

The Effect of International Law on Preferences and Beliefs

Michael Tomz
Stanford University
tomz@stanford.edu

March 2007

Preliminary. Comments welcome!

For financial support, I thank the National Science Foundation (CAREER grant SES-0548285), the Center for Advanced Study in the Behavioral Sciences, and the Vice Provost for Undergraduate Education at Stanford. This project would not have been possible without the outstanding research assistance of Caroline Andersen, Lauren Falcao, and Merrit Kennedy, who conducted interviews with members of the British House of Commons. I am also grateful to the staff at Knowledge Networks—especially Mike Dennis and Sergei Rodkin—for fielding the public opinion poll described in this paper.

Abstract: How does international law affect preferences and beliefs about foreign policy? I investigate this question by offering the first-ever experimental analysis of treaty commitments. The experiments, embedded in interviews with US voters and British policymakers, reveal three patterns. First, international law transforms policy preferences. Individuals 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. Second, international law shapes expectations. Many individuals, including expert policymakers, believe that signatories to international treaties will behave differently from non-signatories. Third, the effect of international law is additive, not absolute. If the material and moral case for violating international law is sufficiently strong, large numbers of voters and policymakers will advocate breaking the law and will expect foreign leaders to do the same. Thus, the experiments reported here reveal both the power and the limits of international law.

How does international law affect preferences and beliefs about foreign policy? Few questions are as central to the study and practice of global politics. According to the United Nations Treaty Series, more than 50,000 bilateral and multilateral international agreements are currently in force. Although these agreements are “legally binding,” their credibility remains a matter of academic and practical debate. Without a higher power that compels sovereign states to respect international law, it is not obvious when—if at all—treaties change the incentives and expectations of key actors in world affairs.

Largely for methodological reasons, 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.” Problems of endogeneity and measurement have made it difficult to draw reliable inferences about the effectiveness of international law.

In this paper, I advance our understanding of international law 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 treaties and others to abstain, is clearly out of the question, but survey-based experiments are both feasible and informative. By embedding experiments in interviews with voters and policymakers, it is possible to estimate the effect of international law while avoiding problems of endogeneity and measurement that have hampered previous research.

My analysis, based on a unique set of interviews with U.S. voters and British policymakers, supports three conclusions. First, international law transforms policy preferences.

Individuals 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. Second, international law shapes expectations. Observers—including elite policymakers—anticipate that the behavior of signatories to international agreements will differ systematically from the behavior of non-signatories. Third, the effect of international law is additive, not absolute. If the material and moral case for violating international law is sufficiently strong, large numbers of voters and policymakers will advocate breaking the law, and will expect foreign leaders to do the same. Thus, the data in this paper reveal both the power and the limits of international law.

1. The Limits of Existing Evidence

International relations scholars disagree about the effect of formal international agreements. Some say that treaties have little constraining 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 remain parties only insofar as the agreements serve their interests. If a treaty ever becomes inconvenient, countries can either withdraw from or break the commitment. Most treaties contain exit clauses that make withdraw legal, and even outright violations can occur without fear of punishment by a world government. Thus, to the skeptic, international agreements have little effect on policy preferences.

By implication, agreements should not affect policy expectations, either. If treaties are merely scraps of paper that can be shredded or disregarded at no cost, even countries with no desire to abide by a treaty could nonetheless opt to sign. Much like cheap talk, the act of signing a treaty would not reveal information about preferences or capabilities. Unless treaties are costly

to break—unless leaders pay a domestic or international price for legalizing their commitments and not following through—treaties should not sway expectations about the way countries are likely to behave.

Other analysts argue that formal agreements can be profoundly consequential (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. One plausible mechanism is reputation: putting an agreement in writing may increase the reputational cost of renegeing. A second mechanism concerns norms, rather than interests: perhaps citizens and policymakers feel it would be morally wrong to break formal promises to foreigners. Under either mechanism, international agreements could shape the way countries behave. By extension, such agreements would also shape expectations, since countries that are predisposed to follow the treaty would sign at a higher rate than countries pre-inclined against the treaty.

Due to limitations of existing data, it has been extremely difficult to resolve the debate. We currently do not know the conditions under which international law matters, 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 law. 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

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

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 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. No country can be observed in the treatment and control regimes simultaneously, a fact that Holland (1986) called “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 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

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 (2007) 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

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).

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 on policy preferences. Later in the paper, I use experiments to estimate the influence of international agreements on policy expectations.

2. An Experiment-Based Analysis of Preferences

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 economic interests. 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 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.”

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

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 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

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.

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.

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 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 noninformative Jeffreys prior, $\text{Beta}(.5,.5)$. When an international agreement exists, this proportion $\pi_1|\text{data}$ is distributed as $\text{Beta}(b_1+.5, N_1 - b_1 +.5)$, where the subscript 1 indicates the treatment regime, N_1 is the number of respondents that received the treatment, and b_1 is the number of people that thought it would be bad to cut trade with Burma. Without an international agreement, the proportion of naysayers is $\pi_0|\text{data}$ and distributed as $\text{beta}(b_0+.5, N_0 - b_0 +.5)$, 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.⁶

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.

⁶ 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.

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]

At the same time, the evidence in Table 2 underscores the limits of international law. When told that a trade embargo against Burma would violate international law, only 44 percent of respondents who took a side felt that the embargo remained a bad idea. The remaining 56 percent, a majority, concluded that the illegal embargo would be a good idea. Although international law has a powerful ability to sway voters, economic and moral considerations trump the law for a significant portion of the electorate.

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.

[TABLE 4 ABOUT HERE]

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.

The evidence in this section has shown that international law constrains the foreign policy preferences of voters. Do policymakers—the President, members of Congress, and the bureaucracy—feel similarly constrained? One cannot know for sure without administering a similar set of experiments to leaders in Washington, DC. Nonetheless, three factors suggest that the findings may hold not only for voters, but also for policymakers. First, in a democracy, mechanisms of electoral control and accountability increase the likelihood that the preferences of leaders will match the preferences of the electorate. Second, it is not obvious why politicians would be any less sensitive than citizens to the issue of international law. Politicians do differ from the average voter in certain respects, including education and income. Table 4 shows, however, that the impact of international law is at least as large for people with college degrees and high incomes, as for people with less education and wealth. Third, evidence from the United Kingdom (discussed in the next section) shows that British policymakers regard international law as highly consequential. It seems likely that many—though not all—US leaders would approach international law in the same way.

3. An Experiment-Based Analysis of Expectations

Having found, through experiments, that international law exerts a substantial effect on policy preferences, I collaborated with Caroline Andersen, Lauren Falcao, and Merrit Kennedy on a follow-up study about the impact of treaties on expectations. The experiment, which concerned the Nuclear Nonproliferation Treaty, was administered through face-to-face interviews with members of the British House of Commons in June-August 2006. All 646 members of the Commons were invited, by email, to participate in this study. A total of 75 agreed to be interviewed. The first eight interviews were devoted to pretesting and refining the questionnaire. The analysis below focuses on the 67 members of parliament who answered the questionnaire in its final form.

Each British MP was told: “There’s much concern these days about the spread of nuclear weapons. I’m going to describe a country that may or may not be pursuing nuclear weapons. For scientific validity, the description involves a general type of country, rather than a specific country in the news today. Some parts of the description may strike you as very important; other parts may seem much less important. When I have finished the description, I will ask—in your best judgment given the limited information available—how likely or unlikely you think it is that the country is pursuing nuclear weapons.”

Each MP then received background information, akin to an intelligence report. The treatment group, comprising roughly half the sample, was presented with the following facts:

- The country borders on an unfriendly nation that has nuclear weapons and has threatened to use them in a future war.
- The country has repeatedly said that it does not want nuclear weapons.
- The country has signed the Nuclear Non-Proliferation Treaty, thereby pledging “not to receive, manufacture, or otherwise acquire nuclear weapons.”

- Recent satellite images show that the country has started enriching uranium, which could be used for either civilian or military purposes.
- The country has a stagnant economy.

The control group received an identical set of facts, but were told that “The country has not signed the Nuclear Non-Proliferation Treaty.”

The experiment was designed not to call undue attention to the NPT. The country’s status as a signatory was included as just one of many facts, and mention of the NPT was inserted in the middle of the list of facts, rather than at the beginning or the end, where its placement might have been more salient. Moreover, participants were told that some of the facts might be unimportant, to reduce if not remove any implication that the country’s status in the NPT should be a key consideration.

After reading these points aloud, the interviewers passed the MP a card that contained the same information in written form, thereby giving the MP a chance to review the facts. The interviewers then asked: “Would you say that it is very likely that the country is pursuing nuclear weapons, that it’s somewhat likely, that there’s a 50-50 chance, that it’s somewhat unlikely, or that it’s very unlikely?” Finally, the interviewers encouraged each MP to elaborate on the reasons behind their assessment.

MP’s who answered these questions were fairly representative of the House of Commons as a whole. Table 5 shows, for example, that the party affiliations and tenures of participants in our study were very close to the House benchmarks. The sample did include a disproportionate share of males, and a larger proportion of MPs that had served on foreign affairs committees or had, in their official biographies, expressed an interest in foreign countries. Given that we had asked MPs to participate in a study about “how British leaders think about important foreign

policy issues,” it is likely that some foreign-affairs-minded leaders self-selected into the study. This kind of sample may be especially relevant for understanding the foreign policy, since those who expressed an interest in foreign affairs or had served on such committees would be likely to be most influential in the shaping of British foreign policy. Through a battery of parametric and nonparametric tests, I confirmed that treatment and control groups were balanced according to the demographic and political variables in Table 5.

[TABLE 5 ABOUT HERE]

Data from this experiment show that even the oft-maligned Nuclear Nonproliferation Treaty has a powerful effect on the expectations of policymakers (see Table 6). When told about a country that had not signed the NPT, a sizeable majority (approximately 61 percent) of British MPs deemed it either somewhat likely or very likely that the country was pursuing nuclear weapons. In contrast, only 35 percent of MPs thought it likely that a signatory was following the nuclear path. The difference between these two estimates is roughly 25 percent, implying that, the mere signing of a treaty would have changed the policy expectations of roughly one-fourth of the British parliament. Another way to express the power of international law is via the relative risk statistic, 1.8, which means that, in this experiment, the perceived likelihood of pursuing nuclear weapons was approximately 80 percent when the country had not signed the NPT than when it was party to the agreement.

Analysis of open-ended comments from British MPs supports the same conclusion. Roughly one-fourth MPs in the treatment group explained that they judged proliferation less likely *because* the country had signed the NPT. Even more impressively, 39 percent of MPs in the control group said they were swayed in their estimates by the country’s failure to sign the NPT. Apparently, countries that signed the treaty reassured British policymakers, whereas

countries that failed to sign aroused suspicion. Both patterns are consistent with the idea that treaties constrain, such that decisions to sign (or not) convey information about preferences and abilities.

As with the experiment about international trade, the experiment about nuclear proliferation demonstrates not only the power but also the limits of international law. More than one-third of British MPs deemed it likely that the country was pursuing nuclear weapons, even though it had signed the Nuclear Nonproliferation treaty. Moreover, six participants in the study explicitly noted that the treaty had no effect. One parliamentarian, for example, described the treaty as “irrelevant,” and another emphasized that signing or not signing the treaty would “not actually make any difference.” Nonetheless, the data show that, in the eyes of a significant portion of British policymakers, international agreements transform expectations about the behavior of other countries.

4. Conclusions

In at least two ways, the experiments in this paper shed have new light on the effects of international legal agreements. First, the experiments have overcome problems of endogeneity and measurement 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 have strengthened the microfoundations of international relations theory by revealing how legal agreements affect preferences and beliefs. Previous researchers have found correlations between international agreements and policy outcomes. Some contend that the relationship is spurious, whereas others argue that international agreements change cost/benefit calculations. The evidence in this paper supports the second interpretation. Individuals 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. And individuals—including expert policymakers—anticipate that signatories to treaties will exhibit systematically different behavior from non-signatories.

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.

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. The evidence in this paper complements a growing body of high-quality research that others have done with historical data, by demonstrating that international law has a large and important effect on preferences and beliefs about foreign policy.

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. 2007. "Membership Has Its Privileges: Understanding the Effects of the GATT and the WTO on World Trade." Forthcoming, *American Economic Review*.

Tomz, Michael. 2007. *Reputation and International Cooperation: Sovereign Debt Across Three Centuries*. 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 on Policy Preferences

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

	<u>Estimate</u>	<u>Confidence Interval</u>
Violates International Law	44	(38 to 51)
Does not Violate	27	(23 to 31)
Difference (Effect of Law)	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)	

Table 5: Representativeness of the Sample

The table compares the political and demographic attributes of the sample (N=67) versus the House of Commons as a whole (N=646)

	House of Commons	NPT Sample
Political Party (%)		
Labour	55	51
Conservative	31	33
Liberal Democrat	10	15
Other	5	1
Experience (years in office)		
Median	10	10
Mean	12	11
Gender (% Male)	80	87
Foreign Affairs (%)		
Stated interest	62	72
Committee service	20	30

Table 6: Effect of International Law on Policy Expectations

Table gives the percent of respondents who deemed it likely that the country was pursuing nuclear weapons. Bayesian 95% confidence intervals appear in parentheses

	Estimate	Confidence Interval
Did not Sign NPT	61	(42 to 77)
Signed the NPT	35	(21 to 52)
Difference	25	(1 to 47)
Relative Risk	1.8	(1.0 to 3.0)

Note: There were 28 valid responses in the control group (did not sign NPT) and 34 valid responses in the treatment group (signed the NPT).