Supplementary Appendix: Audience Costs and the Dynamics of War and Peace

Casey Crisman-Cox* Michael Gibilisco†

Abstract

This document contains the Appendices for "Audience Costs and the Dynamics of War and Peace." In Appendix A, we derive the equilibrium constraint, and we characterize an example equilibrium in Appendix B. In Appendix C, we discuss our implementation of the CMLE. Appendix D presents Monte Carlo Experiments, and Appendix E contains a discussion of time-invariant covariates. In Appendix F, we provide a table form of the audience cost parameters and their associated standard errors, and Appendix G illustrates the distribution of audience costs in autocratic regimes. We discuss the procedure for substantive effects in Section H. Appendices I and J contain additional robustness checks and substantive effects, respectively. Finally, Appendix K contains model fit exercises.

A Characterizing Equilibria

As in Aguirregabiria and Mira (2007), we characterize equilibria with dynamic expected utilities. Let $v_i(a_i, s)$ denote i's net-of-shock expected utility from choosing action a_i in state s and continuing to the play the game for an infinite number of periods, and write $v_i = (v_i(a_i, s))_{(a_i,s)\in A^2}$ for every country i. In other words, given a vector of expected values v_i and a vector of random shocks ε_i, country i chooses action a_i in state s if and essentially only if

$$a_i = \underset{a_i \in \{1,2,3\}}{\text{argmax}} \{v_i(a_i, s) + \varepsilon_i(a_i)\}.$$

Thus, v_i is identical to a cut-off strategy for country i. Because ε_i is distributed type 1 extreme value, i chooses a_i in state s with probability $P(a_i, s; v_i)$, where

$$P(a_i, s; v_i) = \frac{\exp(v_i(a_i, s))}{\sum_{a_i'} \exp(v_i(a_i', s))}\,. \tag{5}$$

*Texas A&M University Email: c.crisman-cox@tamu.edu

†California Institute of Technology. Email: michael.gibilisco@caltech.edu

If g is the distribution of ε_i, described above, we write country i's average expected utility in state s as $G(s, v_i)$, which takes the form

$$G(s, v_i) = \int \max_{a_i} \{v_i(a_i, s) + \varepsilon_i(a_i)\} \, g(\varepsilon_i) d\varepsilon_i,$$

and simplifies to

$$G(s, v_i) = \log \left(\sum_{a_i} \exp(v_i(a_i, s)) \right) + C$$

where C is Euler's constant (McFadden 1978, Corollary p. 82). Consider a profile $v = (v_i, v_j)$ of action-state values. Then country i's iterative value of action a_i in state s, denoted $\Phi_{ij}(a_i, s, v; \theta)$, is written as

$$\Phi_{ij}(a_i, s, v; \theta) = \sum_{a_j} \underbrace{P(a_j, s; v_j)}_{\substack{\textbf{expectation} \\ \textbf{over } j\textbf{'s actions}}} \Big[\underbrace{u_{ij}(a_i, a_j, s; \theta)}_{\substack{\textbf{today's} \\ \textbf{payoff}}} + \delta \underbrace{G\left(\max\{a_i, a_j\}, v_i\right)}_{\substack{\textbf{expectation over} \\ \textbf{tomorrow's payoff}}} \Big]. \tag{6}$$

In words, Equation 6 takes a profile of values v, supposes countries play according to the associated choice probabilities, and then returns new expected values of each action in each state. The iterative value, $\Phi_{ij}(a_i, s, v; \theta)$, is comprised of three components. First, country i weights its opponent's actions by the corresponding choice probabilities, $P(a_j, s; v_j)$. Second, country i receives an immediate payoff, $u_{ij}(a_i, a_j, s; \theta)$. Finally, country i receives a discounted expected future payoff, $\delta G\left(\max\{a_i, a_j\}, v_i\right)$. The profile v is an equilibrium if and only if $\Phi_{ij}(a_i, s, v; \theta) = v_i(a_i, s)$ for every country i, every action a_i, and every state s. Hence, v is an equilibrium profile if and only if it is a fixed point of these iterative value functions. Formally, write $\Phi_{ij}(v; \theta)$ as $\Phi_{ij}(v; \theta) = \times_{a_i} \times_s \Phi_{ij}(a_i, s, v; \theta)$ and $\Phi(v; \theta) = \Phi_{ij}(v; \theta) \times \Phi_{ji}(v; \theta)$. Then an equilibrium is a profile v such that $\Phi(v; \theta) = v$.

Notice the that function $\Phi_{ij}(a_i, s, v; \theta)$ is a weighted sum of current stage utilities and discounted expected payoffs. When the latter are sufficiently bounded, the continuous function Φ maps a convex and compact set into itself, so an equilibrium exists. Formally, define $B \geq 0$ as $B = \max_{i,j,a,s} \{|u_{ij}(a, s; \theta)|\}$, and B^- and B^+ where $B^- = -(1+\delta)B + \delta \log(3)$ and $B^+ = (1+\delta)B + \delta \log(3)$. Thus, B^- and B^+ represent the smallest and largest possible expected action-state values, respectively, in any equilibrium. In addition, for all $v \in [B^-, B^+]^{18}$, $\Phi(v; \theta) \in [B^-, B^+]^{18}$ because the iterative expected utility of each action-state can be no larger or smaller than B^+ or B^-, respectively, when the $v \in [B^-, B^+]^{18}$. Thus, the continuous function Φ maps a convex and compact set into itself, so Φ admits a fixed point, an equilibrium.

Table 4: Structural Parameters and Values in the Example Equilibrium

Parameter	$x_{ij} \cdot \beta(2)$	$x_{ij} \cdot \beta(3)$	$z_i \cdot \kappa(2)$	$z_i \cdot \kappa(3)$	$\gamma(1)$	$\gamma(2)$	$\gamma(3)$	α_i
Value	15	15	-16	-16	-5	5	7	-15

B Example Equilibrium

We consider a version of the game parameterized by the values in Table 4. We choose the values for two reasons. First, they are similar in direction and magnitude to those we estimate in the data. Second, under these values, a symmetric equilibrium exists, and we characterize a symmetric equilibrium to simplify the exposition. As discussed in the paper, we still maintain the normalization that $x_{ij} \cdot \beta(1) = 0$ and $z_i \cdot \kappa(1) = 0$.

Given the values in Table 4, the second column in Table 5 reports a solution to the equation $\Phi(v) - v = 0$, where Φ is defined in Equation 6. When v takes on these (non-rounded) values, the equilibrium constraints are satisfied below a tolerance of $1e^{-10}$, that is,

$$\max_{(i,a_i,s)} \{|\Phi_{ij}(a_i, s; v) - v_i(a_i, s)|\} < 1e^{-10}.$$

Although we characterize an equilibrium for specific parameter values, the equilibrium discussed here will change in a continuous manner for sufficiently small perturbations of the underlying parameters that enter the constraint equation, $\Phi(v) - v = 0$, in a continuously differentiable manner. To show this, we verify that the Jacobian of the constraint equation, $\Phi(v) - v$, has full rank at the equilibrium of interest.

The Table's third column reports the corresponding choice probabilities. Using the equilibrium choice probabilities and the per-period transition function $s^{t+1} = \max\{a_i^t, a_j^t\}$, we construct the equilibrium transition matrix, Q^v where the entry $Q^v[s, s']$ denotes the probability that game transitions from s to state s' given equilibrium v. Specifically, for any state of hostilities s, we have

$$\begin{aligned}
Q^v[s,1] &= P(1, s; v_i) \cdot P(1, s; v_j), \\
Q^v[s,2] &= P(3, s; v_i) + P(3, s; v_j), \text{ and} \\
Q^v[s,3] &= 1 - Q^v[1, s] - Q^v[3, s].
\end{aligned}$$

Table 5: Equilibrium Expected Utilities and Choice Probabilities

(a_i, s)	$v_i(a_i, s)$	$P(a_i, s; v_i)$
$(1,1)$	20.30	0.88
$(2,1)$	18.14	0.10
$(3,1)$	16.12	0.01
$(1,2)$	34.43	0.02
$(2,2)$	38.37	0.84
$(3,2)$	36.61	0.15
$(1,3)$	36.03	0.80
$(2,3)$	34.44	0.16
$(3,3)$	33.31	0.04

In this example, the equilibrium transition matrix takes the following form:

$$Q^v = \begin{bmatrix} 0.78 & 0.19 & 0.03 \\ 0.00 & 0.73 & 0.27 \\ 0.63 & 0.28 & 0.09 \end{bmatrix}.$$

Notice these values are rounded to the second digit. Without rounding, $Q^v[2,1] > 0$.

An invariant or stationary distribution π^v describes the distribution of states in the long run, where $\pi^v(s)$ is the probability of observing state s from the path of play given equilibrium v. Given the transition matrix Q^v, π solves the equation

$$\pi^v Q^v - \pi^v = 0. \tag{7}$$

Because $P(a_i, s; v_i) > 0$ for all actions a_i and all states s, every state s' can be reached from any state s. Thus, the Markov chain described by transition matrix Q^v is aperiodic, so a unique invariant distribution exists. In this example, π^v takes on the following values:

$$\pi^v = (0.41, 0.44, 0.14).$$

With the invariant distribution in hand, we can compute the probability that i initiates a dispute along the path of play given equilibrium v:

$$f_i(v) = \pi^v(1) \cdot (1 - P(1, 1; v_i)). \tag{8}$$

In other words, i initiates a dispute when the current state is s and i plays a non-peaceful action $a_i > 1$. In our numerical example, $f_i(v) = 0.05$.

Effects on Conflict Initiation. We examine how i's probability of initiating a conflict changes as we make the country's audience cost parameter more negative. To ease exposition, we focus on country 1, while noting that the results are symmetric. Specifically, we compute

$$-\frac{\partial f_1}{\partial \alpha_1} = -\frac{\partial f_1}{\partial v} \cdot \frac{\partial v}{\partial \alpha_1}.$$

The gradient $\frac{\partial v}{\partial \alpha_1}$ can be computed using the Implicit Function Theorem and the equilibrium constraint $\Phi(v) - v = 0$, where

$$\frac{\partial v}{\partial \alpha_1} = -\frac{\partial \Phi}{\partial \alpha_1} \cdot \left(\frac{\partial \Phi}{\partial v} - I_{18} \right)^{-1}$$

and I_k is the $k \times k$ identity matrix. Finally, the gradient $\frac{\partial f_1}{\partial v}$ can be computed using the product rule, where $\frac{\partial P(1,1;v_i)}{\partial v}$ has a closed form solution and $\frac{\partial \pi^v(1)}{\partial v}$ can also be computed using the Implicit Function Theorem and the equation $\pi Q^v - \pi = 0$ although we use numerical derivatives here in our analysis. In the example under consideration, $-\frac{\partial f_1}{\partial \alpha_1} = 0.022$. Using linear interpolation, increasing the magnitude of country 1's audience costs by 25% changes 1's probability of initiating from 0.05 to 0.14. The direction of this effect generally matches the equilibria estimated from the data and reported in Table 2.

Likewise, we can consider how country 2's audience costs affect country 1's initiation probability, that is,

$$-\frac{\partial f_1}{\partial \alpha_2} = -\frac{\partial f_1}{\partial v} \cdot \frac{\partial v}{\partial \alpha_2}.$$

where we compute the component gradients as described above. In our equilibrium of interest, $-\frac{\partial f_1}{\partial \alpha_2} = -0.023$. Again, using linear interpolation, decreasing α_2 to -16 from -15 suggests that $f_1(v)$ drops to 0.03 from 0.05, and this effect matches the equilibrium effects found in the data (see Table 2).

Competing Effects on Peace. On the one hand, larger audience costs (more negative) for country i encourages the country to initiate conflict in the peace state, i.e., $-\frac{\partial f_1}{\partial \alpha_1} > 0$. On the other hand, if country i has enhanced audience costs, it's rival is less like to initiate conflict. What are the total effects of larger audience costs on peace? We compute $\frac{\partial \pi^v}{\partial \alpha_i}$, and find that $-\frac{\partial \pi^v}{\partial \alpha_i} < 0$, so increasing country i's audience costs discourage peace in this equilibrium. The effect is relatively small, however, where $-\frac{\partial \pi^v}{\partial \alpha_i} = -0.009$. Notice that audience costs discouraging peace is an effect we do not generally see in the equilibria estimated in the data, i.e., those effects reported in Table 2.

Standing Firm in Crises. We examine as to whether larger audience costs increase or decrease a country's probability of standing firm in a crisis. Conditional on the path of play beginning in state $s = 2$, i's probability of maintaining or escalating the dispute is

$$g_i(v_i) = 1 - P(1, 2; v_i).$$

In our example equilibrium, i's probability of maintaining or escalating the crisis is $g_i(v_i) = 0.98$. As above, we can compute the effects of larger (more negative) audience costs on this probability for country 1. In our example, $-\frac{\partial g_i(v_i)}{\partial \alpha_1} = 0.012$, demonstrating that larger audience costs for country 1 encourage it to stand firm in crises. And similar results hold when looking at the war state $s = 3$. When increasing country 2's audience costs, country 1's conditional probability of maintaining or escalating the crisis also increases, where $-\frac{\partial g_i(v_i)}{\partial \alpha_1} = 0.004$, but the effect is substantially smaller.

Probability of Receiving Audience Costs. In an equilibrium v, country i's long-term probability of backing down and receiving an audience cost can be computed as

$$h_i(v) = \pi^v(2) \cdot P(1, 2; v_i) \cdot (1 - P(1, 2; v_j)) + \pi^v(3) \cdot (1 - P(3, 3; v_i)) \cdot P(3, 3; v_j). \tag{9}$$

In Equation 9, $\pi^v(2)$ denotes the probability that the equilibrium path of play is in the crisis state, and $P(1, 2; v_i) \cdot (1 - P(1, 2; v_j))$ is the probability that country i receives an audience cost in the same state. Likewise, $\pi^v(3)$ denotes the probability that the path of play is in the war state, and $(1 - P(3, 3; v_i)) \cdot P(3, 3; v_j)$ is the probability that country i receives an audience cost war. In our example equilibrium, $h_i(v) = 0.01$. In addition, larger audience costs (more negative) for country i decreases this probability even further, that is, $-\frac{\partial h_i(v)}{\partial \alpha_i} < 0$, which matches the effects from the estimated equilibria.

C Implementation

In this Appendix, we detail our implementation of the CMLE. Our data contains 125 countries in 179 games, and solving the constrained optimization in Equation 4 requires estimating more than $3,347$ parameters, where $3,222$ are expected utility constraints. The high-dimensionality raises some questions about feasibility. To perform the optimization we use the program IPOPT (Interior Point OPTimizer), which is an open-source, industrial optimizer used to solve problems with potentially hundreds of thousands of variables (Wächter and Biegler 2006). IPOPT is particularly well suited to the large problem here. In time trials, IPOPT had better convergence and performance properties than other optimizers such as KNITRO and a version of the Augmented Lagrangian Method. Throughout, we set our convergence criterion to IPOPT's default at $1e^{-6}$.

A general drawback of interior-point methods is that they require accurate representations of the Hessian of the Lagrangian for the problem in Equation 4. In our experiments, numerical approximations using finite differences substantially inflate the estimator's variance. To work around this, we compute the Hessian and all other derivatives using the program ADOL-C which implements an algorithmic differentiation (AD) routine (Griewank, Juedes and Utke 1996). In our set-up, we supply only the log-likelihood and constraint function, and the AD program produce the derivatives by repeatedly applying the chain rule to the supplied functions. We implement the estimator using Python 2.7 on Xubuntu 14.04 using the pyipopt software developed by Xu (2014) to call IPOPT within Python and the pyadolc package developed by Walter (2014) to use the AD routines discussed above. Asymptotic standard errors are estimated using Silvey (1959, Lemma 6, p. 401).

For comparison, we also simulate standard errors using a parametric bootstrap from Davison and Hinkley (1997). Overall, the bootstrapped standard errors closely match the analytical ones; however, countries involved in only one dyad with few (one or two) non-peaceful states have, on average, larger bootstrapped standard errors, associated with their audience cost parameter, than analytical ones. These countries, e.g., Ghana, comprise only 10% of those in our sample.

D Monte Carlo Experiments

In this Appendix, we describe a Monte Carlo experiment in which we use simulated data to evaluate the performance of the method and our implementation as a function of the number of dyads and time periods. The results of this experiment are important for two reasons. They demonstrate that firstly, the parameters of interest are identified, and secondly, the estimation procedure accurately recovers the model's parameters for numbers of dyads and time periods that are similar to our dataset used in the study.

In this experiment, $x_{ij}^s = (1, x_{ij}^1)$ for all s and $z_i = (z_i^1)$, where x_{ij}^1 and z_i^1 are random variables. In addition, we vary the number of countries N to be values in $\{10, 20, 30\}$ and T to be values in $\{20, 80, 150, 250\}$. We consider every possible combination of countries, that

Table 6: Coefficients used in the first Monte Carlo experiment analyzing the performance of the CMLE as a function of N and T.

Coefficient	$\beta(2)$	$\beta(3)$	$\kappa(2)$	$\kappa(3)$	$\gamma(1)$	$\gamma(2)$	$\gamma(3)$	α_i
Value	$(-1,1)$	$(-2,2)$	-0.5	-1	-0.5	0	0.5	$-2+\frac{2(i-1)}{N-1}$

is, $\mathcal{D} = \{\{i,j\} \mid i,j \in \{1,...,N\}, i \neq j\}$ and $D = \binom{N}{2}$. These values capture those in the real-world application below in which $T = 180$ and $D \approx \binom{20}{2}$.

The experiment is conducted as follows. We fix the coefficients used throughout to those in Table 6. For each fixed value of N and T, we first generate control variables $x_{ij}^1 \sim N(0,1)$ and $z_i^1 \sim U(0,1)$. Then for each unordered dyad k, we compute an equilibrium v^k by solving the system of equations generated from Eq. 6. Next, we generate T periods of data using the computed equilibrium, the associated conditional choice probabilities in Eq. 5, and the transition $s^{kt+1} = \max\{a_{i^k}^{kt}, a_{j^k}^{kt}\}$. The initial state s^{k1} is drawn from $\{1,2,3\}$ with equal probability. After a suitable burn-in period, we combine the generated data and estimate the parameters of the model by solving the constrained optimization problem in Eq. 4 using the tools described in the previous section. The procedure is repeated 50 times for each value of N and T.

Table 7: Summary of Monte Carlo Experiment

N	T	$\hat{\beta}(2)_1$	$\hat{\beta}(2)_2$	$\hat{\beta}(3)_1$	$\hat{\beta}(3)_2$	$\hat{\kappa}(2)_1$	$\hat{\kappa}(3)_1$	$\hat{\gamma}(1)$	$\hat{\gamma}(2)$	$\hat{\gamma}(3)$
10	20	-1.10 (0.24)	0.98 (0.15)	-2.02 (0.28)	2.02 (0.16)	-0.43 (0.20)	-1.02 (0.28)	-0.62 (0.46)	0.07 (0.32)	0.45 (0.19)
10	80	-1.05 (0.11)	1.00 (0.08)	-2.05 (0.11)	2.01 (0.08)	-0.45 (0.09)	-0.95 (0.12)	-0.45 (0.22)	0.02 (0.16)	0.48 (0.08)
10	150	-1.02 (0.10)	1.00 (0.06)	-2.00 (0.08)	2.02 (0.06)	-0.50 (0.06)	-1.01 (0.10)	-0.50 (0.19)	0.03 (0.11)	0.50 (0.07)
10	250	-1.01 (0.06)	0.99 (0.04)	-1.99 (0.08)	2.00 (0.04)	-0.50 (0.04)	-1.01 (0.07)	-0.57 (0.16)	0.03 (0.11)	0.49 (0.05)
20	20	-1.03 (0.13)	1.02 (0.09)	-1.98 (0.13)	2.01 (0.07)	-0.48 (0.11)	-1.02 (0.14)	-0.49 (0.29)	-0.01 (0.19)	0.49 (0.10)
20	80	-1.01 (0.06)	0.99 (0.03)	-2.00 (0.06)	1.99 (0.04)	-0.48 (0.04)	-0.99 (0.08)	-0.50 (0.09)	-0.02 (0.06)	0.49 (0.05)
20	150	-1.00 (0.04)	1.01 (0.02)	-2.00 (0.04)	2.01 (0.03)	-0.50 (0.03)	-1.00 (0.03)	-0.50 (0.08)	0.00 (0.04)	0.50 (0.03)
20	250	-1.00 (0.03)	1.00 (0.02)	-2.00 (0.03)	2.00 (0.03)	-0.50 (0.03)	-1.00 (0.04)	-0.50 (0.04)	0.00 (0.04)	0.50 (0.03)
30	20	-1.01 (0.08)	0.99 (0.05)	-2.01 (0.12)	1.99 (0.05)	-0.49 (0.08)	-0.97 (0.13)	-0.57 (0.12)	0.02 (0.10)	0.49 (0.07)
30	80	-1.00 (0.04)	1.00 (0.02)	-2.00 (0.04)	2.00 (0.02)	-0.50 (0.03)	-1.00 (0.04)	-0.50 (0.07)	0.01 (0.05)	0.50 (0.03)
30	150	-1.00 (0.02)	1.00 (0.02)	-1.99 (0.02)	2.00 (0.02)	-0.50 (0.02)	-1.01 (0.03)	-0.51 (0.05)	0.00 (0.05)	0.50 (0.02)
30	250	-1.00 (0.02)	1.00 (0.02)	-2.00 (0.03)	2.00 (0.01)	-0.50 (0.02)	-1.00 (0.03)	-0.51 (0.03)	0.00 (0.02)	0.50 (0.01)

Table 7 reports the means and standard errors in parentheses for the parameters $\beta(s)$, $\kappa(a_i)$, and $\gamma(s)$. Due to space concerns, we do not report the audience cost parameters. In addition, Figures 5 and 6 summarize the results. In these figures, we graph the CMLE's bias and variance, respectively, averaged over the four different sets of coefficients. More specifically, to produce the upper-left graph of Figure 5, we first compute the expected bias of $\hat{\beta}(s)$ for each $s = 1, 2$, and then we averaged these values for each specification of N and T. The upper-left graph of Figure 6 averages the variance of $\hat{\beta}(s)$. The remaining graphs are produced in a similar manner. Most importantly, the bias and the variance of the constrained ML estimator decreases as we increase N and T. This monotonic relationship is especially

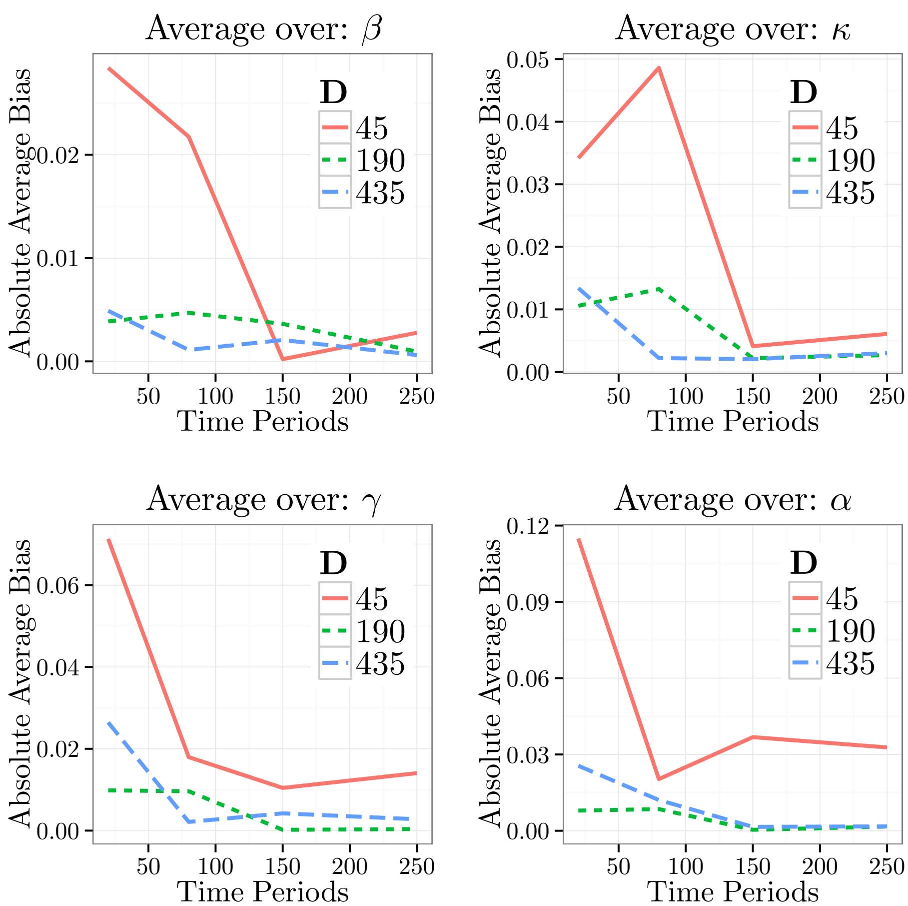

Figure 5: The average bias of the constrained ML estimator by four different types of coefficients as functions of the number of countries N and time periods T. In the analysis, $D = \binom{N}{2}$. Note that the average bias over β is $\frac{1}{4}\sum_{s=2}^{3} |\beta(s) - \mathrm{E}[\hat{\beta}(s)]|_1$.

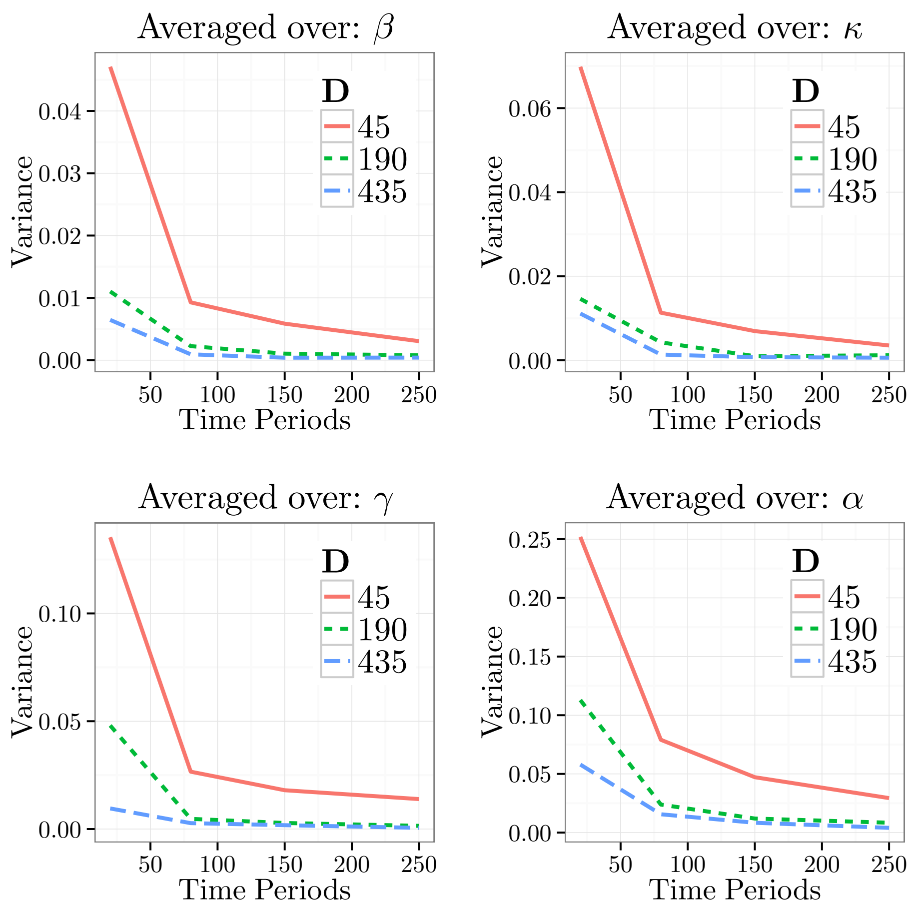

Figure 6: The average variance of the constrained ML estimator by four different types of coefficients. In the analysis, $D = \binom{N}{2}$.

pronounced with the estimator's variance. Even though increasing N means estimating an additional audience cost parameter and more equilibrium constraints, the additional information still attenuates the estimator's bias and variance. With a very small number of countries, i.e. $N = 10$, increasing the number of time periods in the observation may increase the estimator's bias, especially concerning the action-specific cost parameters, $\kappa(a_i)$. However, with a larger number of countries or dyads, this non-monotonicity disappears.

Finally, the experiment provides some quality control on our specific implementation. When $T = 20$, the convergence rate of the procedure is approximately 50%. This is the same across values of N and D. In contrast, when $T > 20$, the convergence rate is 100%, and this is consistent across values of N and D. In addition, Figure 7 graphs the time until convergence. There is exponential growth in computational time as we increase N or the number of ordered dyads. (Recall that adding an unordered dyad means estimating an additional 18 auxiliary

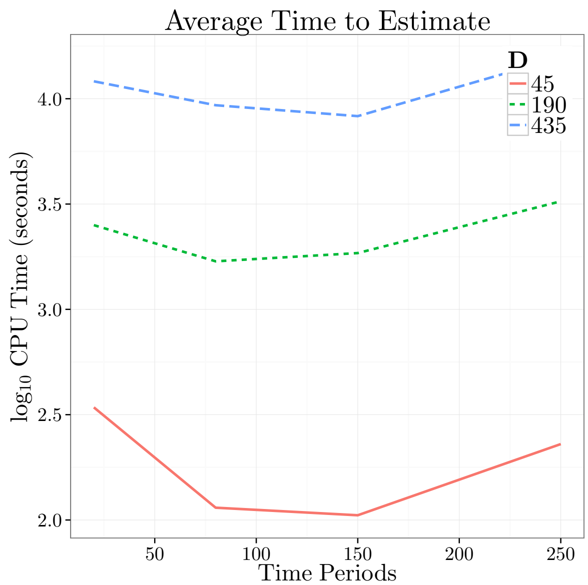

Figure 7: CPU time of the CMLE.

parameters.) Nonetheless, even with 435 dyads the average estimation time in Monte Carlos is approximately three hours.

E Time-Invariant Covariates

Our model and subsequent estimation procedure do not allow for time-varying covariates. More precisely, we have constructed utility functions that are dependent on the state of conflict and actions taken, and we do not incorporate observed variables into the state space. Readers may be concerned that the independent variables included in the model exhibit considerable or even moderate fluctuations over time or that even smaller changes are correlated with observed actions and states. Both of these concerns are unwarranted given our data, however. Namely, we observe very few and very minimal changes in our independent variables, and the changes that do exist are not correlated with the actions chosen and states of conflict.

First, we first examine whether our independent variables change over time. Our independent variables come in two types, and the first type varies by country. These variables include polity2, military personnel per capita, GDP per capita, CINC score and population. For each variable, we compute its means in each country between 1993 and 2007, and then compute country-year deviations from these mean values. The second type varies by dyad, and these variables include trade dependence and whether the dyad has an alliance. We repeat the same process for these variables except we use direct-dyads as observations. Figure 8 displays

Figure 8: Country-Year Deviations from Mean Values of Independent Variables

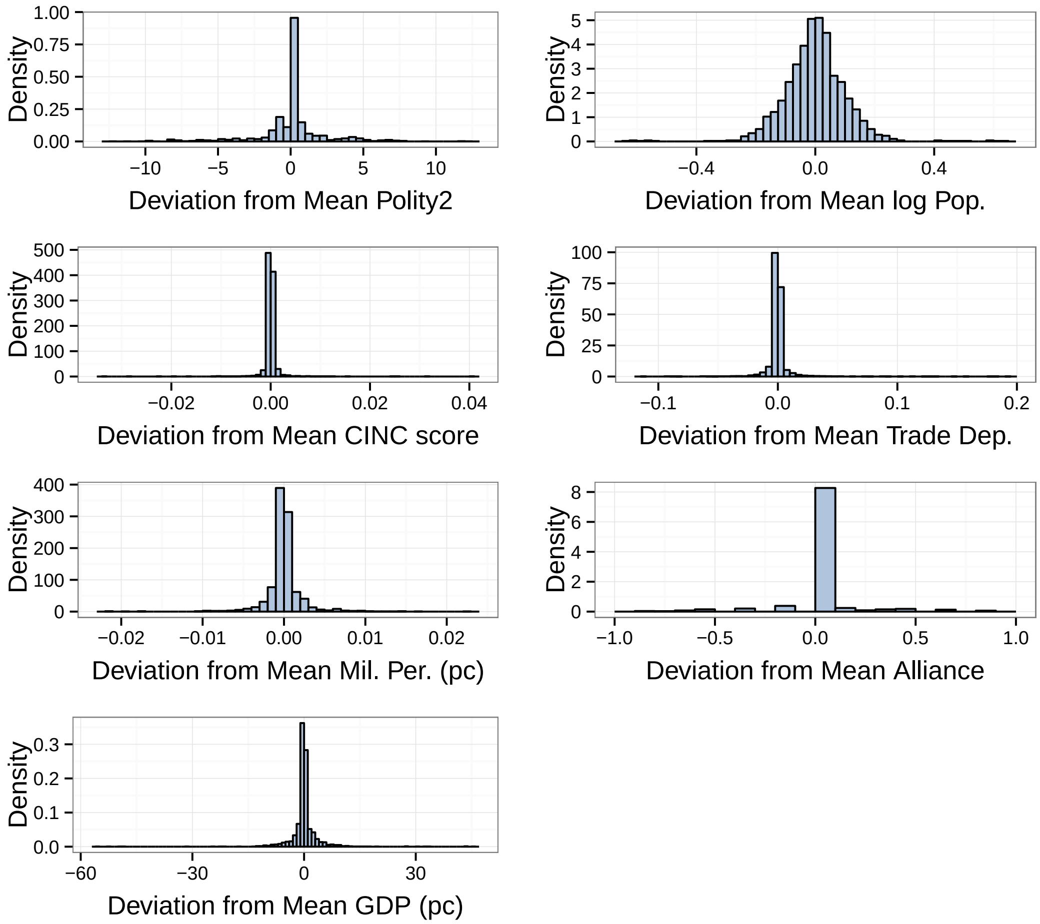

histograms of these deviations for each variable and illustrates that observed deviations from the mean are relatively small across our dataset.

Second, we then attempt to explain our independent variables using observed states and actions in a panel data analysis. Specifically, we regress the country-specific variables on the number of conflict states ($s^t > 1$) in which a country is involved in a given year, the number of hostile actions a country takes in a given year ($a_i^t > 1$), and the number of hostile actions other countries take against it ($a_j > 1$). We also include a lagged dependent variable, lags of these observed actions and states, and country and year fixed effects. We repeated the same process with trade-dependence and alliance presence for the directed dyads in our data except we include dyad and year fixed effects.

Models 1-4 in Table 8 display regression results with country-year observations, and Models

5-6 displays a similar regressions with directed dyad-year observations.[1] The main takeaway should be and that observed actions and states have very little if any influence on our key independent variables. Furthermore, the coefficients on the lagged values are close to 1, which is to be expected if these variables do not change. While one out of the 30 coefficients of interest are significant at the $p < .1$ level, this is not robust to various model and standard-error specifications.

[1]The number of observations vary across models due to missing data in the dependent variable. This is not a problem in the main analysis because we average across years 1993-2007.

Table 8: Predictors of Independent Variables

	Polity2 Model 1	CINC Model 2	Mil. Per. (pc) Model 3	GDP (pc) Model 4	Log Pop. Model 5	Trade Dep. Model 6	Ally Model 7
Dep. Var., lag	0.71^{***}	0.97^{***}	0.74^{***}	0.91^{***}	0.90^{***}	0.92^{***}	0.80^{***}
	(0.03)	(0.05)	(0.04)	(0.02)	(0.02)	(0.08)	(0.02)
Confl. states	0.09	0.00	0.00	−0.04	0.00	0.00	−0.01
	(0.08)	(0.00)	(0.00)	(0.04)	(0.00)	(0.00)	(0.01)
Confl. states, lag	−0.01	0.00	0.00	−0.04	0.00	0.00	0.00
	(0.07)	(0.00)	(0.00)	(0.03)	(0.00)	(0.00)	(0.00)
Conf. Acts against	0.00	0.00	0.00	$0.03^{\dagger}$	$0.00^{\dagger}$	0.00	0.01
	(0.02)	(0.00)	(0.00)	(0.01)	(0.00)	(0.00)	(0.00)
Conf. Acts against, lag	−0.01	0.00	0.00	0.02	0.00	0.00	0.00
	(0.03)	(0.00)	(0.00)	(0.01)	(0.00)	(0.00)	(0.00)
Conf. Acts taken	−0.04	0.00	0.00	0.00	0.00	0.00	0.01
	(0.03)	(0.00)	(0.00)	(0.01)	(0.00)	(0.00)	(0.00)
Conf. Acts taken, lag	−0.02	0.00	0.00	0.00	0.00	0.00	0.00
	(0.02)	(0.00)	(0.00)	(0.01)	(0.00)	(0.00)	(0.00)
N	1736	1750	1750	1730	1750	5004	5012

Notes: $^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$; $^{\dagger}p < 0.1$

Clustered Standard Errors in Parenthesis

Summary statistics of the variables used in the analysis are found in Table 9.

Table 9: Summary of main variables

	Min	Mean	Median	St. Dev.	Max
		Directed Dyadic Variables X_{ij}			
min(Polity 2)	-10.00	-1.03	-1.38	5.30	10.00
log(Cap. Ratio)	-7.49	0.00	0.00	2.39	7.49
sqrt Trade Depend.	0.00	0.07	0.04	0.07	0.46
Alliance	0.00	0.31	0.00	0.46	1.00
log(distance)	1.61	6.65	6.53	1.19	9.23
		Country Specific Variables Z_i			
log(GDP pc +1)	0.14	2.00	1.71	1.21	4.74
log(Mil. Per. pc +1)	0.00	0.01	0.00	0.01	0.05
log(Pop.)	6.11	9.45	9.29	1.47	14.05

F Audience Cost Parameters

This section contains the point estimates on the 125 audience cost parameters we estimated.

Table 10: Audience Cost Estimates

Country	Audience Cost	St. Err.	p
Nepal	-0.87	7.61	0.91
Bangladesh	-1.86	1.12	0.10
Nigeria	-2.18	2.15	0.31
Cambodia	-2.32	6.53	0.72
Lebanon	-2.43	0.77	< 0.01
Benin	-2.46	8.80	0.78
Azerbaijan	-2.64	0.93	< 0.01
Armenia	-2.90	0.93	< 0.01
Pakistan	-3.61	0.90	< 0.01
India	-3.80	0.89	< 0.01
Cameroon	-4.03	2.36	0.09
Djibouti	-4.52	10.85	0.68
Saudi Arabia	-4.75	2.01	0.02

Syria	-5.04	4.65	0.28
Myanmar (Burma)	-5.17	2.17	0.02
Yemen	-5.25	1.95	< 0.01
Iraq	-5.62	0.91	< 0.01
Afghanistan	-7.91	1.85	< 0.01
Albania	-8.95	0.87	< 0.01
Greece	-9.22	0.82	< 0.01
Eritrea	-9.37	5.61	0.09
Turkey	-9.43	0.77	< 0.01
Korea North	-9.54	0.95	< 0.01
Thailand	-9.59	2.38	< 0.01
Korea South	-9.74	1.07	< 0.01
Bosnia	-10.04	2.32	< 0.01
Vietnam	-10.61	5.32	0.05
Georgia	-10.90	3.23	< 0.01
Israel	-11.09	1.71	< 0.01
Japan	-11.33	2.30	< 0.01
Bulgaria	-11.42	5.56	0.04
Latvia	-11.72	4.20	< 0.01
Qatar	-11.97	6.62	0.07
Poland	-12.04	3.61	< 0.01
Jordan	-12.08	5.15	0.02
Yugoslavia	-12.83	1.10	< 0.01
Congo Kinshasa	-13.37	1.99	< 0.01
Croatia	-13.39	1.90	< 0.01
Chile	-13.52	4.88	< 0.01
South Africa	-13.59	17.62	0.44
Romania	-14.08	3.49	< 0.01
Tajikistan	-14.24	4.85	< 0.01
Iran	-14.26	2.38	< 0.01
Ukraine	-14.29	2.20	< 0.01
Uganda	-14.50	2.07	< 0.01
Zimbabwe	-14.60	5.95	0.01
Kuwait	-14.62	2.52	< 0.01
Sudan	-14.75	1.85	< 0.01

Mozambique	-14.79	9.15	0.11
Cuba	-14.80	6.62	0.03
Lithuania	-14.88	3.90	< 0.01
Ethiopia	-14.90	2.00	< 0.01
Malaysia	-14.99	1.78	< 0.01
Uzbekistan	-15.07	4.17	< 0.01
Morocco	-15.14	5.36	< 0.01
Moldova	-15.15	1.72	< 0.01
Central African Republic	-15.17	3.07	< 0.01
Chad	-15.26	2.30	< 0.01
China	-15.29	1.88	< 0.01
Kenya	-15.40	3.31	< 0.01
Australia	-15.41	5.01	< 0.01
Germany	-15.58	1.82	< 0.01
Namibia	-15.73	17.58	0.37
Mongolia	-15.79	2.36	< 0.01
Egypt	-15.79	2.85	< 0.01
Venezuela	-15.80	1.68	< 0.01
Burundi	-16.10	3.28	< 0.01
Mauritania	-16.12	5.73	< 0.01
Singapore	-16.20	3.96	< 0.01
Russia	-16.22	1.10	< 0.01
Belgium	-16.29	3.95	< 0.01
United Kingdom	-16.30	1.95	< 0.01
El Salvador	-16.33	2.99	< 0.01
Guyana	-16.35	4.29	< 0.01
Papua New Guinea	-16.42	2.24	< 0.01
Portugal	-16.43	4.13	< 0.01
Zambia	-16.45	2.66	< 0.01
Indonesia	-16.46	1.49	< 0.01
Netherlands	-16.47	3.80	< 0.01
Slovenia	-16.48	2.43	< 0.01
Nicaragua	-16.55	2.30	< 0.01
Canada	-16.76	2.81	< 0.01
Bhutan	-16.89	4.10	< 0.01

United States	-16.92	1.12	< 0.01
Philippines	-17.00	3.26	< 0.01
Angola	-17.07	2.35	< 0.01
Kyrgyzstan	-17.13	3.86	< 0.01
Haiti	-17.22	2.83	< 0.01
Spain	-17.23	3.56	< 0.01
Hungary	-17.23	2.12	< 0.01
Liberia	-17.27	1.62	< 0.01
Costa Rica	-17.55	3.98	< 0.01
Libya	-17.58	3.02	< 0.01
Botswana	-17.69	12.60	0.16
Belarus	-17.81	3.42	< 0.01
Tanzania	-17.81	1.84	< 0.01
Rwanda	-17.89	2.46	< 0.01
France	-17.92	1.06	< 0.01
Algeria	-17.94	2.36	< 0.01
Italy	-18.15	2.04	< 0.01
Sierra Leone	-18.18	1.48	< 0.01
Ghana	-18.31	13.35	0.17
Gambia	-18.47	4.56	< 0.01
Congo Brazzaville	-18.52	3.31	< 0.01
Argentina	-18.54	1.57	< 0.01
Mali	-18.63	1.91	< 0.01
Togo	-18.66	9.01	0.04
Niger	-18.71	2.13	< 0.01
Denmark	-18.75	3.22	< 0.01
Turkmenistan	-18.84	4.00	< 0.01
Swaziland	-18.99	3.76	< 0.01
Ivory Coast	-19.28	1.38	< 0.01
Peru	-19.34	4.44	< 0.01
Lesotho	-19.41	16.71	0.25
Ecuador	-19.42	2.85	< 0.01
Brazil	-19.43	3.01	< 0.01
Guinea	-19.69	1.32	< 0.01
Colombia	-19.88	3.08	< 0.01

Somalia	-19.91	6.04	< 0.01
Senegal	-19.92	1.52	< 0.01
Honduras	-19.94	2.17	< 0.01
Suriname	-20.22	3.39	< 0.01
Norway	-21.47	1.34	< 0.01
Dominican Rep	-22.34	2.37	< 0.01
Cyprus	-33.87	3.02	< 0.01

G Audience Costs in Autocratic Regimes

Table 11 reports the estimates from two regressions where we regress our audience cost estimates the on regime types from Weeks (2008) and Weeks (2012), respectively.

H Substantive Effects

In this Appendix, we detail how to compute substantive effects in the model using a simple homotopy predictor-corrector method. We use this method to examine the effect of changing the parameters from $\hat{\theta}$, which are estimated when solving Eq. 4, to $\tilde{\theta}$, which are chosen by the researcher, on the estimated equilibria $\hat{\mathbf{v}}$. Appendix B discusses how to use marginal effects, or instantaneous effects, for comparative statics, but in this section, we examine how equilibria change when we move or vary the data or parameters by substantial degrees.

Because there are multiple equilibria, we cannot just vary the parameters, compute a new equilibrium, and compare choice probabilities under the old and new parameters values. Doing so would not guarantee that the new equilibrium bears any resemblance to the estimated one. For example, it is possible to uncover drastic differences in choice probabilities when selecting among different equilibria although the data may not change. Likewise, we do not consider substantive effects in an average dyad because there is no information on what equilibrium such a dyad would play. To alleviate these concerns, we focus on two substantively interesting dyads, Lebanon-Israel and North-South Korea, and implement a method from Aguirregabiria (2012) that maps equilibria as functions of data or parameters.

The method proceeds as follows. First, define the vectors of data x^k and z^k as $(x_{i^k j^k}, x_{j^k i^k})$ and (z_{i^k}, z_{j^k}), respectively. Then, comparative statics refer to how the equilibrium v^k changes as we vary the parameters and data from (θ, x^k, z^k) to $(\tilde{\theta}, \tilde{x}^k, \tilde{z}^k)$. For example, how would

Table 11: Regime classification and audience costs

	WLS: Audience Costs	
	Weeks (2008)	Weeks (2012)
Personalist	4.89^{**}	
	(1.67)	
Single-party	-1.11	
	(2.28)	
Military	2.39	
	(3.97)	
Machine		-1.43
		(1.63)
Junta		1.93
		(2.11)
Boss		3.94^{*}
		(1.58)
Strongman		1.15
		(2.00)
Other non-democracy	1.69	2.13^{*}
	(1.10)	(1.01)
Population	0.34	0.36
	(0.36)	(0.28)
Constant	-18.35^{**}	-18.32^{**}
	(3.80)	(2.85)
R^2	0.10	0.09
N	89	125

Notes: $^{**}p < 0.01$; $^{*}p < 0.05$
Standard errors in parenthesis. Observations are weighted by the number of dyads in which each country appears.

the equilibrium, and potentially the probability of war, between Russia and the United States change if Russia were to become a democracy holding all other variables constant and given our estimate θ? Likewise, how would the equilibrium between Lebanon and Israel change if we were to increase Lebanon's audience cost parameter? When the Jacobian of $\Phi^k(\cdot; \theta \mid x^k, z^k)$ (with respect to v^k) is not vanishing at equilibrium v^k, then small changes in the data and parameters result in small changes to the equilibrium by the implicit function theorem. This condition can be verified at the estimated parameters and equilibria. Nonetheless, we also need a behavioral assumption: If we vary the data continuously, countries play the equilibrium corresponding to the smooth change in the original equilibrium.

Unfortunately, even with this assumption, we cannot simply solve the system of equations

$\Phi^k(v^k;\tilde{\theta}|\tilde{x}^k,\tilde{z}^k)-v^k=0$ for $\tilde{v}^k$ because multiple equilibria potentially exist, and it is therefore possible to not even change the data but uncover a very drastic change in behavior by selecting among different equilibria. To alleviate this problem, we implement a homotopy method proposed in Aguirregabiria (2012) to trace equilibria from v^k to $\tilde{v}^k$. More specifically, for small changes in the parameter vector and the data, we first approximate changes in v^k using the implicit function theorem and linear approximation as in Aguirregabiria (2012). Next, we use these approximations as starting values in a Newton or quasi-Newton method that computes new equilibria resulting from the small changes in the data. Finally, we repeat this procedure until reaching the final vector of data. A minor difference between this routine and the specifics discussed in Aguirregabiria (2012) is that ours requires the computation of equilibria at each step. While this does increase the computational burden of the procedure, it is feasible given our small state space and will return an equilibrium upon convergence.

Algorithm 1: COMPARATIVE STATICS (CS) using a homotopy

Input: A coefficient vector θ, control variables for dyad k x^k and z^k, an equilibrium v^k, i.e., $\Phi^k(v^k \mid U(\theta,x^k,z^k))=v^k$, new values for the parameter vector $\tilde{\theta}$ and control variables $\tilde{x}^k$ and $\tilde{z}^k$, and a tuning parameter $n\in\mathbb{N}$. To pass to the BROYDEN solver, a convergence criterion $\varepsilon>0$, and a number of maximum iterations $m\in\mathbb{N}$.

Output: An equilibrium $\tilde{v}^k$ under new parameters $\tilde{\theta}$ and data $\tilde{x}^k$ and $\tilde{z}^k$.

$U_{\texttt{old}} \leftarrow U(\theta,x^k,z^k)$
$\tilde{v}^k \leftarrow v^k$
for $i\leftarrow 1$ ***to*** n **do**
    $\lambda \leftarrow \frac{i}{n}$
    $U_{\texttt{new}} \leftarrow (1-\lambda)U(\theta,x^k,z^k)+\lambda U(\tilde{\theta},\tilde{x}^k,\tilde{z}^k)$
    $\texttt{slope} \leftarrow -\left(D_U\Phi^k(\tilde{v}^k \mid U_{\texttt{old}})\right)'\left(D_{\tilde{v}^k}\Phi^k(\tilde{v}^k \mid U_{\texttt{old}})\right)^{-1}$
    $\texttt{start} \leftarrow \tilde{v}^k+[U_{\texttt{new}}-U_{\texttt{old}}]\,\texttt{slope}$
    $(\tilde{v}^k,\texttt{success}) \leftarrow \text{BROYDEN}(\texttt{start},\Phi^k(v \mid U_{\texttt{new}})-v,\varepsilon,m)$
    if *success* **then**
        $U_{\texttt{old}} \leftarrow U_{\texttt{new}}$
    else
        $\tilde{v}^k \leftarrow$ "Warning: Convergence Problems."
        break
return $\tilde{v}^k$

Algorithm 1 presents the specifics of the procedure for reference and is an implementation of predictor-corrector method. Let $U(\theta,x^k,z^k)$ denote the vector of actor-state-action-profile utilities given a parameter vector θ and data x^k and z^k. Likewise, let $\Phi^k(v \mid U)$ denote the equilibrium conditions from Eq. 6 with utilities U. In line 6, we must compute $D_U\Phi$ and $D_v\Phi$. These can be computed using automatic differentiation or are relatively straightforward to due by hand. In line 7, we use linear interpolation to predict how the equilibrium changes. In line 8, we call a Broyden solver, which returns a pair consisting of a solution and an indicator of a successful convergence. In our experiments, we save the Broyden call in each iteration to produce a continuous representation of the equilibrium, which we use to graphically verify that the output has indeed continuously traced the equilibrium.

I Additional Results Robustness Checks

Table 12 reports the coefficient estimates associated with the independent variables that determine a country's state- and action-specific payoffs. Notice that the table does not report five models, but rather reports the output from one model across five columns. The first two columns describes the estimates concerning state-specific payoffs that country i gets from its current state of affairs with country j. As with the standard multinomial logit, these estimates are interpreted as the relative increase or decrease in utility compared to being in the peaceful state. Columns three and four show our estimates of the country-specific costs to country i from taking action a_i. The last column contains the estimates of the structural parameters $\gamma(s)$ which describe whether conflict deters or spirals in state s.

We report the results of several robustness checks. The first robustness check restricts the 125 audience cost parameters to a single estimated value, these results are reported in Table 13. As we can see the results are largely unchanged from the main model.

In Table 14 we add a dummy that records if the pairs of countries have a formal alliance in more than half of the time periods in our data. The results roughly match the results from the main part of the paper, the only noticeable differences are the dyadic variable in crisis. In this model we see that the democratic peace appears to turn up in both crisis and war and allies prefer to be in peace.

The next check is in Table 15, and it uses our original variables plus a control for distance. Distance is measured as the logged distance between each country's capital. We find no real effect of distance, but the rest of the results match the previous model.

The next model, Table 16, includes all of our original variables plus both alliance and distance. As before we find no real effect of distance, but we do find the intuitive result that allies prefer peace to both crisis and war. In this, and the last two models the liberal peace of trade and democracy are definitely present in the war state, but there is some indication

Table 12: Structural estimates, omitting α_i.

	β(Crisis)	β(War)	κ(Crisis)	κ(War)	γ
Joint Democracy	0.00	-0.03^*			
	(0.01)	(0.01)			
Cap. Ratio	0.00	-0.05^*			
	(0.02)	(0.02)			
Trade Depend.	0.19	-3.06^*			
	(0.49)	(0.79)			
GDP pc			0.13^*	-0.15^*	
			(0.03)	(0.03)	
Mil. Per. pc			7.54	-8.15	
			(4.39)	(5.15)	
γ(Peace)					-48.24^*
					(2.65)
γ(Crisis)					9.32^*
					(0.60)
γ(War)					13.84^*
					(0.59)
Constant	19.23^*	19.08^*	-21.18^*	-20.96^*	
	(1.10)	(1.40)	(1.02)	(1.27)	
Log L			-45.50		
Dyads			179		

Notes: $^*p < 0.05$; Standard errors in parenthesis

Table 13: Structural estimates, single α model.

	β(Crisis)	β(Conflict)	κ(Crisis)	κ(Conflict)	γ
Joint Democracy	0.01^{*}	-0.03^{**}			
	(0.00)	(0.01)			
Cap. Ratio	0.03^{**}	-0.02^{**}			
	(0.01)	(0.01)			
Trade Depend.	-0.19	0.10			
	(0.17)	(0.10)			
GDP pc			0.01	-0.01	
			(0.01)	(0.00)	
Mil. Per. pc			0.47	-0.26	
			(1.09)	(0.72)	
γ(Peace)					-33.96^{**}
					(2.87)
γ(Crisis)					7.64^{**}
					(0.65)
γ(Conflict)					10.72^{**}
					(0.99)
Constant	8.17^{**}	8.48^{**}	-10.87^{**}	-12.42^{**}	
	(1.09)	(2.14)	(1.00)	(2.18)	
Log L			-48.92		
Dyads			179		

Notes: $^{**}p < 0.01$; $^{*}p < 0.05$; $^{\dagger}p < 0.1$
Standard errors in parenthesis

from these first three models that this effect keeps democracies and trading partners out of crisis as well.

The fourth check, shown in Table 17, looks similar to the last few. We continue to add the dyadic variables of distance and alliance, but now we have removed all the country specific cost variables except for the constants. These results largely match the last few tables.

Our final robustness check is in Table 18. It is the fullest model we consider as it includes all the dyadic variable mentioned thus far, the same cost variables from the original model, plus an additional cost variable that is logged population. These results look the most similar to our original regression results in the main text. Note that the effect of democracy on preferring crisis to peace is now statistically significant, but the sign is now positive. This is additional evidence for our conjecture above that there may not be a real pacifying effect of democracies when it comes to crisis level interaction. The only major change that appears in this model is that γ(Peace) is now positive, however, this result shows up in none of the

Table 14: Adding alliance, still omitting α_i.

	β(Crisis)	β(Conflict)	κ(Crisis)	κ(Conflict)	γ
Joint Democracy	-0.02^{**}	-0.05^{**}			
	(0.01)	(0.01)			
Cap. Ratio	0.04^{*}	-0.09^{**}			
	(0.02)	(0.02)			
Trade Depend.	-1.01	-6.12^{**}			
	(0.65)	(0.81)			
Alliance	$-0.11^{\dagger}$	-0.33^{**}			
	(0.06)	(0.11)			
GDP pc			0.09^{**}	-0.13^{**}	
			(0.02)	(0.03)	
Mil. Per. pc			3.75	-4.39	
			(3.91)	(4.66)	
γ(Peace)					-39.14^{**}
					(2.38)
γ(Crisis)					11.07^{**}
					(0.62)
γ(Conflict)					15.18^{**}
					(0.61)
Constant	19.11^{**}	18.28^{**}	-21.22^{**}	-20.37^{**}	
	(1.00)	(1.20)	(0.92)	(1.10)	
Log L			-46.12		
Dyads			179		

Notes: $^{**}p < 0.01$; $^{*}p < 0.05$; $^{\dagger}p < 0.1$
Standard errors in parenthesis

other models.

In order to compare the 125 α_i parameters across the original model and the 5 robustness checks we conduct some simple correlation analysis. When we compare the pair-wise correlations across the 6 models, the median correlation is about 0.37 a reasonable but not phenomenally high correlation. On examination we noticed that the audience costs from the largest model, the one in Table 18, appear to exhibit some major separation. When we drop that model from the audience costs analysis we find that the parameters have a median correlation across models of about 0.82. This is a much stronger relationship that gives us confidence that the model parameters exhibit only mild change across specification.

Table 15: Adding distance, still omitting α_i.

	β(Crisis)	β(Conflict)	κ(Crisis)	κ(Conflict)	γ
Joint Democracy	-0.01^{*}	-0.03^{**}			
	(0.01)	(0.01)			
Cap. Ratio	0.01	-0.10^{**}			
	(0.02)	(0.02)			
Trade Depend.	-2.26^{**}	-6.74^{**}			
	(0.64)	(0.84)			
Distance	-0.02	0.05			
	(0.03)	(0.04)			
GDP pc			0.10^{**}	-0.15^{**}	
			(0.03)	(0.04)	
Mil. Per. pc			3.85	-6.98	
			(4.06)	(5.84)	
γ(Peace)					-42.39^{**}
					(2.45)
γ(Crisis)					11.22^{**}
					(0.63)
γ(Conflict)					14.75^{**}
					(0.60)
Constant	19.05^{**}	16.54^{**}	-20.98^{**}	-19.07^{**}	
	(1.05)	(1.28)	(0.94)	(1.14)	
Log L			-46.18		
dyads			179		

Notes: $^{**}p < 0.01$; $^{*}p < 0.05$; $^{\dagger}p < 0.1$
Standard errors in parenthesis

J Additional Substantive Effects

In this section, we use the estimated coefficients and equilibria to conduct counterfactual experiments on the remaining parameters. Throughout this section, we expand our focus to include four theoretically interesting dyads: the United States and Iran, Cyprus and Turkey, Lebanon and Israel, and North and South Korea. These experiments are similar to our analysis of audience costs. In the first three experiments, we vary the dyad-specific variables x_{ij} to $\tilde{x}_{ij}$, trace the equilibrium (as before), and compute the new invariant distribution. In Experiment 1, we increase every country's trade dependence by one-half a standard deviation. In Experiment 2, we increase the dyad's minimum polity score to 10. In Experiment 3, we set the capability ratio of the dyad to 1. We choose these values because they ensure that the equilibria do not vanish as we change the stage utility function and they are consistent with

Table 16: Structural estimates, restricted κ, still omitting α_i.

	β(Crisis)	β(Conflict)	κ(Crisis)	κ(Conflict)	γ
Joint Democracy	-0.01^{*}	-0.05^{**}			
	(0.01)	(0.01)			
Cap. Ratio	0.07^{**}	-0.10^{**}			
	(0.02)	(0.02)			
Trade Depend.	-0.66	-5.93^{**}			
	(0.64)	(0.85)			
Alliance	-0.22^{**}	-0.41^{**}			
	(0.08)	(0.12)			
Distance	0.01	-0.04			
	(0.02)	(0.04)			
γ(Peace)					-37.63^{**}
					(2.26)
γ(Crisis)					11.10^{**}
					(0.61)
γ(Conflict)					15.46^{**}
					(0.61)
Constant	19.14^{**}	18.87^{**}	-21.03^{**}	-20.98^{**}	
	(0.99)	(1.22)	(0.90)	(1.08)	
Log L			-46.27		
Dyads			179		

Notes: $^{**}p < 0.01$; $^{*}p < 0.05$; $^{\dagger}p < 0.1$
Standard errors in parenthesis

the domain of the independent variables.

Table 19 reports the results from the first class of the experiments. In the far-right columns we report the old and new invariant distributions, respectively. Here we see that increasing Lebanon and Israel's trade dependence by one-half of a standard deviation leads to an approximate 10% increase in the probability of peace. Given the small change made to trade dependence this effect is quite large especially in light of substantive effects described in other conflict work. In contrast, a similar change to the dyad including North and South Korea leads to a 2.5% decrease in the probability of peace. Likewise, increasing the minimum democracy score in the Lebanon–Israel dyad increases peace by more than 10%, but a similar change with North and South Korea decreases the probability of peace by 6% although the latter change is not statistically significant. To better illustrate the effect of minimum democracy in the Israel-Lebanon dyad, Figure 9 graphs the probability of peace and war as we vary the dyad's

Table 17: Structural estimates adding distance and alliance, still omitting α_i.

	β(Crisis)	β(Conflict)	κ(Crisis)	κ(Conflict)	γ
Joint Democracy	-0.02^{**}	-0.05^{**}			
	(0.01)	(0.01)			
Cap. Ratio	0.04^{*}	-0.10^{**}			
	(0.02)	(0.02)			
Trade Depend.	-0.84	-6.10^{**}			
	(0.65)	(0.82)			
Alliance	-0.14^{*}	-0.35^{**}			
	(0.07)	(0.11)			
Distance	-0.02	-0.02			
	(0.03)	(0.04)			
GDP pc			0.10^{**}	-0.13^{**}	
			(0.03)	(0.04)	
Mil. Per. pc			2.85	-4.75	
			(3.93)	(4.69)	
γ(Peace)					-38.66^{**}
					(2.40)
γ(Crisis)					11.12^{**}
					(0.62)
γ(Conflict)					15.29^{**}
					(0.61)
Constant	19.40^{**}	18.58^{**}	-21.35^{**}	-20.55^{**}	
	(1.02)	(1.23)	(0.92)	(1.09)	
Log L			-46.11		
Dyads			179		

Notes: $^{**}p < 0.01$; $^{*}p < 0.05$; $^{\dagger}p < 0.1$
Standard errors in parenthesis

minimum polity score from the smallest to the largest possible values.[2] As minimum polity varies from –10 to 10, the probability of peace increases from 50% to 80%, approximately, and the probability of war decreases from 45% to 25%, approximately. In contrast, there are no such noticeable changes in any dyad when looking at the ratio of military capabilities.

In our final experiment, we reduce the degree to which conflict is mutually reinforcing in the crisis and war states by decreasing both $\gamma(2)$ and $\gamma(3)$ by 1.0. Such a situation reflects an increase in the mutual destructiveness of war which can theoretically be imposed by an international organization. The procedure for this experiment is exactly the same as our

[2] Of the four dyads considered, this is the only one that allows such large reductions in the polity score before the estimated equilibrium vanishes.

Table 18: Structural estimates including distance, alliance, and population, still omitting α_i.

	β(Crisis)	β(Conflict)	κ(Crisis)	κ(Conflict)	γ
Joint Democracy	0.04^{**}	-0.08^{**}			
	(0.01)	(0.01)			
Cap. Ratio	-0.12^{**}	-0.01			
	(0.03)	(0.04)			
Trade Depend.	-2.68^{**}	-5.06^{**}			
	(0.83)	(0.91)			
Alliance	0.02	0.93^{**}			
	(0.10)	(0.14)			
Distance	-0.31^{**}	-0.11			
	(0.04)	(0.07)			
GDP pc			0.21^{**}	-0.11^{*}	
			(0.03)	(0.05)	
Mil. Per. pc			31.76^{**}	28.85^{**}	
			(4.21)	(5.08)	
Population			0.27^{**}	0.07	
			(0.03)	(0.05)	
γ(Peace)					21.74^{**}
					(0.95)
γ(Crisis)					3.96^{**}
					(0.17)
γ(Conflict)					13.35^{**}
					(0.52)
Constant	10.34^{**}	23.11^{**}	-14.76^{**}	-26.25^{**}	
	(0.60)	(1.24)	(0.65)	(1.20)	
Log L			-45.68		
Dyads			179		

Notes: $^{**}p < 0.01$; $^{*}p < 0.05$; $^{\dagger}p < 0.1$
Standard errors in parenthesis

audience cost experiment, and the results are found in Table 20.[3]

In Experiment 4, decreasing the destructiveness of mutual conflict leads to a 10% increase in the amount of time spent in peace in the Korean dyad. This makes sense because in this experiment we have decreased the mutual profit to be gained from continuing to a fight. In contrast, we observe that despite making mutual continuance of conflict less profitable, the probability of peace between Cyprus and Turkey decreases by about 12%. This can occur because as we make mutual conflict more destructive, one country begins to attack less which

[3]In Experiment 4, we could not continuously trace the equilibrium between Lebanon and Israel.

Table 19: Counterfactual Experiments: Dyadic Variables

	Dyad	Invariant Dist. In Data	Invariant Dist. Counterfactual
Experiment 1 (Trade)	United States and Iran	(94.75, 3.81, 1.44)	(95.14, 3.77, 1.09)
	Cyprus and Turkey	(88.69, 9.19, 2.12)	(86.80, 10.78, 2.41)
	Lebanon and Israel	(54.86, 6.44, 38.71)	(64.78, 4.91, 30.31)*
	North and South Korea	(86.08, 9.33, 4.59)	(83.79, 12.14, 4.07)
Experiment 2 (Democracy)	United States and Iran	(94.75, 3.81, 1.44)	(95.66, 3.41, 0.94)
	Cyprus and Turkey	(88.69, 9.19, 2.12)	(87.88, 9.84, 2.27)
	Lebanon and Israel	(54.86, 6.44, 38.71)	(66.45, 4.90, 28.65)**
	North and South Korea	(86.08, 9.33, 4.59)	(80.89, 13.34, 5.77)
Experiment 3 (Capability Ratio)	United States and Iran	(94.75, 3.81, 1.44)	(94.60, 3.83, 1.57)
	Cyprus and Turkey	(88.69, 9.19, 2.12)	(87.21, 10.23, 2.57)
	Lebanon and Israel	(54.86, 6.44, 38.71)	(55.59, 6.05, 38.35)
	North and South Korea	(86.08, 9.33, 4.59)	(86.08, 9.33, 4.59)

Notes: $^{**}p < 0.01$; $^{*}p < 0.05$; $^{\dagger}p < 0.1$

Table 20: Counterfactual Experiments: γ

	Dyad	Invariant Dist. In Data	Invariant Dist. Counterfactual
Experiment 4 $\gamma(s)$	United States and Iran	(94.75, 3.81, 1.44)	(94.89, 3.71, 1.40)
	Cyprus and Turkey	(88.69, 9.19, 2.12)	(76.31, 17.66, 6.03)**
	Lebanon and Israel	(54.86, 6.44, 38.71)	NA
	North and South Korea	(86.08, 9.33, 4.59)	(88.31, 3.67, 8.02)**

Notes: $^{**}p < 0.01$; $^{*}p < 0.05$; $^{\dagger}p < 0.1$

means the other can attack with more impunity. Such a comparative static is common in games of chicken when one player finds standing firm less and less attractive.

K Model Fit

This Appendix contains several model fit exercises. First, we examine aggregate predictions using our estimated equilibria. Aggregating across all dyads, our model predicts 93.9% of states should be peace and 2.86% of states should be crisis.[4] This compares to rates of 95.4% and 2.52% observed in the data, respectively. Likewise, Table 21 illustrates our predictions concerning the nine different transitions and should be compared to Table 1, which reports

[4]These are computed by averaging the equilibrium invariant distributions across dyads.

Figure 9: The effect of minimum democracy.

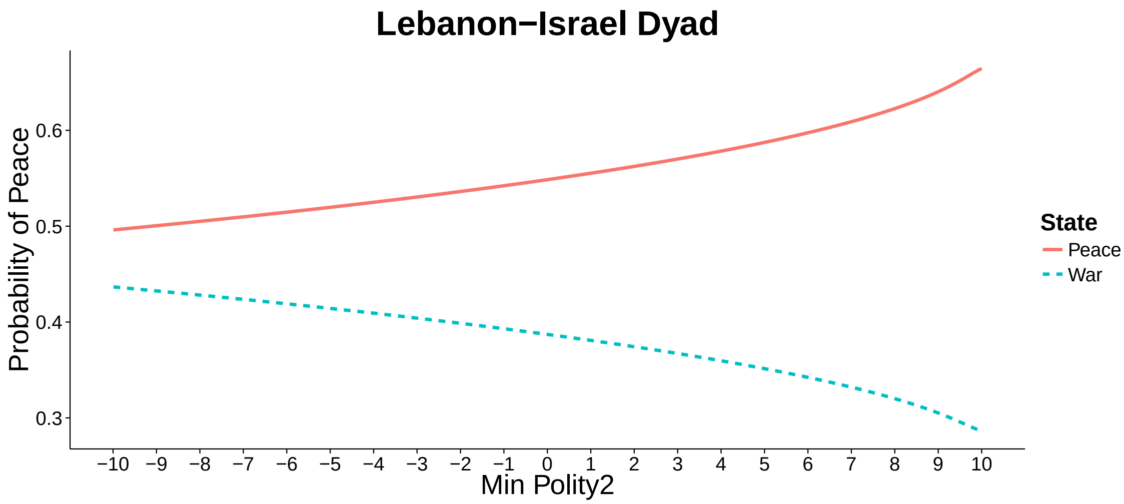

transitions observed in the data. Again, the model's predictions match the data. A notable exception occurs in the war state, where the model over predicts the percent of war transitions conditional on the dyads being in war by 8% and under predicts the percent of war-to-crisis transitions by 6%. Nonetheless, the model's predictions essentially mimics their real-world counterparts.

Next, we examine the degree to which our estimated equilibria fit the data *within each dyad*, and we use two non-strategic models as our benchmarks. Although our structural model is relatively parsimonious, the results suggest that it not only fits the data quite well but also fits the data at least as well if not better than the more standard models.

For the non-strategic models, we use multinomial logits to estimate the probability that a country i^k chooses action a_i in dyad k and include observed covariates $x_{i^k j^k}$, z_{i^k}, and z_{j^k} as predictors, where an observation is the ordered-dyad month. The first logit does not condition this probability on the observed state, but the second one adds this flexibility. Similar approaches are the standard in MIDs analyses (Buhaug and Gleditsch 2008; Huth and Allee 2002; Reiter and Stam 2003). To determine the degree to which these models fit the data, we use several power divergence statistics, which include a χ^2 statistic and a likelihood ratio or G^2 statistic. In addition, we follow the recommendation of Read and Cressie (1988) and include divergence statistics with parameter $\lambda = \frac{2}{3}$ and $\lambda = 5$. The former is robust across several types of null models, and the latter is particularly suited for ones that overwhelmingly predict a single category, for example, peace in our application.

Table 22 presents the results across the three different models (rows) on three types of observations (columns). In columns 3–5, the corresponding cells report the percent of dyads

Table 21: Predicted aggregate transitions.

Transition	Predicted transitions	Predicted within state
Peace → Peace	90.1%	95.9%
Peace → Crisis	2.27%	2.42%
Peace → War	1.54%	1.64%
Crisis → Peace	2.12%	74.2%
Crisis → Crisis	0.45%	15.7%
Crisis → War	0.29%	10.1%
War → Peace	1.69%	51.6%
War → Crisis	0.14%	4.29%
War → War	1.44%	44.1%

Caption: The middle column displays the probability distribution over expected transitions and the far-right column presents the conditional distribution in each state, both of which are computed from equilibrium estimates $\hat{\mathbf{v}}$. This table matches Table 1, which records the percent of transitions observed in the data.

Table 22: Divergence statistics and comparison to non-strategic models

		Observed states			Conditional transitions			Transition matrix		
		$p < .1$	$p < .05$	$p < .01$	$p < .1$	$p < .05$	$p < .01$	$p < .1$	$p < .05$	$p < .01$
Estimated equilibria	χ^2	46.93	28.49	19.55	23.85	18.35	12.39	25.70	20.67	18.99
	G^2	72.63	63.69	34.08	37.61	30.05	11.47	39.11	24.58	15.64
	$\lambda=\frac{2}{3}$	59.78	39.11	18.99	29.13	18.58	11.70	24.02	22.91	17.32
	$\lambda=5$	22.91	19.55	17.32	24.31	22.94	21.56	32.40	31.84	30.73
Multinomial model	χ^2	78.21	55.87	27.93	51.83	40.60	24.31	49.72	42.46	34.64
	G^2	87.15	83.24	65.92	50.00	45.41	34.63	81.56	71.51	41.34
	$\lambda=\frac{2}{3}$	82.12	70.39	31.84	52.06	42.20	24.77	56.98	44.69	34.64
	$\lambda=5$	31.28	26.26	22.35	33.94	31.42	29.13	41.34	40.22	38.55
Multinomial cond. on state	χ^2	85.47	76.54	41.34	21.10	16.06	9.40	30.73	23.46	15.64
	G^2	91.06	87.71	75.42	37.84	28.67	11.47	49.16	31.28	17.32
	$\lambda=\frac{2}{3}$	88.27	79.89	49.16	26.15	16.97	9.17	31.28	22.35	15.08
	$\lambda=5$	38.55	29.61	22.35	27.75	24.54	22.02	38.55	37.43	34.64

Caption: The results of four power divergence tests for three different models, which include the estimated equilibria (rows 1–4), a multinomial logit predicting whether an ordered-dyad engages in peace, crisis and war (rows 5–8), and a similar multinomial logit conditioning on the observed state (rows 9–12). Numbers in cells correspond to percent of observations in which we reject the null that the respective model generates the observed distribution of states (columns 3–5), distribution of transitions conditional on state (columns 6–8), and transition matrix (columns 9–11). Smaller numbers indicate better fit.

in which we reject the null hypothesis that the model's invariant distribution generates the observed states for three standard p-vales. Essentially, this is a joint hypothesis test where we reject the null if the equilibrium does not generate the states or the path of play has not converged to its invariant distribution.[5] Thus, smaller numbers indicate better fit, and the tests suggest that the equilibria explain 70%–80% of dyads quite well, which is better than both multinomial models. Furthermore, this suggests that the data generating process has converged to its stationary distribution. Upon further inspection, there does not appear to be an obvious pattern as to why our model performs poorly on certain dyads. For example, the model predicts 99 peaceful states in the Lebanon-Israel dyad, but the data contain 52. In other words, we over predict peace. In contrast, the model under predicts peace in the India-Pakistan dyad, with 65 peaceful states predicted compared to 110 in the data.

In columns 6–8, we examine the degree to which expected transitions match the distribution of observed transitions within each dyad and state. Here, the cells report the percentage of dyad-state pairs in which we reject the null hypothesis that the model under consideration generates the observed transitions.[6] The results suggest that the estimated equilibria explain between 80%–90% of observations, indicating strong fit. Furthermore, the estimated equilibria fit the data better than the standard multinomial model and as good as the multinomial model conditioning on state. As in the previous test, the estimated equilibria perform poorly for a variety of reasons. For example, it over predicts peace transitions by approximately 20 observations in the UK-Iraq-war pairs, but under predicts peaceful transitions in the India-Pakistan-war pair by a similar amount. In a similar vein, columns 9–11 report the percentage of dyads in which we reject the null hypothesis that the observed transition counts were generated from the predicted Markov-transition matrix (Billingsley 1961). Overall, the tests suggest our equilibria fit the data well, where we reject the null in 20%–30% of dyads. As before, the equilibria appear to fit the data better than the standard models.

In addition, we compute our equilibria's Kullback-Leibler distance using the three different data types in Table 23. And we compare these distances to those from the two estimated multinomial logits described above. Notice that we avoid traditional model tests in Vuong (1989) and Clarke (2007) because the data's Markovian structure and serial correlation potentially violate the necessary IID assumptions. Both tests, however, motivate their analysis with Kullback and Leibler (1951), and similar conclusions hold if we were to report either test. Table 23 reports the results. On all three types of observations, the estimated equilibria have smaller distances than either multinomial model on average. In addition, the equilibria have

[5] Given the large number of time periods and the small number of states, this latter possibility is unlikely.

[6] There are 436 dyad-state pairs in the data, because we lose some pairs when the dyad never enters the relevant state, e.g., the U.S. and Canada never go to war.

Table 23: Kullback-Leibler distance and comparison to non-strategic models

	Mean distance		
	Observed states	Conditional transitions	Transition matrix
Estimated equilibria	0.030	0.157	0.061
Mutinomial model	0.048	0.221	0.110
Mutinomial cond. on state	0.053	0.275	0.090

Caption: Cells report the mean Kullback-Leibler distance from the observed empirical distribution for three models (rows) and three types of observations (columns). Smaller numbers indicate better fit. For observed states, the equilibria have smaller distances than the multinomial and conditional multinomial models in 74.3% and 80.4% of observations, respectively. For conditional transitions and the transition matrix types, the corresponding comparisons are 57.8% and 68.1% and 82.1% and 83.8%, respectively.

smaller distances than the multinomial and conditional multinomial models in the majority of observations.

References

Aguirregabiria, Victor. 2012. "A Method for Implementing Counterfactual Experiments in Models with Multiple Equilibria." *Economics Letters* 114(2):190–194.

Aguirregabiria, Victor and Pedro Mira. 2007. "Sequential Estimation of Dynamic Discrete Games." *Econometrica* 75(1):1–53.

Billingsley, Patrick. 1961. "Statistical Methods in Markov Chains." *The Annals of Mathematical Statistics* 32(1):12–40.

Buhaug, Halvard and Kristian Skrede Gleditsch. 2008. "Contagion or Confusion? Why Conflicts Cluster in Space1." *International Studies Quarterly* 52(2):215–233.

Clarke, Kevin A. 2007. "A Simple Distribution-Free Test for Nonnested Model Selection." *Political Analysis* 15(3):347–363.

Davison, A.C. and D.V. Hinkley. 1997. *Bootstrap Methods and their Application.* Cambridge Series in Statistical and Probablistic Mathematics New York: Cambridge University Press.

Griewank, Andreas, David Juedes and Jean Utke. 1996. "Algorithm 755: ADOL-C: A Package for the Automatic Differentiation of Algorithms Written in C/C++." *ACM Transactions on Mathematical Software (TOMS)* 22(2):131–167.

Huth, Paul K. and Todd L. Allee. 2002. "Domestic Political Accountability and the Escalation and Settlement of International Disputes." *Journal of Conflict Resolution* 46(6):754–790.

Kullback, Solomon and Richard A. Leibler. 1951. "On Information and Sufficiency." *The Annals of Mathematical Statistics* pp. 79–86.

McFadden, Daniel. 1978. Modelling the Choice of Residential Location. In *Spatial Interaction Theory and Planning Models*, ed. Anders Karlqvist et al. Amsterdamn: Noth-Holland pp. 75–96.

Read, Timoth R.C. and Noel A.C. Cressie. 1988. *Goodness-of-Fit Statistics for Discrete Multivariate Data*. New York: Springer-Verlag.

Reiter, Dan and Allan C. Stam. 2003. "Identifying the Culprit: Democracy, Dictatorship, and Dispute Initiation." *American Political Science Review* 97(2):333–337.

Silvey, S. D. 1959. "The Lagrangian Multiplier Test." *The Annals of Mathematical Statistics* 30(2):389–407.

Vuong, Quang H. 1989. "Likelihood Ratio Tests for Model Selection and Non-nested Hypotheses." *Econometrica: Journal of the Econometric Society* 57(2):307–333.

Wächter, Andreas and Lorenz T. Biegler. 2006. "On the Implementation of an Interior-Point Filter Line-Search Algorithm for Large-Scale Nonlinear Programming." *Mathematical Programming* 106(1):25–57.

Walter, Sebastian F. 2014. "PyADOLC: a Wrapper for ADOL-C." GitHub repository.
URL: *http://github.com/b45ch1/pyadolc*

Weeks, Jessica L. 2008. "Autocratic Audience Costs: Regime Type and Signaling Resolve." *International Organization* 62(1):35–64.

Weeks, Jessica L. 2012. "Strongmen and Straw Men: Authoritarian Regimes and the Initiation of International Conflict." *Amercian Political Science Review* 106(2):326–347.

Xu, Eric. 2014. "Pyipopt." GitHub repository.
URL: *https://github.com/xuy/pyipopt*