International Law for an Uncertain Environment

Barbara Koremenos
Assistant Professor of Political Science
UCLA

1301 Franklin #10
Santa Monica CA 90404
310 613 6544
koremeno@polisci.ucla.edu

Abstract

International cooperation is plagued by uncertainty. While states conclude the best agreements possible using available information, things happen after agreements are signed that are beyond states' control and extremely difficult to predict. States may not commit themselves at all to an agreement if they anticipate that circumstances will alter their expected stream of benefits. I argue that flexibility provisions can insure states in this context. Whether states incorporate flexibility depends positively on the degree of uncertainty and their relative risk aversion. Working in the opposite direction is the cost of flexibility. Unlike existing work on international cooperation, this paper presents evidence on flexibility provisions using a random sample of agreements drawn from the United Nations Treaty Series (UNTS). Most important, formally-derived hypotheses are subjected to empirical test using the sample. The results strongly suggest that international agreements and institutions are consequential and their specific design provisions are, in great part, what make them stable and hence consequential.

1 Introduction

For decades, the theoretical literature on international cooperation focused on overarching questions about whether cooperation is possible and how important it is. The seminal contributions of the 1980s increased our theoretical understanding of the possibility of cooperation. Yet we know empirically that cooperation is pervasive. Hundreds of multilateral agreements are signed each year. If we count bilateral agreements as well, the number jumps to thousands. This is not to say that cooperation is easy. In fact, given the challenges of successful cooperation, it is time for the theoretical literature to focus not on whether cooperation can occur at all, but on more subtle questions regarding how the actual institutions of cooperation work and through what means they have their impact on state behavior.

Cooperative international relations are typically organized through agreements and the institutions they create. These agreements and institutions display a wide range of forms across different parties, issues, and circumstances. It is astonishing that after over fifty years of intense debate in international relations theory, much of which turns on the question of the value and function of international agreements, no one has systematically collected data on important dimensions of international agreements. This is especially puzzling given that such agreements have increased substantially since World War II. This paper is part of a research program that addresses both the theoretical and empirical gaps in the literature on international cooperation by investigating core questions of agreement and institutional design.

International agreements and institutions *are* consequential and their specific design features are, in great part, what make them *stable* and hence consequential. These features vary in systematic and important ways and are deserving of focused research. Thus, like scholars in international law, I take seriously the details provisions of international agreements and organizations. Importantly, however, I go beyond the descriptive work that characterizes much of international law to show theoretically that the careful choice of these provisions makes international cooperation both more likely and more robust.

The optimal design provisions depend on characteristics of both the environment and the parties trying to cooperate. One pervasive and important environmental factor is uncertainty. In this paper, I ask

how states structure the institutions of international cooperation to deal with ongoing uncertainty and rapid change? To capture the rapidly changing environment within which many international agreements are concluded, I assume that states face persistent uncertainty about the distribution of gains from their agreements. I examine states' choices between an indefinite (inflexible) agreement and a series of regularly renegotiated (flexible) agreements.

This article adds to a small but growing theoretical literature on how states choose the flexibility provisions of international agreements. Koremenos (2001) considers the choice of flexibility in an environment where the distribution of gains from an agreement is fixed but initially unknown, and where players learn about the distribution over time by observing their outcomes under the agreement. The model sheds light on important security agreements like the Nuclear Non-Proliferation Treaty. A small set of papers considers an environment in which a state's gain from an agreement is affected by infrequent, irregular, transitory negative shocks. The classic examples of such shocks are wars and banking crises under a gold standard. Bordo and Kydland (1995) show how states used an "escape clause" consisting of temporary suspensions of convertibility to deal with these shocks. Downs and Rocke (1995) and Rosendorff and Milner (2001) consider similar models in a trade context, where the shocks represent sudden mobilization of domestic interest groups in response to particular aspects of trade agreements such as the GATT and WTO. The escape clauses in these agreements allow states to temporarily opt out of particular aspects of the agreement without leaving the agreement as a whole.

What is missing from all of these works is an empirical evaluation based on a large-n data set. This paper fills that void by presenting evidence on flexibility provisions using a random sample of agreements drawn from the United Nations Treaty Series (UNTS). This evidence illustrates the extent of

heterogeneity in the incidence and types of flexibility provisions. Most important, formally-derived hypotheses are subjected to empirical test using the random sample.[1]

The remainder of the paper is organized as follows. Section 2 presents the theory, focusing on the type of flexibility available and on the independent variables that influence that choice. Section 3 reviews the empirical evidence on flexibility provisions in random samples of international agreements and tests the predictions of the theoretical framework. Section 4 concludes.

2 Explaining the Structure of International Agreements

International cooperation is plagued by uncertainty. While states conclude the best agreements they can using available information, things happen after their agreements are signed that are beyond their control and often extremely difficult to predict. When the stakes are large, states may be unable to commit themselves to an agreement in the first place if they anticipate that circumstances will alter the stream of benefits they were expecting to receive. But what if they could somehow insure themselves? The right scheme could make possible many commitments that would otherwise have been precluded.

I focus on a particular kind of uncertainty -- that surrounding the distribution of future gains from an agreement. The parties select an initial distribution of gains based on their relative bargaining power, but this gain then evolves over time under the agreement due to external shocks. The precise nature of the shocks depends on the issue area but could include climactic fluctuations (as in the case of commodity agreements), exogenous private sector shocks (as in the case of macroeconomic coordination or exchange rates), or political shocks (as in the case of refugee agreements). I assume that the parties to an agreement always know the distribution of gains in the current period, but know only the probability distribution for the distributions of gains in all future periods. The distribution of gains in each period

[1] While this is perhaps the first study to use a random sample and to examine systematically issues of agreement design, Simmons (2000), Guzman and Simmons (2002), and Martin (1992; 2000) are fine examples of careful large-n empirical work in the field of international cooperation.

equals the distribution of gains in the previous period plus a random shock. These random shocks are cumulative so that the distribution of gains, left to its own devices, follows a random walk.

Put simply, a shock to the distribution of gains occurs in each period. As a result, either the gains to one state increase relative to that of the other or, possibly, there is no change, if the shock equals zero. These shocks have two important features. First, they are completely random. Each shock is independent of every other, both before and after. They are also independent of everything else that the states can observe, which implies that they cannot be predicted, either at the time an initial agreement is made or later on. Second, the shocks add up over time. That is, the gain for a state in any one period is the sum of its initial gain plus each of the subsequent changes. Because the random shocks are cumulative (rather than, say, simply affecting the distribution of gains in a given period and not thereafter), it is possible for the states' relative shares to evolve in a way that differs quite substantially from the assumed equal division at the start of the agreement. How much different depends both on how much time has passed and on the specific sequence of realized shocks. The more time that has passed, the larger the potential divergence from the initial equal division because the passage of time allows, but does not require, the random walk process to walk farther away from its starting point. Hence, once a policy course is chosen, without subsequent redirection, the system may continue down particular paths longer than the parties originally intended. I refer to this environment as one characterized by "persistent uncertainty."

The extent or importance of the *uncertainty* about the distribution of gains is summarized by the variance of the shocks. The larger the variance, the greater and more important is the uncertainty. Moreover, the larger the variance, the farther the distribution of gains will depart from its initial value, on average, over a given time.

States are averse to such uncertainty. The idea that they may gain far less than they anticipated from an agreement could make some of them too nervous to commit to one. Formally, I employ a standard assumption in international relations that states are *risk averse*. I allow them to vary with respect to their relative level of risk aversion.

Types of Agreements

Suppose states determine that there are indeed gains from cooperation and decide to negotiate an agreement. They have two options to choose from in a given agreement context:

(1) An agreement of indefinite duration (an inflexible agreement).

(2) A series of agreements renegotiated at regular intervals in order to adjust the distribution of gains for the effects of the shocks that cumulate during each agreement (a flexible agreement).

It is important to note two options not available. The first is a completely contingent agreement. I exclude this option both because the theoretical models in economics discussed below exclude it and because we observe empirically that no agreements attempt to condition on more than a very few easily measurable and verifiable contingencies.[2] The second is to create and then delegate "cooperation" to an international organization (IO). This second option is found in Koremenos (2000), which models explicitly the creation of an IO as a choice variable. It is left out here because data limitations do not allow for testing the fuller model. In random samples, there is a predominance of bilateral agreements; ongoing data collection will allow testing the fuller model in the future.

Consider first an agreement of indefinite duration. States pay initial negotiation costs to set up the agreement, but no additional costs after that. The result of this cost saving is that there is no provision for renegotiation of the distribution of gains in response to shocks, which implies that the distribution may wander well away from its initial level.

Given there is no international authority to enforce agreements, should the distribution of gains move sufficiently far enough away from the initial level, one state may break the agreement. The state that elects to do this pays a cost. The costs of reneging consist of any sanctions imposed by the other parties to the agreement as well as reputational costs. These costs may be substantial. I assume that once reneging costs are paid, the states begin again with a new agreement of nominally indefinite duration. Put

[2] See the discussion in Hart and Holmstrom 1987 for a justification of this assumption in the context of the economics literature on incomplete contracts.

differently, I assume that the reneging state incurs no cost in terms of reputation or refusal to bargain, thereby avoiding problems of renegotiation-proofness.

Instead of one indefinite duration agreement, states can choose to conclude a series of finite duration agreements of equal duration, with renegotiation taking place between each pair of agreements in the series. I consider only series of agreements of the same duration because of my assumption that the environment facing the parties does not change over time. As a result of this assumption, the states will make the same choice regarding agreement duration at the start of the first agreement and every subsequent agreement because the choice problem they face in each case does not change.

States choosing to conclude a series of finite duration agreements incur two kinds of costs. First, the states incur the negotiation costs required to reach the initial agreement in the series. I assume that these are the same as in the case of a single agreement of indefinite duration, although they could be less due to lower bargaining costs resulting from a shorter shadow of the future. Second, states incur *renegotiation costs* at the start of each new agreement in the series. Like negotiation costs, renegotiation costs will vary depending on the states in the agreement and the environment. For example, I would expect both costs to increase with the number of parties to the agreement.[3]

In some contexts, renegotiation costs may differ substantially from the costs incurred to negotiate the initial agreement in the series. For example, renegotiation may consist only of adjusting a few provisions of an existing agreement while leaving the bulk of the agreement framework unchanged. In such cases, renegotiation costs are lower than the original negotiation costs. In other contexts, such as the annual G-7 macroeconomic accords, negotiation and renegotiation costs may be roughly equal, due to the rapidly changing agenda across years and the lack of formal structure to the agreements.

[3] Conybeare (1986) makes this argument in relation to trade talks, noting that the "Kennedy Round of multilateral negotiations from 1963 to 1967 took considerably longer than the previous Dillon Round of bilateral negotiations in 1960-61."

The advantage that states derive from concluding a series of renegotiated agreements rather than one agreement of indefinite duration is flexibility. With renegotiation, the division of gains can be reset to the initially desired level at regular intervals. States choose to reset the division of gains to its original level because I assume that bargaining power and other factors that influence their choice of distribution do not change over time.[4]

The planned readjustment to the division of gains that occurs under a series of renegotiated agreements greatly reduces the chance that either state will want to incur the costs of reneging or be forced to endure an unsatisfactory division of gains for long periods. Put differently, for risk-averse states, the opportunity periodically to reset the distribution of gains back to its initial level increases the ex ante value of a series of renegotiated agreements relative to a single indefinite agreement. It does so by reducing the variance of the discounted expected gain from the agreement. States that choose to renegotiate must decide how often to do so. Doing so involves trading off the costs of more frequent renegotiation against the costs of living with a distribution of gains that differs from that initially chosen by the parties in light of their relative bargaining power and other factors.

Foreshadowing the formal results below, states will choose to renegotiate more often (have shorter duration agreements) when *renegotiation costs* fall. Similarly, states renegotiate more often when the *variance of the shocks* to the increases because this increases the expected deviation between the realized and desired distributions of gains and when their level of *risk aversion* increases because the variance is more costly to them then.

Suppose that the variance of the shocks to the distribution of gains is very low so the benefits from flexibility are small and the costs of flexibility (renegotiation costs) are high. In this case, states may choose an indefinite agreement. The 1967 Outer Space Treaty, which forbids the placement of

[4] This implies in particular that they do not change as a result of the shocks to the distribution of gains under the agreement. While this assumption could be relaxed without changing the main results, doing so would add a great deal of additional complexity to the model.

weapons of mass destruction into orbit, onto celestial bodies, or in outer space, is not subject to great uncertainty of this kind. There are no supply and demand shocks, and no resource or territorial issues are at stake. Rather the goal is to lock in the status quo and prevent the militarization of space. The states involved therefore chose not to incur renegotiation costs; the agreement is of indefinite duration.

Now suppose that the variance of shocks to the distribution of gains is large so that it drifts quickly away from its preferred level. In this case, states will want to incorporate flexibility into their agreement. The annual G-7 macroeconomic accords are an example of a series of renegotiated agreements. In this case, negotiation and renegotiation costs are very low given that most of what is agreed is not legally binding nor formally ratified and given the nature of the parties involved: a small group of states with relatively similar beliefs and goals. These facts, combined with a high variance of shocks to the distribution of gains due to a rapidly changing world economy, led the G-7 to choose a series of short agreements. In the case of the G-5 finance ministers, costs are lower, the variance of the shocks is higher, and agreements are even shorter.

Formal Model

I now consider a formal model of agreement type choice.[5] I consider two states, State 1 and State 2.[6] Let

$$Y_{1,t} = b_{1,t}, \; Y_{2,t} = b_{2,t},$$

[5] A few economists have asked how the parties to a series of finite duration contracts choose the optimal duration of those contracts in the presence of uncertainty similar to the type I analyze here. The relevant papers in economics are Gray (1978), Dye (1985), and Harris and Holmstrom (1987). Although these three papers focus on different contexts (labor contracts and loan contracts) the basic setup is similar. In each case, there is some underlying random variable that affects the payoffs to the parties from the contract. All assume the parties cannot make the contract contingent on this variable.

[6] The UNTS data at the time my sample was drawn consisted of 32,939 bilateral agreements and only 2330 multilateral ones.

denote the outcomes for the two states in the absence of an agreement. The first subscript indicates the state and the second indicates the time period in each case. These outcomes depend on the specific issue area under consideration, but may include things like GDP, trade levels, or exchange rate stability, appropriately measured. For simplicity, I assume that the base outcomes do not change over time.

In the presence of an agreement, the outcomes attained by the two parties to an agreement in each period t are given by

$$Y_{1,t} = b_{1,t} + m_t, \;\; Y_{2,t} = b_{1,t} + (g - m_t),$$

where g is the total gain from the agreement (the size of the pie) and m_t is the portion received by State 1.

For simplicity, I assume a fixed total gain g. Rather than having the total gain vary over time, I only have the division of gains between the parties vary over time. While this sacrifices some realism, it simplifies the model and allows me to focus on the key issue of how states structure agreements in the face of distributional uncertainty.

To take account of the rapidly changing context of international agreements, I assume that in the absence of renegotiation, the distribution of gains evolves according to a random walk, with

$$m_t = m_{t-1} + e_t,$$

where m_0 is chosen under the agreement and where e_t has density function $h(e_t)$. I assume, without loss of generality, that the states are symmetric. This implies equal bargaining power and, under a Nash bargain solution, equal initial shares so that $m_0 = 0.5g$.[7]

I assume that $E(e_t) = 0$ and that e_t is independently and identically distributed across periods. The parties observe m_t in each period. Let $t = 0$ denote the first period of an agreement. By choosing m_0, the parties to the agreement are choosing the expected value, as of $t = 0$, of m_t for all future periods. Because the e_t are independent across periods and have a common variance, we have

$$m_t = m_0 + e_1 + e_2 + \ldots + e_t,$$

[7] This cooperative solution corresponds to the Rubinstein alternating offers non-cooperative solution when ? is close to one. See Osborne and Rubenstein (1990).

and

$$var(m_t) = var(g - m_t) = t * var(e_t).$$

That is, because the shocks are independent and cumulative, the set of possible distributions of gains from the agreement fans out over time as the agreement continues. Put differently, the probability that the realized value of the distribution of gains differs by any fixed amount from the initial choice of m_0 increases over time. In the general case considered here, the gain for one party (either m or $(g - m)$) from the agreement may become negative (though an agreement whose initial gain was negative for one or both parties would, obviously, not elicit much interest).

I denote the cost of negotiating either a single, infinite duration agreement or the first in a series of renegotiated agreements by k_n. When an agreement gets renegotiated, states pay a renegotiation cost k_r. These costs are incurred by all of the parties to an agreement. I allow the costs of the negotiation and subsequent renegotiations to differ because we observe in the real world that they often do.

In general, I expect that $k_n > k_r$ because renegotiation often simply changes parameters within an agreement framework, while the original negotiation involves the more difficult task of establishing the framework. I interpret both k_n and k_r much more broadly than the economics literature. In addition to the costs of sitting down and talking, I think of both as including the costs of assembling or reassembling domestic political coalitions – often necessary for agreement ratification. In addition, renegotiation costs may be affected by changes in membership in multilateral agreements.

As noted above, when states renegotiate an agreement, they reset the distribution of gains under the agreement to that originally agreed upon. This follows from my assumption that the bargaining power of the states in the agreement remains unchanged over time (and in particular is not affected by the realized distribution of gains under the agreement). In my notation, this means that $m_t = m_0$ in periods immediately following a renegotiation.

To focus on the basic tradeoffs, consider a simple two-period case. The parties must choose between one two-period agreement (an inflexible agreement) and two one-period agreements (a flexible agreement in that it can be renegotiated). If they choose two one-period agreements rather than one two-

period agreement, they each gain a reduction in the variance of their gain from $2 * var(e_t)$ to $var(e_t)$, but they lose the additional renegotiation cost. The expected value (gross of negotiation and renegotiation costs) is the same in both cases, which implies that risk neutral parties would opt for the two-period agreement, as variance does not matter to them.

Formally, in the two-period case, the parties compare their expected utilities with one two-period agreement, given by

$$E(u(b_1 + m_0 + e_1 - k_n) + \delta u(b_1 + m_0 + e_1 + e_2)),$$

$$E(u(b_2 + (g - (m_0 + e_1)) - k_n) + \delta u(b_2 + (g - (m_0 + e_1 + e_2)))),$$[8]

to their expected utilities from two one-period agreements, given by

$$E(u(b_1 + m_0 + e_1 - k_n) + \delta u(b_1 + m_0 + e_2 - k_r)),$$

$$E(u(b_2 + (g - (m_0 + e_1)) - k_n) + \delta u(b_2 + (g - (m_0 + e_2)) - k_r)),$$

The parties choose the agreement type that provides the highest expected utility.

Comparative Statics

Below I present the comparative static predictions of my model. In all cases the hypotheses implicitly hold everything else equal so they indicate the effect of changing one parameter of the model holding the others constant. The proofs are in Appendix A. These predictions represent testable hypotheses about the effects of variation in states' characteristics and in the agreement context on the type of agreement selected. I also present some hypotheses regarding the choice of agreement duration, should states choose to conclude a series of finite duration agreements. I have simulated the model with the

[8] This formulation implicitly assumes neither party reneges in the two-period case. I assume the possible values of the agreement shock are such that, in only two periods, the distribution of gains never moves far enough away from the initial value to make paying the cost of reneging worthwhile. This is a reasonable assumption except for (hypothetical) states on the extreme margin that completely discount future benefits. Still, the complete case with reneging is described in Appendix B.

possibility of reneging using a variety of alternative values for the key parameters: the level of risk aversion, renegotiation costs, and the variance of the shocks to the distribution of gains. The simulations support the comparative statics (see Appendix B).

(CS-1) As **renegotiation costs** increase, the probability that the parties will choose finite, renegotiated agreements decreases.

(CS-2) If the parties choose to conclude a series of renegotiated agreements, then as **renegotiation costs** increase, they will choose to make each agreement in the series longer.

An increase in renegotiation costs (CS-1) raises the costs of choosing a series of renegotiated agreements relative to an indefinite duration. With respect to (CS-2), the costs of renegotiation increase while the benefits of renegotiation remain unchanged; hence the parties will renegotiate less often by making each agreement in the series longer.

(CS-3) As the **variance of the shocks to the distribution of gains (uncertainty)** increases, the probability that the parties will choose finite, renegotiable agreements to adjust for shocks increases.

(CS-4) If the parties choose to conclude a series of renegotiated agreements, then as **the variance of the shocks to the distribution of gains (uncertainty)** increases, the parties will choose to make each agreement in the series shorter.

An increase in the variance of the shocks (CS-3) makes the parties value flexibility more. Put differently, it increases the variation in realized outcomes under an indefinite duration agreement, which makes it less attractive relative to the alternative. With respect to (CS-4), the value to the parties of renegotiating more often to undo the shocks increases; hence, they conclude shorter agreements.

(CS-5) As the **risk aversion** of the parties increases, the probability that they will choose finite, renegotiable agreements to adjust for shocks increases.

(CS-6) If the parties choose to conclude a series of renegotiated agreements, then as the **risk aversion** of the parties increases, they will choose to make each agreement in the series shorter.

In (CS-5), an increase in the risk aversion of the parties increases the value to them of having some form of flexibility in their agreement to reduce the variation of the realized outcomes. In the case of (CS-6), the intuition is that as the level of risk aversion increases, so do the costs of putting up with an agreement whose distribution of gains has moved far from that originally agreed upon. As I hold the costs of renegotiation constant, states will choose in this case to renegotiate more frequently so that the variance in the realized outcomes falls.

Subject to two caveats, the two-period model generalizes very readily to an infinite horizon framework, particularly under the assumption of a time homogeneous environment. States face the same two choices and again compare their discounted expected utility from each one, choosing that with the highest value.[9]

The first caveat to this generalization concerns the alternative of a series of finite duration agreements. In the two-period model, these agreements can only be of length one. In an infinite horizon framework, the states must choose their preferred duration from the set of all possible finite durations, by comparing the discounted expected utilities associated with series of finite duration agreements of various lengths. Once they have identified the finite duration agreement length they like best, they can then compare it to their discounted expected utility from the other three agreement type choices.

The second caveat concerns the possibility of reneging. Over an infinite time horizon, the distribution of gains under an indefinite agreement could move very far away from the parties' preferred division. If it moved far enough, one party might be losing more from the agreement than the sum of the expected future gains and the costs of reneging. The same condition could hold in the midst of a relatively long finite duration agreement. Thus, reneging looms larger when we extend the time horizon.

[9] Generalizing the model for an environment that changes over time would be conceptually simple but notationally burdensome. It would not add anything to the substantive results.

The generalization to additional states is straightforward, subject to the following caveats. First, adding additional states will have effects on certain parameters of the model. As a result, holding the other parameters constant, agreement type choice, as well as the choice of whether to have an agreement at all, may change with an increase in the number of parties. For example, increasing the number of states will increase renegotiation costs. Moreover, generalizing the model both across time and to multiple parties calls attention to the relationship between potential changes in membership and the anticipated costs of renegotiation. If the agreement is such that membership is likely to dramatically increase, then states may not choose renegotiation even though the initial membership is quite low. Given that membership rules themselves also greatly affect the number of future signatories to an agreement, this also raises the issue of the interaction between two aspects of institutional design: flexibility *and* membership rules.

3 Empirical Results

Data

According to Article 102 of the Charter of the United Nations, "every treaty and every international agreement entered into by any Member of the United Nations after the present charter comes into force shall as soon as possible be registered with the Secretariat and published by it" (http://www.un.org/Overview/Charter/contents.html). Given the almost universal membership of the United Nations and its stature among international organizations, its list of international agreements is the most comprehensive to be found. All international agreements registered or filed and recorded with the Secretariat since 1946 are published in the United Nations Treaty Series (UNTS). The Internet collection at the time the sample was drawn contained over 34,000 international agreements "which have been published in hard copy in over 1,450 volumes, which corresponds to all treaties and subsequent actions registered up to December 1986" (http://www.un.org/Depts/Treaty/). The key exception to this completeness is the absence of informal agreements from the UNTS. By implication, my data set contains only formal agreements.

The UNTS Internet site provides a range of subject terms that can be used in searching for international agreements. The primary data for this paper are drawn from the monetary matters, finance, investment, agricultural commodities, environment, human rights, security, and disarmament subject headings. Within each issue area, I generate a random sample of the total set of agreements by printing out a list of the agreements, numbering them, and then using a random number generator to select the agreements to be coded. Once an agreement is selected, it is checked against a series of screening criteria for exclusion from the data set. I exclude agreements that only establish procedures for other agreements or designate the host for international conferences, that do not include at least two states among the parties, and that do not prescribe, proscribe, or authorize behavior that is observable at least in principle. I also have rules to avoid double counting. (TABLES S1-S7 LIST THE AGREEMENTS INCLUDED IN THE SAMPLE. THESE ARE FOR THE REVIEWERS. THEY WILL BE ON A WEBSITE, NOT IN THE PAPER.)

The characteristics of the selected agreements not excluded are recorded using a special coding form. Detailed information on various agreement flexibility provisions such as duration and renegotiation provisions, withdrawal and escape clauses, and the presence of quasi-legislative institutions designed to adapt to shocks as well as a number of other design features like membership and forms of delegation is coded. .The sample includes 74 economic agreements, and 25 each in the environment, human rights, and security issue areas. Within the economics issue area, there are 19 agricultural commodity agreements, 19 finance agreements, 25 investment agreements, and 9 monetary agreements.

Variables

The dependent variables, whether an agreement is finite or not and, if so, how long it is, are measured directly by coding the agreements. Still, it is necessary to operationalize the following independent variables: renegotiation costs, uncertainty, and risk aversion. One reason there has been negligible scientific testing of any formal models of international cooperation is the difficulty of finding measures for commonly-used variables, like uncertainty and risk attitude. Still, we must rise to the challenge or some of the claims made by critics of the formal approach will be substantiated (Brown

2000). Although I qualify my measures somewhat below, they represent earnest attempts to capture the theoretical concepts and thereby serve as useful starting points for this much-needed work in the subfield.

I use the number of original signatories to the agreement as a proxy for *renegotiation costs* since, according to bargaining theory, increasing the number of states involved is likely to make the process of reaching a consensus on the negotiation of a new agreement more lengthy given multiple equilibria.

For *uncertainty*, the coding rules are based on categories of agreements in each issue and sub-issue area. In the model, the uncertainty surrounds the variance of shocks to the distributions of gains from the agreement. Certain kinds of agreements are far more subject to these shocks than are others. For example, agreements whose distribution of gains is governed by the forces of supply and demand are coded as high uncertainty whereas those primarily about coordinating policies to avoid suboptimal outcomes are coded as low uncertainty. Put differently, agreements for which changes in the environment can cause the distribution of gains to vary substantially over time, even while being welfare-enhancing in the aggregate, are high uncertainty agreements whereas those for which efficiency concerns dominate and period-to-period changes in the distribution of gains are not expected are low uncertainty agreements.

Beginning with the economics issue area, the sub-issue area of *monetary* agreements essentially contains exchange rate agreements. These are subject to supply and demand shocks that could dramatically alter the distribution of gains from period to period. They are all coded as having high uncertainty. All *trade* agreements are coded as high uncertainty observations for the same reason. *Finance* agreements are of two types: those that resemble monetary or trade agreements – high uncertainty – and those that are about coordinating policies (e.g., Convention for the avoidance of double taxation and the prevention of fiscal evasion with respect to taxes on income between Australia and Italy) – low uncertainty. Almost one third of finance agreements are high uncertainty. *Investment* agreements concern the promotion of and protection of investments against nationalization and expropriation and

thereby are subject to political shocks that could alter the distribution of gains; hence they are high uncertainty agreements.[10]

For the *environmental* issue area, the following rules were followed: Those agreements addressing plant and bird protection or scientific cooperation on subjects like insect control or soil are coded as low uncertainty agreements (e.g., Exchange of Notes Constituting an Agreement on the Project Soil Management and Conservation in East Amazonia between Brazil and the Federal Republic of Germany). Such agreements are predominantly about coordinating policies and address topics with relatively low economic or political stakes. On the other hand, agreements about pollution abatement or about fishing or other sea resources are coded as high uncertainty agreements (e.g., the multilateral International Convention for the Conservation of Atlantic Tunas). Distributional concerns are weighty in such agreements. Restraint is costly in both sub-issues, while dependence on the resource varies (fish) as does technological development (e.g., in pollution control, expected positive developments may not be forthcoming); hence shocks can have nontrivial distributional consequences. Just over half of the environmental agreements are coded with high uncertainty.

For *human rights*, universal declarations (e.g., the multilateral Convention on the prevention and punishment of the crime of genocide) are coded as low uncertainty agreements. Such agreements serve to capture or establish ethically-based international norms. On the other hand, agreements about refugees and detailed labor standards are subject to distributional shocks. For instance, political shocks may dramatically change the flow of refugees and thereby change the distribution of gains. Just over half of human rights agreements are coded with high uncertainty.

Security agreements fall into two basic categories: universal prohibitions (e.g., the multilateral Agreement governing the activities of states on the moon and other celestial bodies) and those related to

[10] If I were to restrict my interpretation of uncertainty and not include political shocks, only one *investment* agreement is high uncertainty. Still, none of the test results are substantively altered; only the significance of a few variables changes with some improving and others becoming less significant.

mutual security. For the same reasons I argue above in human rights, the prohibitions are low uncertainty while the mutual security agreements are high uncertainty. Almost two-thirds of security agreements are characterized by high uncertainty.

While my measure of uncertainty is far from perfect, looking at sub-issue areas is a reasonable operationalization for two reasons: First, it captures much of the uncertainty modeled in this paper. Second, subcategories are treated systematically based on conversations with specialized scholars, thereby avoiding any agreement-by-agreement subjectivity.

Although *risk aversion* has been a standard assumption in theoretical models of international relations, i.e., anarchy makes states risk averse, very little attention has been paid to developing measures and incorporating it into analyses of international cooperation. I used three different proxies, all of which capture at least some aspect of relative risk aversion.

I first turn to Bueno de Mesquita's (1985) risk attitude measure, which has become a "*de facto* academic standard" (Bennett and Stam, 2000a, 541) in the international conflict literature. The measure is based on the premise that the closer a state is to the alliance portfolio that maximizes its security, the more risk-averse it is (Bennett and Stam 2000b, 464).[11] To generate risk attitude scores, I used the *EUGene* risk attitude variable (Bennett and Stam 2000c).[12] Risk scores are region based, ranging from –1 (very

[11] More specifically, the variable is constructed in three steps: (1) define state *i*'s "security level" (the sum of all other states' expected utilities vs. *i*); (2) identify the hypothetical alliance portfolio that would minimize and the portfolio that would maximize *i*'s security; (3) how proximate *i*'s **actual** policies are to its **hypothetical** policies may be interpreted as an indication of *i*'s willingness to take risks. Bueno de Mesquita assumes that "*i*'s risk acceptance increases as *i*'s security score approaches its level of greatest vulnerability, and that *i*'s risk aversion increases as its security approaches the level possessed by its safest policy portfolio (Bueno de Mesquita, 1985, 157).

[12] According to Bennett and Stam (2000a, 466), *EUGene* uses an improved and more accurate algorithm to generate risk attitude scores than does Bueno de Mesquita (1985). I employed *S* rather than Tau-b as a

risk-averse) to +1 (very risk acceptant). For bilateral treaties, I use regional risk scores. For example, if states *i* and *j* sign a treaty, I calculate *i*'s risk attitude toward state *j*'s region, and *j*'s risk attitude toward *i*'s region. For multilateral treaties, which include up to 119 signatories, I compute global risk scores (the mean of each signatory's regional risk scores) for each signatory. For each treaty, I assume the resultant agreement to be a function of the risk attitude of the *most* risk-averse participant.[13] Finally, for ease of interpretation, I convert the risk attitude score into a risk aversion measure by inverting the scale. Risk aversion therefore runs from –1 (least risk-averse) to +1 (most risk-averse).

Bueno de Mesquita's risk attitude variable is not without criticism. One problem is that it assumes that alliance is equal to security (and is therefore related to risk aversion); it fails to acknowledge the possibility that an alliance *can* enable autonomy (and therefore may indicate that a state is less, rather than more, risk-averse).[14] The measure could be improved if data on national issue positions and the status quo on those positions were accessible (Morrow 1987, 436). Unfortunately, since these data are not available, we cannot tell whether the failure to distinguish between the security benefits and the autonomy benefits of alliances introduces error into Bueno de Mesquita's risk attitude measure.

Despite this shortcoming, Bueno de Mesquita's risk attitude is far from subjective and appears to be a fairly good measure of the concept I wish to operationalize – i.e., states would be relatively more risk averse in international cooperation with those states they fear the most from a security standpoint.

---

measure of alliance portfolio similarity for the risk attitude scores. Signorino and Ritter (1999) provide good reason to believe that Tau-b can seriously misrepresent the degree to which two states' alliance portfolios are similar; their measure *S* appears to measure better the similarity of foreign policy portfolios.

[13] This "weak-link assumption" is common in quantitative research on the causes of international conflict. See, for example, Dixon 1994; and Oneal and Russett 1997.

[14] Hence, as Morrow points out, according to Bueno de Mesquita's calculations, Hitler's Germany appears slightly risk-averse. The risk attitude "assigns to German security the autonomy benefits that Germany derived from its alliances with other revisionist powers" (Morrow 1987, 436).

To create a second risk attitude variable, I use Gartzke and Jo's *Affinity of Nations Index*, which measures preference similarity among states. Preference similarity does not measure directly the level of risk aversion. Still, it is reasonable to assume that states would be relatively more risk averse to unanticipated changes in the distribution of gains when dealing with partners with very different preferences; in contrast, if a partner with similar preferences gained more, it can be more safely assumed that the gains will be used for shared priorities. That is, it is likely that *state A* will be increasingly concerned about the distribution of gains with *state B* the more divergent their preferences are.

The Affinity indicator reflects the similarity of state preferences based on their voting positions in the United Nations General Assembly (Gartzke and Jo, 2002) and is calculated using Signorino and Ritter's "S" procedure (Signorino and Ritter 1999). Possible values range from –1 (least similar interests) to 1 (most similar interests). For ease of interpretation, I invert this variable so that –1 indicates the least risk-averse state and +1 indicates the most risk-averse state. I use Gartzke and Jo's interpolated SUN2CATI variable, which indicates whether a state voted "yes" or "no."[15] Because the Affinity data are dyadic, I simply take the dyadic Affinity value for each bilateral agreement. For the multilateral agreements, the procedure is not as straightforward. I first create a dyad for each pair of signatories. Hence, if there are three signatories, there are three dyads; if there are four signatories, there are six dyads, and so on. For each multilateral agreement, I use the "weakest link assumption," taking the Affinity value of the dyad with the least similar interests.

The Affinity data's goal of measuring the similarity of state preferences is far from unproblematic, as Gartke and Jo (2002, 1-2) acknowledge. The key difficulty is that preferences are not directly observable, and hence we must assume that a state's behavior reveals its preferences. The Affinity measure does have two considerable advantages over Bueno de Mesquita's risk attitude measure. First, it is based on an information source that is less distorted than are alliance portfolios. Alliances are

[15] I also performed a probit analysis using SUN3CATI as a proxy for risk aversion. The results are similar to those obtained with the SUN2CATI variable, though slightly less statistically significant (p = .129).

costly; hence, states that do not have strong motives or threats may not pursue one, even though their preference may be to have an alliance (Gartke and Jo, 2). Second, there is much more variation in UN voting behavior than in alliance formation (particularly during the Cold War).

I now turn to a third proxy for risk-aversion: a signatory's GDP growth. I argue that states are least risk-averse when they have either very low or very high growth levels, and that they are most risk-averse when growth levels lie in a middle-range. The relationship between GDP growth and risk aversion is therefore believed to follow an inverted-U shape. The motivation for the U-shaped function stems from the literature on electoral incentives and risk attitude, in particular, those who argue that diversionary foreign policy will be seen in times of trouble (Downs & Rocke 1993; Levy 1989; Richards, Wilson & Schwebach 1993; Smith 1996, 1998). The argument is that leaders with domestic problems who anticipate being removed often undertake adventurous foreign policies that they would not have otherwise attempted. It is also widely argued that economic variables, like growth rates, are good predictors of election results. Thus leaders facing either very low or very high growth rates are likely to be relatively more risk acceptant (they are either gambling for resurrection or extremely secure) than those with middle levels; the latter do not wish to rock the boat.

Using GDP growth as a proxy for risk aversion is particularly appealing because, unlike the two previous proxies, it is a characteristic of the individual state (as in economics where risk attitude is conceived as part of an actor's utility function); a country's growth rate is the same regardless of whom the other parties to the agreement are. To operationalize this variable, I used the Penn World Table, and calculated each signatory's change in real gross domestic income (adjusting for terms of trade changes) from the year prior to signing. I again used the weakest link assumption, taking the value of the signatory with the lowest growth rate. Next, I separated the growth rate variable into three equal categories, based on the variable's distribution. States with growth levels from the 0 to 33rd percentile of the distribution were considered low-growth states, states with growth levels from the 33rd to the 66th percentile of the distribution were considered mid-growth states, and finally states with growth levels from the 66th to the

100th percentile of the distribution were considered high-growth states. I then created a dummy variable to designate risk-averse states. This variable equals one for mid-growth states and zero otherwise.

Though my three risk indicators differ in numerous ways, an iterated principal component factor analysis reveals moderately strong factor loadings, ranging from about .27 to .44.[16] While this by no means indicates that they are substitutes, these factor loadings strongly suggest that the three indicators are significantly related to one another.

Patterns of Flexibility in International Agreements

Perhaps the most striking feature of Table 1 (top panel), which presents the intended duration of my random sample of agreements, is that there is in fact a lot of flexibility: about two thirds of the agreements have a finite duration. As the bottom panel of Table 1 illustrates, the average duration for a finite agreement is about ten years.

TABLE 1 AROUND HERE

In the top panel of Table 1, the p-value of 0.000 indicates that whether an agreement is finite or not depends strongly on issue area. For example, agreements in the economics issue area are almost twice as likely human rights agreements to be of finite duration. That this kind of flexibility varies in important ways across the broad issue areas suggests that states rationally design their agreements to match the broad features of different issue areas. Yet, given that they also vary in important ways within issue areas, suggests states also tailor their agreements to meet their individual preferences and to reflect the unique aspects of particular agreement contexts.

The bottom panel of Table 1 displays the mean and standard deviation of the intended durations of the finite duration agreements.[17] Again, we see differences across issue area, although they are not as significant as those above. Of course, the sample size is also smaller given that I am conditioning on an agreement being finite.

[16] The Bueno de Mesquita and Gartzke and Jo measures are correlated at .121; the Bueno de Mesquita and growth measures are correlated at -.072; and the Gartzke and Jo and growth measures are correlated at -.144.
[17] I exclude those finite duration agreements whose duration is contingent (N = 36).

Subjecting the Model to Empirical Testing

Using the absence or presence of a finite (flexible) agreement as the dependent variable, I conduct three probit analyses. In all specifications, I include the variable indicating the number of participants (logged) and the uncertainty measure, as well as the dummies designating the human rights, economics, and environmental issue areas as discussed above. I test three separate models – one for each measure of risk aversion. Table 2 displays the results of these analyses, and Table 3 shows predicted probabilities.[18]

The empirical analyses provide strong support for the comparative statics in section 2.8. There is robust evidence that as renegotiation costs increase, states become less likely to choose finite agreements **(CS-1)**.[19] Indeed, in all three models, an increase in the number of participants (logged) significantly

[18] Scrutiny of the agreements in my sample suggests that certain states that conclude a large number of agreements, such as the United States, may follow a "template" when concluding similar agreements with different states. To the extent that there is such a template, and to the extent that multiple agreements based on the template appear in my sample, it can be argued that I should not be treating my agreements as entirely independent observations. To deal with this potential problem, I ran all three models using robust standard errors. The results do not differ significantly from those displayed in Table 2.

[19] I also operationalize the renegotiation cost variable in an additional way. Some agreements contain language requiring official ratification by the governments of the participating states; others do not. If we were to assume ratification adds to negotiation and renegotiation costs, then agreements requiring ratification are more likely to be of indefinite duration and less likely to take the form of a series of renegotiated finite duration agreements **(CS-1)**. I find that not having the requirement of ratification increases the probability of a finite agreement by about 0.32, with the underlying coefficient being strongly statistically significant. While this test lends support to my hypothesis, the measure may be problematic. If decisions about which kinds of agreements need ratification and which do not are dictated by customary international law, then it may be plausible to treat these requirements as exogenous to any given agreement. On the contrary, if states make ratification decisions at the same time they choose

decreases the probability of a finite agreement ($p < .001$). All else equal, agreements with only two participants are between 48.8% and 74.4% more likely (depending on the operationalization of risk aversion employed) to be finite than are agreements with the highest number of participants in the sample (119).[20] The empirical results also provide support for the argument that uncertainty increases states' propensity to choose a flexible agreement (**CS-3**). Indeed, in all three model specifications the uncertainty variable was in the expected direction and highly statistically significant ($p < .001$). All else equal, agreements with high levels of uncertainty are between 38.5% and 39.6% more likely (depending on the operationalization of risk aversion employed) to be finite than are agreements with low levels of uncertainty.

All three operationalizations provide evidence that risk aversion increases a state's propensity to sign a finite agreement, yielding support for **CS-5.** The Bueno de Mesquita risk variable is statistically significant at conventional levels ($p < .05$). According to this model, the most risk-averse participants are 43.9% more likely to create finite agreements than are the least risk-averse participants, *ceteris paribus*. The signs of the Affinity risk variable and the growth risk variable are consistent with the theory and their magnitude suggest important effects. However, perhaps due to the small sample size, they do not attain statistical significance at conventional levels ($p = 0.21$ and $p = 0.12$, respectively).[21] While none of the

---

flexibility provisions, then the test proposed here is rendered meaningless by endogeneity problems. A further complication arises from the fact that ratification may have different meanings from an international law point of view and from a domestic point of view. It may also mean different things to different states. Still, this measure does capture some aspect of renegotiation costs with positive results.

[20] Unless otherwise stated, I used *Clarify* (Tomz, Wittenberg, and King 2001) to generate all predicted probabilities in this paper, varying the independent variable of interest and holding all other independent variables at their mean values.

[21] An argument can also be made that states become less risk-averse as their citizens become richer. This would imply that states with higher GDPs per capita should be more likely to sign an indefinite

risk measures is ideal, each is capturing some aspect of what the concept means at the level of states operating in anarchy. Moreover, two of the measures are relational, that is, a state's level of risk aversion depends on its particular partner. The other (growth) is an individual attribute of the state that does not vary with partnership. While the field of international relations is fond of using relational measures, the other measure I have created is closer to the economist's conception.

As a final test of the appropriateness of including the renegotiation costs, uncertainty, and risk variables in Models 1 through 3, I perform three Wald tests of the null hypothesis that the *joint* effect of the three variables equals zero.[22] The Wald tests yield strong support for my arguments: the joint effect of renegotiation costs, uncertainty and risk is significantly different from zero at $p < .005$ when the Bueno de Mesquita risk measure is used, at $p < .05$ when the Affinity risk measure is employed, and at $p < .01$ when the growth risk measure is used.

TABLES 2 AND 3 AROUND HERE

Hence the renegotiation costs, uncertainty, and risk aversion variables are performing as expected.[23] With regard to the other variables in the analyses, the following findings are of note. In all three models, agreements in the economics issue area are found to be considerably more likely to be of finite duration than are security agreements (p is between .007 and .038, depending on model

---

agreement. In an alternate specification, I used GDP per capita (taken from the Penn Table) as a proxy for risk-aversion. The results provided some evidence that states with higher GDPs per capita are more likely to sign indefinite agreements, although this result was not significant at standard levels ($p = .317$).

[22] I use a two-tailed test, which is conservative given that my alternative for each variable is in fact one-sided.

[23] In additional analyses, I included in Models 1, 2, and 3 a variable indicating whether a superpower was a signatory to assess the argument that power affects the outcomes of international negotiations. This variable had a positive coefficient, but it was very far from standard levels of statistical significance. The inclusion of the superpower variable does not affect the results displayed in Table 2.

specification). Agreements in the economics issue area do not, however, differ significantly from those in the human rights or environment issue area. Human rights and environmental agreements are also more likely to be of finite duration than are security agreements, but the differences are not uniformly statistically significant at conventional levels in all three model specifications.

My model also predicts that, if the parties choose to conclude a series of renegotiated agreements, increases in renegotiation costs (proxied by the logged number of signatories) will lead to agreements with longer intended durations (**CS-2)**. Conversely, increases in uncertainty and increases in relative risk aversion will lead states to choose agreements with a shorter duration (**CS-4** and **CS-6**, respectively). To test this theory, simply performing a regression analysis in which the dependent variable is the intended duration is problematic because the intended duration of some agreements is indefinite. If I limit the sample to only those agreements with a finite intended duration, I am selecting on the dependent variable and quite possibly introducing bias. A more appropriate way of testing is to set the intended duration of all indefinite treaties at some value greater than the longest finite treaty. I therefore set the intended duration of all indefinite treaties at four hypothetical values: 42, 70, 100 and 200 years. Because the data are right censored (that is, there is a maximum possible value), tobit regression analysis is used.

I conduct separate tobit analyses using the three risk aversion proxies. The results provide strong evidence for **CS-2**. At all four values used, the intended duration of finite agreements increases as renegotiation costs increase (*p* is between .001 and .057, depending on model specification). The results also provide strong evidence in support of **CS-4**. At all four values used, as uncertainty increases, the intended duration of finite agreements decreases notably (*p* is between .0001 and .001, depending on specification). At all four values used, the analyses provide some support for **CS-6,** suggesting that as risk aversion increases, the agreement's intended duration decreases. When the Bueno de Mesquita variable is used, the result is significant at between .028 and .058. When the *Affinity* measure is used, the result is significant at between .069 and .204. The coefficient on the growth risk variable is in the expected direction, but is far from standard levels of statistical significance (*p* is between .374 and .681). Table 4 displays the results of the tobit analyses in which all indefinite agreements are set at 100.

Wald tests on Models 4 through 6 provide considerable evidence that renegotiation costs, uncertainty, and risk *jointly* have a significant impact on the intended duration of agreements. The three variables' combined effect is statistically different from zero at $p < .005$ when the Bueno de Mesquita risk measure is used, at $p < .01$ when the Affinity risk measure is employed, and at $p < .05$ when the growth risk measure is used. These results add further and significant support for **CS-2**, **CS-4,** and **CS-6.**

Once again, agreements in the security issue area are significantly more likely to be indefinite than agreements in the other issue areas. This means that issue area is correlated with some variable(s) not explicitly included in my model that also affects agreement length. In any event, my theoretical predictions are supported by these results.

TABLE 4 AROUND HERE

4. Conclusion

How exactly do states manage to cooperate given the uncertainties plaguing the international environment? This paper, joining others like Downs and Rocke, offers one answer: States integrate flexibility into their cooperation so that they can respond to distributional shocks. A model of a particular kind of flexibility (duration provisions) is presented that highlights three independent variables: renegotiation costs, the degree of uncertainty, and relative risk aversion. Of course, other variables affect the duration provisions of international agreements, but there is no systematic theory identifying such variables. The model is then subjected to empirical testing which turns out to be quite supportive.

The model and the empirical testing help refine the literature on international cooperation in the following ways. The dependent variable is refined by taking it as a given that cooperation is possible and focusing instead on the more detailed and subtle issue of the form it takes. The independent variables are refined by looking beyond two-by-two games, like Prisoners' Dilemma, to distributional uncertainty and factors like risk attitude. Finally, and perhaps most important, scientific testing is introduced.

Duffield (2003:426) criticizes certain rationalist approaches to international cooperation for their reliance on case studies. Quite rightly he argues that nothing can be "proved" with a case study.[24] This paper meets the Duffield challenge by employing a large, scientifically-generated sample to test particular theoretical hypotheses of international cooperation. Nonetheless, there are those who may argue that some of the actual agreements in the random sample are trivial; hence such testing may be meaningless. In fact, many of the agreements are not the most exciting examples of international cooperation; nor are they worthy of case studies. Yet, (since the sample *is* random) such less important agreements govern most of the day-to-day cooperation between states. If cooperation is to be explained, such agreements deserve attention as a category despite their relative unimportance individually. Our intuitions about the forms of international cooperation are often shaped by a few high-profile agreements. High-profile agreements get that way because they have big effects; thus, it is important to be able to make statements about their characteristics. Future work will build a sample of important agreements, and then we will be able to make a comparison between the two samples. Only then can we say with any confidence whether past reliance on case studies was somewhat justified.

This analysis also sheds light on some broader issues of international politics. The descriptive statistics alone force us to reconsider some of the conventional wisdom in international cooperation, i.e., states tie their hands in order to make their commitments credible. It is hard to accept such a generalization given that such a large percentage of agreements in this sample – two thirds – are designed *not* to last forever. The theory suggests that flexibility may *enhance* credibility in those issue areas subject to shocks; reneging, which can be quite damaging to international cooperation, can often be avoided if renegotiation is an option instead.

[24] That is not to say that case studies are not important. I would argue they are essential for two reasons. First, they reveal the historical importance of our theoretical propositions. Second, they provide us with much-needed help regarding how to operationalize our theoretical variables.

Finally, two of the major and long-standing debates in IR – whether states can cooperate and whether international agreements and institutions matter – are addressed. IR, for decades, assumed an anarchic environment inside which states were behaving on the basis of their interests alone and free of any institutional constraints. Coalition building in the face of a common opponent was the only form of cooperation emerging in such an anarchic environment. Later, the possibility of cooperation was researched exhaustively. However, the study relied heavily on Prisoner's Dilemma games and, as a result, focused almost exclusively on the shadow of the future and the credibility of commitments.

My study, instead of affirming that cooperation is or is not possible, takes cooperation as a starting point and studies under what conditions this cooperation can be most successful. Instead of an anarchic international environment, or an environment where the shadow of the future is the only independent variable, we observe an environment structured by international agreements as institutions. Instead of random variations among agreements or automatic replication of the same agreement provisions over and over, we observe that the detailed provisions of international agreements are chosen in ways that increase the incidence and robustness of cooperation. These agreements are themselves important and consequential because they regulate cooperation, and the fact that they obey law-like regularities indicates that serious efforts are made for them to be able to regulate interactions in lasting and successful ways. So, *because* institutions matter, they are designed in rational ways, and the fact that people take time and make efforts to design them in such ways corroborates their significance.

References

Bennett, D. Scott, and Allan Stam. 2000. *EUGene*: A Conceptual Manual. *International Interactions* 26:179-204.

Bennett, D. Scott, and Allan C. Stam. 2000b. "A Cross-Validation of Bueno de Mesquita and Lalman's International Interaction Game." *British Journal of Political Science* 30: 541-60.

Bennett, D. Scott, and Allan C. Stam. 2000c. "A Universal Test of an Expected Utility Theory of War." *International Studies Quarterly* 44: 451-80.

Bordo, Michael and Finn Kydland. 1995. "The Gold Standard as a Rule: An Essay in Exploration." *Explorations in Economic History*. 32(4): 423-464.

Brown, Michael E., ed. 2000. *Rational Choice and Security Studies: Stephen Walt and His Critics.* Cambridge: MIT Press.

Bueno de Mesquita, Bruce. 1985. The War Trap Revisited: A Revised Expected Utility Model. *American Political Science Review*, 79: 156-177.

Conybeare, John. 1985. "Trade Wars: A Comparative Study of Anglo-Hanse, Franco-Italian, and Hawley-Smoot Conflicts." *World Politics*, 38(1): 147-172.

Dixon, William J. 1994. "Democracy and the Peaceful Settlement of International Conflict." *The American Political Science Review*, 88(1): 14-32.

Downs, George and David Rocke. 1995. *Optimal Imperfection.* Princeton: Princeton University Press.

Dye, Ronald A. 1985. "Optimal Length of Labor Contracts." *International Economic Review* 26: 233-50.

Duffield, John. 2003 "The Limits of 'Rational Design.'" *International Organization* 57 (2): 411-30.

Gartzke, Erik, and Dong-Joon Jo. 2002. "The Affinity of Nations Index, 1946-1996." *Version 3.0.*

Gray, Jo Anna. 1978. "On Indexation and Contract Length." *Journal of Political Economy* 86: 1-18.

Guzman, Andrew and Beth A. Simmons. 2002. "Litigation and Settlement." *Journal of Legal Studies* 31: S205-S235.

Harris, Milton and Bengt Holmstrom. 1987. "On the Duration of Agreements." *International Economic*

Review 26: 389-406.

Hart, Oliver and Bengt Holmstrom. 1987. The Theory of Contracts. In *Advances in Economic Theory,* edited by T.F. Bewley. Cambridge: Cambridge University Press.

Koremenos, Barbara 2001. "Loosening the Ties that Bind: A Learning Model of Agreement Flexibility" *International Organization* 55: 289-325.

Koremenos, Barbara 2000. "Bending But Not Breaking: Flexibility in International Financial and Monetary Agreements," Working Paper 1.73. University of California, Berkeley: Center for German and European Studies.

Kreps, David M. 1990. *A Course in Microeconomic Theory*. Princeton, N.J. : Princeton University Press.

Levy, Jack S. 1988. "When do Deterrent Threats Work?" *British Journal of Political Science*, 18(4): 485-512.

Martin, Lisa L. 2000. *Democratic Commitments: Legislatures and International Cooperation.*. Princeton, N.J. : Princeton University Press.

Martin, Lisa L. 1992. *Coercive Cooperation: Explaining Multilateral Economic Sanctions.* Princeton, N.J. : Princeton University Press.

Morrow, James D. 1987. On the Theoretical Basis of a Measure of National Risk Attitudes. *International Studies Quarterly*, 31(4): 423-438.

Oneal, John R. and Bruce M. Russett. 1997. "The Classical Liberals Were Right: Democracy, Interdependence, and Conflict, 1950-1985. *International Studies Quarterly*, 41(2): 267-293.

Osborne, Martin J. and Ariel Rubinstein. 1990. *Bargaining and Markets.* Boston: Academic Press.

Richards, Diana, T. Clifton Morgan, R.K. Wilson, Valerie L. Schwebach. 1993. Good Times, Bad Times and the Diversionary Use of Force. *Journal of Conflict Resolution*, 37(3): 504-535.

Rosendorff, Peter and Helen Milner. 2001. "The Optimal Design of International Trade Institutions: The Sources and Impact of Escape Clauses" In Koremenos, Barbara, Charles Lipson, and Duncan

Snidal (eds.) *The Rational Design of International Institutions*. forthcoming in *International Organization*, special issue.

Signorino, Curtis and Jeffrey Ritter. 1999. Tau-b or Not Tau-b: Measuring the Similarity of Foreign Policy Positions. *International Studies Quarterly*, 43: 115-144.

Simmons, Beth A. 2000. "International Law and State Behavior: Commitment and Compliance in International Monetary Affairs." *The American Political Science Review*, 94 (4): 819-835.

Smith, Alastair. 1996. Diversionary Foreign Policy in Democratic Systems. *International Studies Quarterly*, 40(1): 133-153.

Smith, Alastair. 1998. International Crises and Domestic Politics. *American Political Science Review*, 92(3): 623-638.

Tomz, Michael, Jason Wittenberg, and Gary King. 2001. "Clarify: Software for Interpreting and Presenting Statistical Results." http://gking.harvard.edu/clarify.

Appendix A

Assuming State A and B are symmetric, compare State A's expected utility with one two-period (inflexible) agreement, given by $EU_2 = E(u(b_1 + m_0 + e_1 - k_n) + ?u(b_1 + m_0 + e_1 + e_2))$, to its expected utility from two one-period (flexible) agreements, given by EU_1 ? $E(u(b_1 + m_0 + e_1 - k_n) + ?u(b_1 + m_0 + e_2 - k_r))$. The only differences between EU_1 and EU_2 occur in the second period. In that period, with EU_2 the states again live with the first period shock e_1 but do not pay renegotiation cost k_r. The reverse is true for EU_1, where the states avoid having to experience e_1 by paying the renegotiation cost k_r.

With respect to **CS1**, the key point is that EU_1 is decreasing in the renegotiation cost k_r while EU_2 is unaffected by k_r. Thus, suppose that initially EU_2 ? EU_1; that is, suppose that initially an inflexible agreement dominates. Holding the distribution of e_1 fixed, we can then decrease k_r until, at some point, a threshold is reached and EU_1 ? EU_2. States will then change from an inflexible two-period agreement to a flexible series of one-period agreements. A similar story holds in the other direction when we start with the case where two one-period agreements are preferred. Increasing k_r will decrease EU_1 until at some point the two expected utilities are equal, and after which point the states jointly prefer an inflexible two-period agreement because the increased cost of renegotiation makes the one-period agreements no longer optimal.

With respect to **CS3**, holding the form of the utility function constant and assuming it is strictly increasing and concave, a mean-preserving increase in the variance of the shocks to the distribution of gains (i.e., an increase in uncertainty), e, reduces both EU_1 and EU_2, because e_1 and e_2 enter both of them and because concavity of the utility function implies risk aversion, which means that the states prefer a lower variance to the shocks at a given mean (in this case zero). However, because the shocks are independent, in the second period the combined shock $(e_1 ? e_2)$ experienced under the two-period agreement has a variance equal to $2\text{var}(e)$, while the second period shock in the case of two one-period agreements, which is just e_2, has only $\text{var}(e)$. As a result, a given increase in $\text{var}(e)$, holding the mean

constant, increases the variance of the combined shock more for the two-period agreement, and thereby leads to a larger decrease in EU_2 than in EU_1. That is, increasing the variance of the shocks decreases EU_2 relative to EU_1. As a result, if we start from a situation where an inflexible two-period agreement is preferred by the states, and then continually increase the var(e), holding constant k_r and the mean of the shocks, at some point EU_2 will fall sufficiently in relation to EU_1 that the states will prefer a flexible regime of two one-period agreements. At this point, it is worth it to the states to pay k_r in order to reduce the variance of the second period outcomes. The same argument holds in the other direction as well. If we start with a case where the states prefer two one-period agreements, we can decrease var(e), which will now increase EU_2 relative to EU_1, until a point is reached where the states prefer an inflexible two-period agreement. The limiting case, where var(e) ? 0, makes this particularly clear, as then states will clearly prefer a two-period agreement as long as k_r ? 0.

A similar argument holds with respect to **CS5**. Here we hold the variance of the shocks constant while varying the concavity of the utility function. The key point here is that the variance of the outcomes is higher under the inflexible two-period agreement, due to the fact that states experience the combined shock $(e_1 ? e_2)$ in the second period. As a result, increasing the concavity of the utility function (i.e., increasing the degree of risk aversion) reduces both expected utilities, EU_1 and EU_2, but decreases expected utility under the two-period agreement more than it does expected utility under two one-period agreements. Suppose we start from a situation where the two-period agreement is preferred by both states. Now we gradually increase the concavity of the utility function, say by increasing the degree of the exponent, a, when $u(x) ? x^a$ as in the simulations, or by increasing the risk aversion parameter a in a constant absolute risk aversion (CARA) utility function where $u(x) ? -? e^{? ax} ? ?$ (Kreps 1990: 86). In either case, as we increase the level of risk aversion, EU_2 declines relative to EU_1, until at some point $EU_1 ? EU_2$ and the states now both prefer to enter into two one period agreements. As above, the same arguments work in reverse if we start from a situation where the states choose two one-period agreements,

and then decrease the level of concavity (and therefore risk aversion) in the utility function. Eventually the states will become risk tolerant enough that they prefer an inflexible two-period agreement to two one-period agreements, so long as $k_r ? 0$. In the limiting case of a linear utility function and risk neutrality, the states will always prefer the longer agreement if $k_r ? 0$; that is, if renegotiation costs are non-zero.

Appendix B

This appendix describes the construction of the solutions to the discretized version of the two-period model used for the simulations. The value of base outcome, *a*, is set to 40.0. This value is completely arbitrary as long as it is positive; its main function in the simulations is to be large enough to keep the outcome from going negative even after a series of bad agreement shocks. The value of each state's share of the initial distribution of gains is set to $m ? 4$, which implies a total gain from the agreement, given equal ex ante division due to Nash bargaining, of $g ? 2m ? 8$. The general class of utility functions considered in the simulation is polynomial. In particular, I consider $U(X) ? X$, $U(X) ? X^2$, $U(X) ? X^3$, $U(X) ? X^4$ or, in words, a linear utility function, a quadratic utility function a cubic utility function, and a quartic utility function. For the values of *X* considered here, which are all positive and greater than one, an increase in the exponent of the utility function represents an increase in the level of risk aversion.

To keep the programming simple, the distribution of shocks is discrete, not continuous. The shock in the first period is denoted $?_1$ and the shock in the second period is denoted $?_2$. The potential values of the shocks to the distribution of gains under the agreement are denoted by $?_{1j}$ and $?_{2j}$ where $j ? \{1,2,3\}$ denotes one of the three possible values of the shocks. In the simulations in Table B1, the values of the shocks are –2, 0 and 2 so that $?_{11} ? ?2$ and so on. The variance of the distribution of shocks is increased/decreased through a mean-preserving increase/decrease in the spread of the distribution; in this case, there is a symmetric increase in the probability of the two extreme values, with a corresponding reduction in the probability of a zero shock.

I examined values for the costs of negotiation, k_n, and renegotiation, k_r, in the range $[0,1]$. I considered values of the discount rate, $?$, in the set $\{0.6,0.7,0.8,0.9\}$ while the values of the costs of reneging, *c*, ranged over the interval $[0,5]$. For values of *c* near the high end of this range (how high

depends on the shape of the utility function) reneging does not occur given the values of the shocks utilized here.

For each of the variables corresponding to the comparative static results, I examined essentially the entire range of possible values (or distributions) within this framework; all were (not surprisingly given the proofs in Appendix A), consistent with the model. The particular values presented in Table B1 were chosen to represent the wide range of simulation results actually generated.

TABLE B1 ABOUT HERE

The expected utility associated with two one-period agreements is given by

$$E(U_R) = \sum_{j=1}^{3} \Pr(\theta_1 = \theta_{1j}) U(a - m - \theta_{1j}) + \sum_{j=1}^{3} \Pr(\theta_2 = \theta_{2j}) U(a - m - \theta_{2j} - k_r),$$

where $U(\cdot)$ is the utility function.

The equation for a single two-period agreement is more complicated due to the possibility of reneging. Consider the problem from the standpoint of state A. We can define an indicator variable, R_A, for whether or not State A will renege as follows:

$$R_A(\theta_{1j}) = 1[E(U(a - m - \theta_2 - k_r - c)) > E(U(a - m - \theta_{1j} - \theta_2))],$$

where the expectations are over the possible values of θ_2. The expectations are as of the end of the first period, so that θ_1 is known but θ_2 is not. In words, State A compares the expected utility of reneging, given by the first term, with the expected utility of not reneging, given by the second term. The gain from reneging is getting rid of θ_1; the cost is the combined cost of reneging plus renegotiation, given by

$(k_r ? c)$. State B makes a similar calculation; we can define an indicator variable R_B for its decision as follows:

$$R_B(?_{1j}) ? 1[E(U(a ? m ? ?_2 ? k_r ? c)) ? E(U(a ? m ? ?_{1j} ? ?_2))].$$

With these indicator variables in hand, we can now give the formula for State A's expected utility from a two-period agreement. It is

$$E(U_N) ? \sum_{j?1}^{3} \Pr(?_1 ? ?_{1j})\{U(a ? m ? ?_{1j} ? k_n) ? ?[R_A(?_{1j})T_1 ? R_B(?_{1j})T_2 ? (1 ? R_A(?_{1j}) ? R_B(?_{1j}))T_3]\}$$

where

$$T_1 ? \sum_{j?1}^{3} \Pr(?_2 ? ?_{2j})U(a ? m ? ?_{2j} ? c ? k_r),$$

$$T_2 ? \sum_{j?1}^{3} \Pr(?_2 ? ?_{2j})U(a ? m ? ?_{2j} ? k_r),$$

and

$$T_3 ? \sum_{j?1}^{3} \Pr(?_2 ? ?_{2j})U(a ? m ? ?_{1j} ? ?_{2j}).$$

The formulae for State B are symmetric.

TABLE 1
INTENDED DURATION OF AGREEMENTS

	FINITE OR INDEFINITE? (p-value from test of independence: 0.003)	
Issue Area	**Percent Finite**	**Percent Indefinite**
Economics	79.7	21.3
Environment	60.0	40.0
Human Rights	44.0	56.0
Security	48.0	52.0
All Agreements	66.4	33.6

	MEAN DURATION OF FINITE AGREEMENTS (p-value from test of equality of means: 0.107)
Issue Area	**Mean Duration (Standard Deviation of Duration)**
Economics	12.1 (11.6)
Environment	6.6 (3.5)
Human Rights	8.4 (3.2)
Security	5.0 (0.0)
All Agreements	10.0 (9.4)

Source: Author's calculations using data on international agreements. The total sample contains 97 finite agreements and 49 indefinite agreements. Within the economics issue area, the sample includes 19 financial agreements, 17 investment agreements and 9 monetary agreements

TABLE 2. RESULTS OF PROBIT ANALYSES OF THE PRESENCE OF A FINITE AGREEMENT

	MODEL 1	**MODEL 2**	**MODEL 3**
Independent variable	**Coefficient (Std. Errors)**	**Coefficient (Std. Errors)**	**Coefficient (Std. Errors)**
Number of participants (logged)	-.487 (.142)***	-.881 (.230)***	-.341 (.139)***
Uncertainty	1.072 (.276)***	1.074 (.296)***	1.092 (.279)***
Risk aversion (Bueno de Mesquita measure)	.631 (.269) **	-----	-----
Risk aversion (Gartzke and Jo measure)	-----	.393 (.316)	-----
Risk aversion (GDP growth measure)	-----	-----	.602 (.384)
Security issue area	(omitted category)	(omitted category)	(omitted category)
Human rights issue area	.766 (.500)	.957 (.538)*	.613 (.501)
Environmental issue area	.781 (.431)*	.775 (.439)*	.635 (.419)
Economics issue area	.997 (.370)***	.768 (.369)**	.742 (.359)**
Constant	-.503 (.413)	.039 (.509)	-.490 (.415)
$?^2$, Wald test of joint significance of Number of participants, Uncertainty, Risk aversion	7.43****	4.90**	7.17***
Indefinite agreements correctly predicted	69.39%	77.78%	73.91 %
Finite agreements correctly predicted	75.00%	74.44%	69.57%
Number of observations	145	135	138

Source: Author's calculations using data on international agreements. *p < .10; **p < .05; ***p < .01; ****p < .005. The Bueno de Mesquita and Gartzke and Jo risk aversion variables are continuous, and range from –1 (least risk-averse) to +1 (most risk-averse). The GDP growth risk variable is dichotomous, and equals 1 for states that have mid levels of growth (and are therefore believed to be most risk-averse). Percent correctly predicted figures are calculated using the sample mean of the dependent variable as the cutoff point. Predicted probabilities smaller than the sample mean of .664 are coded 0 (indefinite) and predicted probabilities greater than or equal to .664 are coded 1 (finite). This provides a more fine-tuned assessment than would be obtained by simply using .5 as the cutoff point.

TABLE 3. PREDICTED PROBABILITY OF THE PRESENCE OF A FINITE AGREEMENT BASED ON CHANGES IN INDEPENDENT VARIABLES OF INTEREST

	Predicted probability at minimum value of independent variable	**Predicted probability at mean value of independent variable**	**Predicted probability at maximum value of independent variable**
Number of participants (Model 1)	.797	.697	.148
Uncertainty (Model 1)	.414	-----	.798
Risk aversion (Bueno de Mesquita measure) (Model 1)	.388	.697	.827
Risk aversion (Gartzke and Jo measure) (Model 2)	.575	.680	.817
Risk aversion (GDP growth measure) (Model 3)	.652	-----	.825

Predicted probabilities are based on results displayed in Table 2 and were calculated using *Clarify*. All other independent variables were held constant at their mean values.

TABLE 4. RESULTS OF TOBIT ANALYSES OF THE INTENDED DURATION OF FINITE AGREEMENTS			
	MODEL 4	**MODEL 5**	**MODEL 6**
Independent variable	**Coefficient (Std. Errors)**	**Coefficient (Std. Errors)**	**Coefficient (Std. Errors)**
Number of participants (logged)	20.046 (7.125)***	35.704 (10.293)***	14.519 (7.523)*
Uncertainty	-54.146 (15.555)***	-56.247 (16.749)***	-60.430 (16.881)***
Risk aversion (Bueno de Mesquita measure)	-29.636 (13.391)**	-----	-----
Risk aversion (Gartzke and Jo measure)	-----	-29.751 (16.199)*	-----
Risk aversion (GDP growth measure)	-----	-----	-12.512 (18.137)
Security issue area	(Omitted Category)	(Omitted Category)	(Omitted Category)
Human rights issue area	-80.905 (30.958)***	-95.246 (32.959)***	-80.018 (32.626)**
Environmental issue area	-97.156 (28.835)***	-103.811 (30.130)***	-97.199 (30.182)***
Economics issue area	-83.980 (26.918)***	-75.925 (27.747)***	-75.453 (28.421)***
Constant	165.033 (30.230)***	132.276 (33.248)***	168.698 (31.935)***
$?^2$, Wald test of joint significance of Number of participants, Uncertainty, Risk aversion	9.26****	6.34***	5.51**
$?^2$	48.03	47.35	39.37
Number of observations	110 (61 uncensored)	102 (57 uncensored)	104 (50 censored)

Source: Author's calculations using data on international agreements. *$p < .10$; **$p < .05$; ***$p < .01$; ****$p < .005$. The Bueno de Mesquita and Gartzke and Jo risk aversion variables are continuous, and range from –1 (least risk-averse) to +1 (most risk-averse). The GDP growth risk variable is dichotomous, and equals 1 for states that have mid levels of growth (and are therefore believed to be most risk-averse). Indefinite agreements are set at an intended duration of 100 years.

TABLE B1
RESULTS FROM SIMULATION OF TWO-PERIOD MODEL

Utility Function	?	var(?)	k_n	k_r	c	$E(U_R)$	$E(U_N)$	Choice
Base Case								
Square Root	0.9	Low	1.0	0.5	1.0	12.4927	12.5130	One Two-Period (Inflexible)
Decrease Renegotiation Costs								
Square Root	0.9	Low	1.0	0.1	1.0	12.5199	12.5185	Two One-Period (Flexible)
Increase Variance of Shocks (Uncertainty)								
Square Root	0.9	High	1.0	0.5	1.0	12.4900	12.4557	Two One-Period (Flexible)
Increase Risk Aversion								
Cubic	0.9	Low	1.0	0.5	1.0	6.6682	6.6658	Two One-Period (Flexible)