus the control condition (no agreement), we can isolate the effect of international agreements on policy preferences and beliefs.

As a first step toward implementing this approach, I designed an experiment about international trade agreements. The experiment aimed to quantify the effect of trade pacts *in the context* of many other foreign policy considerations, such as humanitarianism and the national interest. This represents a relatively hard test. If trade agreements matter even in the face of competing or redundant concerns, they are likely to be influential in less stringent settings.

The experiment, which was administered over the internet to a nationally representative sample of 1000 U.S. adults in July 2005,³ began as follows: “The next question is about foreign

³ The experiment was fielded by Knowledge Networks, and interviews were conducted over Web TV and the internet.

policy. Some leaders want the United States to prohibit trade with the country of Burma. They say we should neither buy products from Burma nor sell products to Burma. Experts who have studied this proposal agree on several points. Please consider each point carefully, and then tell us what you think.”

I presented each respondent with two or more of the following points.

US economy: “The proposal would help the U.S. economy. Many Americans are getting laid off because of competition from Burma. If we stop trading with Burma, there will be more jobs and higher wages in the United States.”

Human rights: “The proposal would help human rights. In Burma, the government kills political opponents and does not allow free speech. By stopping trade with Burma, we can pressure the government to start respecting basic rights.”

Burmese economy: “The proposal would hurt the Burmese economy. Burma sells \$300 million in products to the United States each year. If we stop trading with Burma, people in that country will lose their jobs, and poverty will rise.”

International law: “The proposal would violate international law. The United States has signed treaties that make it illegal to limit trade with Burma. If we stop trading with Burma, we will be breaking international law.”

Neutral argument: “The proposal would change our trade relations. The United States trades with many countries. If we stop trading with Burma, we will no longer suffer the costs (if any) nor will we get the benefits (if any) of trade with that particular country.”

I randomly assigned each respondent to one of nine groups, each of which confronted a different configuration of arguments (see Table 1). For example, respondents in group 1 learned that the proposal to cut off trade would help the US economy but hurt the Burmese economy. Participants in group 2 received exactly the same considerations, plus they were told that the plan would violate international law by contravening trade agreements the United States had signed. This between-subject design makes it possible to measure the effect of international agreements by comparing the preferences of citizens in group 1 versus group 2. The third group received a “neutral” argument, which I included to quantify the influence of international law while holding the number of arguments constant (e.g. compare group 2 versus group 3).

[TABLE 1 ABOUT HERE]

The remaining entries in Table 1 altered the mix of “pro” arguments. Conditions 4-6 contained no mention of the US economy, but they stressed that the plan would help human rights by putting pressure on the Burmese economy. Finally, I posed a double-challenge to trade agreements by constructing scenarios in which the decision to eliminate trade would both serve US economic interests and advance human rights (groups 7-9).⁴ To guard against any order effects, I randomized not only the group assignments but also the sequence of arguments within

⁴ By design, all respondents heard that the plan would hurt the Burmese economy. I included the factor to balance and fill-out the list of considerations.

each group. For instance, half the participants in group 1 received the argument about the US economy before the argument about the Burmese economy; the other half considered the opposite succession of points.

After presenting these arguments, I asked citizens to express their foreign policy preferences. “How good or bad an idea is it for the United States to prohibit trade with Burma?” The response options were: extremely good, moderately good, slightly good, neither good nor bad, slightly bad, moderately bad, and extremely bad. By analyzing the responses, one can see how the presence or absence of an international legal agreement affected policy preferences.

Before computing the effect of such agreements, it seemed prudent to confirm that the treatment and control groups were balanced on baseline covariates that could affect foreign policy preferences. I estimated a logistic regression in which the dependent variable was the dichotomous treatment (international law mentioned=1, not mentioned=0) and asked whether any demographic or contextual variables predicted membership in the treatment group. Not one of the many variables in the model—ideology, party identification, gender, age, race, education, income, and a variety of other demographic factors—had a statistically significant effect on the probability of being in the treatment group. Based on a likelihood ratio test, we cannot reject the hypothesis that the relationship between the treatment and *all* baseline variables was zero.⁵

Having established that the treatment was random, I proceeded to estimate the effect of international agreements. Of respondents who took a side (all respondents except those who answered “neither good nor bad”), I computed the share who thought it would be a bad idea to sever commercial relations with Burma. To the extent that international agreements matter, this

⁵ The likelihood ratio test statistic, 22.26, was distributed chi-squared with 23 degrees of freedom. If all coefficients were zero, we would observe a test statistic that large roughly half the time.

share should be significantly higher when international law is mentioned (groups 3, 6, and 9) than when it is not.

In all cases, I obtained point estimates and confidence intervals via Bayesian simulation. Specifically, I modeled the proportion of people who opposed the policy proposal as a beta distribution with a uniform prior (see Johnson and Albert 1999, 11). When an international agreement exists, this proportion $\pi_1|data$ is distributed as $\text{beta}(b_1+1, N_1 - b_1 + 1)$, where the subscript 1 indicates the treatment regime, N_1 is the number of respondents who received the treatment, and b_1 is the number of people who thought it would be bad to cut trade with Burma.

Without an international agreement, the proportion of naysayers is $\pi_0|data$ and distributed as $\text{beta}(b_0+1, N_0 - b_0 + 1)$, where the subscript 0 signifies the control regime. By drawing random variates from these independent beta distributions, we can obtain the full posterior distribution (and therefore point estimates and confidence intervals) of $\delta = \pi_1 - \pi_0$, the effect of the international agreement. I employ this approach in the remainder of the paper. Other statistical methods, including ordered probit analysis and comparison of means using the full seven-point scale from extremely bad to extremely good, lead to the same conclusions.

Table 2 shows that international legal agreements powerfully affect the preferences of citizens. When no international agreement was mentioned (conditions 1, 3, 4, 6, 7, and 9), approximately 27 percent of respondents who took a side deemed it bad to prohibit trade with Burma. This percentage jumped 17 points when respondents were told that the policy would violate international law. The 95 percent confidence interval around this effect ranged from 10 to 25, so we can be quite sure that the shift in policy preferences did not arise by chance alone. The data in Table 2 thus provide strong microfoundations for the view that international legal agreements do not simply reflect, but can actually change, preferences about foreign policy.

[TABLE 2 ABOUT HERE]

Table 3 divides the sample into three parts, based on non-legal considerations that respondents weighed. The table supports three inferences. First, on matters of international trade, economic arguments tend to sway citizens more effectively than appeals to human rights. When experts concluded that the trade barriers would improve human rights, 37 percent of respondents who took a side nonetheless opposed the measure (Table 3, column 1). Citizens who heard about the U.S. economy, rather than human rights, were significantly more likely to support the proposal.

Second, appeals to human rights and economic interest reinforce each other. Citizens do not regard these considerations as equivalent, and they are more willing to support a foreign policy that would serve both objectives than a policy that would enhance only one. In fact, only 15 percent of citizens who expressed an opinion actually disapproved when told that we could create U.S. jobs and improve human rights by eliminating trade with Burma.

[TABLE 3 ABOUT HERE]

Third, the table confirms that international law can change foreign policy preferences, even in the face of powerful counterarguments. Consider the first row of Table 3, which pertains to people who heard that the policy initiative would help human rights. For those citizens, the mere insertion of an international agreement converted naysayers from a minority (37 percent) to a majority (54 percent). The estimated effect of international law in this situation was 17 percentage points, with an unambiguously positive confidence interval. The second row of the table summarizes opinion among citizens who heard that the proposal would help the U.S.

economy. Here, international law exerted a smaller but still discernable effect, leading to a 10-percentage-point swing in policy preferences. This estimate, though somewhat imprecise, was nonetheless greater than zero with a Bayesian p-value of 0.06.

Surprisingly, international law remained potent even when counter-arguments were extremely strong (Table 3, row 3). The proposal to create jobs and improve human rights was highly popular among members of the control group, but the same proposal elicited scorn from 37 percent of citizens in the treatment condition. Under this scenario, the presence of an international agreement more than doubled the share of citizens who opposed the foreign policy. The overall effect was 22 points, with a 95-percent confidence interval from 9 to 35. Thus, contrary to the skeptics, it appears that international legal agreements can affect preferences, even when the decision to follow international law might not serve either economic self-interest or humanitarian concerns.

Further analysis suggests that the political effects of international law diffuse widely throughout the population, affecting conservatives as well as liberals and crossing other demographic divides. Table 4 presents a breakdown of estimates by demographic group. Regardless of political ideology, party identification, gender, education or income, international law substantially changes preferences about foreign policy. The effect sizes in the table range from 12 to 25 percentage points, and all are distinguishable from zero with a Bayesian p-value of .001 or better. Moreover, although some demographic groups appear more sensitive to international law than other groups, we cannot affirm these differences with a high level of confidence. Overall, we cannot reject the null hypothesis that international law has similar effects on all demographic groups.

3. Conclusions

In at least two ways, the experiments in this paper shed new light on the effects of international legal agreements. First, the experiments overcame problems of endogeneity that have hampered previous studies. In the experiments, I assigned treatment and control randomly without reference to background features of the situation or the respondent. As a result, there was no significant correlation between the agreement and baseline propensities or sensitivity to treatment. This greatly simplified the problem of inference: I obtained unbiased estimates of the agreement's effect through tabular analysis of the experimental data. Randomization eliminated the need for scores of regressions with control variables, which must be used in observational studies to balance the treatment and control groups.

Second, the experiments helped strengthen the microfoundations of international legal theory by revealing how legal agreements affect the preferences of individuals. Researchers have found correlations between international agreements and policy outcomes. Some contend that the relationship is spurious, whereas others argue that international agreements change the cost/benefit calculations of citizens. The evidence in this paper supports the second interpretation. At an individual level, citizens are more reluctant to pursue policies that would violate international law than to pursue otherwise identical policies that are not enshrined in a legal commitment.

The experimental method in this paper could be extended to answer a wide range of questions about international agreements. The prevailing research strategy, which relies entirely on observational data, is essentially passive. To study a particular variable, researchers must wait for natural processes to generate the variation they need, in quantities large enough to support statistical analysis or in patterns convenient enough to permit controlled case studies. A

passive strategy has significant limitations. Some factors may exhibit minimal variation or be highly collinear with other factors, and some values may occur too rarely to support precise estimates. An experimental approach can overcome these limitations by allowing full control over the explanatory variables.

In particular, subsequent studies could vary not only the presence but also the form of the international agreement, thereby shedding light on the effects of institutional design (Koremenos, Lipson and Snidal 2001). Lipson (1991) hypothesizes that the costs of renegeing increase with the precision of the agreement, the formality by which it was conveyed, and level of government that authorized it. Rosendorff and Milner (2001) add that the penalty for deviating from commitments can be lower in the presence of escape clauses. Finally, my work on international debt (Tomz 2007) shows that lenders will excuse defaults that occur because of a fundamental change in circumstances or widespread noncompliance by other parties. With experiments, we can test whether citizens and elites take similar contingencies into account when thinking about treaties and other international agreements.

Of course, the experimental approach is not infallible. Indeed, experiments are vulnerable on precisely the dimension where observational data is most compelling: external validity. Voters and elites might behave differently in an interview than in real foreign policy situations. Differences could emerge because respondents know they are subjects of a study, because the interviewer can offer only limited background information, or because emotion plays a different role in interviews than in actual politics. It is not clear how serious these problems would be, but it is important to acknowledge them.

To some extent, concerns about external validity can be minimized by making the scenario as convincing as possible and replicating the experiments with different question

wording and sample frames, to increase confidence in the generality of the results. Ultimately, though, the evidence from experiments should be combined with observational data to obtain a fuller understanding of international agreements. Every methodology has its limitations. The best way to make progress on complicated topics is to analyze data from multiple sources. Thus, the evidence in this paper complements a growing body of high-quality research that others have done with historical data.

Works Cited

- Besley, Timothy and Anne Case. 2000. "Unnatural Experiments? Estimating the Incidence of Exogenous Policies." *Economic Journal* 110 (November): F672-94.
- Busch, Marc L. and Eric Reinhardt. 2002. "Testing International Trade Law: Empirical Studies of GATT/WTO Dispute Settlement." In Daniel L. M. Kennedy and James D. Southwick, eds., *The Political Economy of International Trade Law* (New York: Cambridge Univ. Press, 2002): 457-81.
- Fortna, Virginia Page. 2003. "Scraps of Paper? Agreements and the Durability of Peace," *International Organization* 57:337-72.
- Goldstein, Judith, Douglas Rivers, and Michael Tomz. 2007. "Institutions in International Relations: Understanding the Effects of the GATT and the WTO on World Trade." *International Organization* 61: 37-67.
- Hathaway, Oona A. 2002. "Do Human Rights Treaties Make a Difference?" *Yale Law Journal* 111: 1935.
- Hathaway, Oona A. 2005. "Between Power and Principle: An Integrated Theory of International Law." *University of Chicago Law Review* 72: 469-536.
- Holland, Paul W. 1986. "Statistics and Causal Inference." *Journal of the American Statistical Association* 81, no. 396 (December): 945-60.
- Johnson, Valen E. and James H. Albert. 1999. *Ordinal Data Modeling*. New York: Springer.
- Keohane, Robert. 1984. *After Hegemony*. Princeton, NJ: Princeton Univ. Press.
- Koromenos, Barbara, Charles Lipson and Duncan Snidal. 2001. "The Rational Design of International Institutions," *International Organization* 55, no. 4 (Autumn): 761-99.
- Lipson, Charles. 1991. "Why are Some International Agreements Informal?" *International Organization* 45, no. 4 (Autumn): 495-538.
- Mearsheimer, John J. 1994. "The False Promise of International Institutions," *International Security* 19:5-49.
- Mearsheimer, John J. 2001. *The Tragedy of Great Power Politics*. New York: Norton.
- Przeworski, Adam and James Raymond Vreeland. 2000. "The Effect of IMF Programs on Economic Growth." *Journal of Development Economics* 62: 385-421.
- Raustiala, Kal and Anne-Marie Slaughter. 2002. "International Law, International Relations and Compliance." *Handbook of International Relations*, pp. 538-58. NY: Sage.

- Rosendorff, B. Peter and Helen V. Milner. 2001. "The Optimal Design of International Trade Institutions: Uncertainty and Escape." *International Organization* 55, no. 4 (Autumn): 829-57.
- Rubin, Donald. 1974. "Estimating Causal Effects of Treatments in Randomized and Non-Randomized Studies." *Journal of Educational Psychology* 66: 688-701.
- Simmons, Beth. 1998. "Compliance with International Agreements." *Annual Review of Political Science* 1:75-93.
- Simmons, Beth. 2000. "International Law and State Behavior: Commitment and Compliance in International Monetary Affairs." *American Political Science Review* 94, no. 4 (December): 819-35.
- Simmons, Beth. 2004. "International Law Compliance and Human Rights." Working Paper, Harvard University.
- Tomz, Michael, Judith Goldstein and Doug Rivers. 2005. "Membership Has Its Privileges: Understanding the Effects of the GATT and the WTO on World Trade." Working paper, Stanford Univ.
- Tomz, Michael. 2007. *Sovereign Debt and International Cooperation*. Princeton, NJ: Princeton Univ. Press.
- Vreeland, James. 2003. *The IMF and Economic Development*. New York: Cambridge University Press.

Table 1: Experimental Conditions

Each respondent was assigned to one of the following groups and presented with the arguments marked by X's.

Group	U.S. economy	Human rights	Burmese economy	International agreement	Neutral argument	Sample size
1	X		X			128
2	X		X	X		114
3	X		X		X	125
4		X	X			114
5		X	X	X		103
6		X	X		X	99
7	X	X	X			113
8	X	X	X	X		88
9	X	X	X		X	116

Table 2: The Effect of International Law

Table gives the percent of respondents who opposed cutting trade with Burma
Bayesian 95% confidence intervals appear in parentheses.

Violates International Law?		
No	Yes	Effect
27 (23 to 31)	44 (38 to 51)	17 (10 to 25)

Table 3: Effect of International Law, Conditional on Counter-Arguments
 Table gives the percent of respondents who opposed cutting trade with Burma
 Bayesian 95% confidence intervals appear in parentheses.

	Violates International Law?		
	No	Yes	Effect
Human rights only	37 (30 to 45)	54 (43 to 65)	17 (4 to 30)
U.S. economy only	30 (23 to 37)	40 (30 to 51)	10 (-3 to 23)
Both arguments	15 (10 to 21)	37 (26 to 49)	22 (9 to 35)

Table 4: The Effect Cuts Across Demographic Divisions

Table gives the percent of respondents who opposed cutting trade with Burma
 Bayesian 95% confidence intervals appear in parentheses.

	Violates International Law?			Difference
	No	Yes	Effect	
Liberals	28 (21 to 36)	53 (41 to 65)	25 (11 to 40)	} 10 (-10 to 29)
Conservatives	31 (25 to 38)	47 (35 to 58)	15 (2 to 29)	
Democrats	27 (21 to 33)	48 (37 to 59)	21 (9 to 34)	} 4 (-14 to 22)
Republicans	25 (19 to 32)	43 (32 to 53)	17 (5 to 30)	
Females	26 (21 to 31)	45 (37 to 54)	19 (9 to 30)	} -4 (-11 to 20)
Males	28 (22 to 34)	43 (33 to 53)	15 (4 to 26)	
Some College	28 (23 to 33)	47 (38 to 56)	19 (9 to 29)	} 5 (-11 to 20)
No College	26 (20 to 32)	40 (31 to 50)	15 (3 to 26)	
High Income	23 (18 to 28)	46 (36 to 56)	24 (12 to 35)	} 12 (-4 to 27)
Low Income	31 (25 to 36)	43 (34 to 51)	12 (2 to 22)	