War and Peace in the Shadow of the Future*

Casey S. Crisman-Cox† Michael Gibilisco‡

July 13, 2015

Abstract

This paper structurally estimates a dynamic game of international crisis escalation in which countries have infinite-time horizons and the strategic environment, or stage game, evolves as countries threaten and attack their rivals. In the model and its subsequent estimation, countries fully anticipate the future consequences of moving into and out of three states of conflict consisting of peace, crisis, and war, and their expected utilities in these three states are determined by observed variables, equilibrium play, and structural parameters. Most importantly, this approach allows us to nonparametrically estimate country-specific audience costs without specifying proxies and incorporating their dynamic trade-offs. We fit the model to Militarized Interstate Disputes (MIDs) incident level data (1993-2007) using a full-information constrained maximum likelihood estimator, test hypotheses related to the liberal peace and audience costs, and produce counter-factual experiments consistent with estimated equilibria. Our results provide some new support for the liberal peace in war but not crisis, and indicate that conflict spirals only in crisis and war. Furthermore, we find that democracy alone only explains some of the variation of audience costs, among anocracies and autocracies executive constraints helps explain audience costs in ways that suggest that more nuanced understandings of audience costs are required.

Word Count: 9202

*Thanks to Rob Carroll, Mark Fey, Tasos Kalandrakis, Brenton Kenkel, Curt Signorino, and audiences at the 2015 annual meeting of SPSA and University of Rochester. All remaining errors are the fault of the authors.

†PhD candidate, University of Rochester. Email: c.crisman-cox@rochester.edu

‡PhD candidate, University of Rochester. Email: michael.gibilisco@rochester.edu

1 Introduction

Long-standing theories, historical accounts, and common intuitions maintain that states and national leaders are strategic and forward-looking. They are long-lived and decide whether to engage in conflict given their rival's expected actions not only today but also tomorrow. At their very core, theories of conflict and cooperation, such as power transitions, interstate rivalries, audience costs, and wartime demands rely on the shadow of the future (Axelrod, 1984; Goemans, 2000; Powell, 1999). Indeed, models of wars have increasingly involved dynamic games,[1] and there is little reason to expect that world powers such as the United States and Russia interact without anticipating the long-term consequences of their actions. This motivates a unified theoretical and empirical investigation of conflict dynamics that raises the question of how countries navigate war and peace in the shadow of the future.

Such a framework has important implications for international relations and peace research. For example, consider audience costs, which are central to theories of conflict, crisis escalation and other international phenomena.[2] Despite their centrality, scholars have found it quite hard to move beyond traditional proxies, e.g., democracy and free press measures, when quantifying the magnitude and effects of audience costs (Eyerman and Hart, 1996; Kurizaki and Whang, 2014). Even with proxies, audience costs are essentially dynamic in nature and involve incentives that vary over time (Fearon, 1994). Today, audience costs prevent countries from starting new conflicts, but they encourage countries to continue conflicts tomorrow once hostilities begin (Downs and Rocke, 1995). When countries internalize these dynamics, their unobservable strategies create a time dependence between current escalation decisions and the expected path of future conflict, confounding the relationship between audience costs and the probability of escalation and retaliation. Such inference problems are not confined to audience costs, however. Strategic expectations and the current state of hostilities mitigate how countries move "into and out of" peace (Jackman, 2000; Signorino, 1999).

Our analysis addresses these concerns head-on by adopting a structural approach. It starts with the assumption that countries are long-lived and anticipate the decisions of their rivals both today and in the future. Specifically, we construct and subsequently estimate a dynamic game between two countries in which their strategic environment evolves according to past actions, and they transition into and out of three non-absorbing states of hostilities: peace, crisis, and war. Countries fully anticipate the expected evolution of

[1] See Debs and Monteiro (2014), Fearon (1994, 2004, 2013), Jackson and Morelli (2009), Powell (1999), Slantchev (2003), and Spaniel (2014) for recent contributions.

[2] For examples, see Bueno de Mesquita et al. (1999), Dorussen and Mo (2001), Fearon (1994), Goemans (2000), Jensen (2003), Partell and Palmer (1999), Prins (2003), and Weeks (2008), among others.

conflict, which proceeds by *unilateral escalation* (e.g. Bueno De Mesquita, Morrow and Zorick, 1997; Fearon, 1994). For example, if one country declares war on another today, then tomorrow the countries begin their interaction in a state of war.

In constructing the model, we incorporate several important features for empirical research. Countries' payoffs over the states of conflict and the available actions, such as threats, border violations, and declarations of war, are functions of variables and associated coefficients. This allows for the straightforward testing of hypotheses common in the literature, such as the democratic or capitalist peace, while incorporating long-run strategic dynamics. Furthermore, the game allows us to model audience costs as structural parameters independent of observed data.[3] Here, an audience cost is a country-specific fixed-effect that denotes the (dis)utility a country receives from backing down in a conflict when its rival maintains or escalates. These audience costs need not depend on *a priori* determined attributes, and can be uncovered without further parametric or functional form assumptions, which is substantially more flexible than other empirical endeavors. Likewise, we consider another structural parameter that describes how a country's cost of war changes with the action's of its rival. This effect varies depending on whether the countries are in war or peace and represents whether the expectation of conflict deters or encourages future conflict.

We fit the model to the Militarized Interstate Dispute (MIDs) dataset that records escalation deicsions at the monthly, if not, weekly level. To do this, we invoke a technique from Su and Judd (2012) who equate the estimation of a game to solving a constrained optimization problem. Specifically, we maximize a full-information likelihood function subject to constraints that describe equilibrium play. This constrained maximum likelihood (CMLE) approach tractably handles multiple equilibria and estimates a different, and basically independent, equilibrium for each dyad. This means that there is no need to assume that a very hostile dyad (e.g. Israel and Lebanon) and a very peaceful dyad (e.g. US and Canada) play the same equilibrium conditional on observed covariates.

Our initial results provide additional support for the democratic peace theory in which pairs of democracies are unlikely to go to war. More substantively, we find that increasing Lebanon's polity score by 6 points increases its probability of peace with Isreal by 12%, but the size and direction of the effect can vary across dyads. However, there is an important caveat; the analysis demonstrates that democratic dyads are actually at least as likely to engage in a crisis than other dyads. An additional finding in this vein is that we only find a pacifying effect of trade on war, not on minor disputes or conflicts. In fact as with democracy, we find that, if anything, trade dependence may encourage states to enter crisis.

[3] We are not the first to use a structural model to uncover audience cost. Kurizaki and Whang (2014) estimate audience costs using democracy as a proxy and a signaling model; however, such an approach does not incorporate the dynamics and flexibility of the one here.

This means that our results support a more fine tuned version of the liberal peace than most scholars examine: democracy and trade deter wars, but they do not prevent threats or crises. Together, this suggests that democrats and trading partners use interstate crisis to hammer out their disputes but do not want to see them escalate to armed conflict.

Furthermore, we find that whether joint escalation increases or decreases an individual country's cost of conflict varies with its current relations. In peace, when a country expects its rival to attack, it is discouraged from attacking as well, but in crisis and war, it is encouraged. In terms of systemic models of conflict (e.g. Braumoeller, 2008), we find that conflict deters in peace but spirals in crisis and war, and this result suggests that the security dilemma is important for dyads only after entering an international crisis. Such a nuance explains why conflicts tend to exhibit temporal clustering and peace appears to be self-enforcing (Beck, Katz and Tucker, 1998; Gleditsch and Ward, 2000; Jackman, 2000).

Our estimates of country-specific audience costs, find strong evidence that "gambling for resurrection" occurs in international conflict (Downs and Rocke, 1995). Gambling theories contend that countries face two opposing incentives from large audience costs. On one hand, these costs discourage countries from starting a crisis or conflict, but on the other hand they encourage countries to continue conflict once in an ongoing dispute. Because of this, the effect of audience costs on the likelihood of peace is unclear, and we find varying effects across dyads. For example, we find that a small increase in Lebanon and Israel's audience costs decreases the probability that the countries will be in peace by 10%, but a similar change increases this probability in the North Korea – South Korea dyad by about the same.

Furthermore, we use our estimated audience costs to examine a host of possible predictors, and they correlate with common proxies in the literature such as polity scores. However, we also find that simply using polity likely neglects other interesting factors that appear to determine audience costs. As others have pointed out, autocrats also face domestic constraints when navigating a crisis (Goemans, 2000; Weeks, 2008). More concretely, past work has established that the particular institutions within autocratic regimes can induce audience costs that are as high as in some democracies. We find additional evidence for these autocratic audience costs. First, we are able to replicate results from Weeks (2012) comparing audience costs across types of authoritarian regimes. Second, we find autocracies with legal provisions for executive removal have audience costs that are as high as democracies with weak participatory institutions, such as indirect executive election and low polity scores.

We subject the model to a series of fitness and robustness checks. When analyzing the predicted number of states and transitions in the data, goodness of fit measures along with anecdotal examination suggests good fit. In addition, our model out-performs a naive model using pooled averages, i.e., a non-strategic, all constants model. Furthermore, we estimate the model with a single audience cost parameter and find that the our substantive conclusions concerning the liberal peace and conflict spirals do not change. These conclusions also remain robust to other controls common in the literature.

Our paper relates to the literature on crisis escalation that has an established history of estimating parameters using likelihoods derived from equilibrium play (Carter, 2010; Esarey, Mukherjee and Moore, 2008; Kurizaki and Whang, 2014; Leemann, 2014; Lewis and Schultz, 2003; Signorino, 1999; Signorino and Tarar, 2006). Like our model, this literature analyzes why countries attack, threaten, or ignore their rivals and traditionally rationalizes observed data with action-specific shocks as in McKelvey and Palfrey (1995). (For an exception see Lewis and Schultz, 2003). Unlike our model, the literature considers finite games in which war and peace are absorbing outcomes and essentially end the game. Furthermore, our estimation strategy explicitly accounts for multiple equilibria, and we find substantial evidence that multiple equilibria exist under our estimated parameters. As discussed in Jo (2011), not accounting for this multiplicity leads to inconsistent estimates and incorrect counterfactuals which has been a roadblock in the estimation of previous models (e.g. Bas, Signorino and Whang, 2014; Kurizaki and Whang, 2014; Whang, McLean and Kuberski, 2013).

2 Dynamic Model of Crisis Escalation

In this section, we present a discrete-time, infinite-horizon dynamic game that models crisis and conflict between two countries. Its key feature is that countries begin a period in one of three states of conflict: peace, crisis, and war. In each period, countries simultaneously decide to engage in various levels of hostility, which potentially include making demands, standing firm, backing down or engaging in war. The most hostile of these choices determines the game's transition to a more or less severe state of conflict, which affects countries' future payoffs. Thus, a country may wish to attack its rival to gain resources or a strategic advantage today, but such a choice means the game enters a state of war tomorrow, which carries very different strategic incentives. In addition, we specify the model with empirical estimation as an end goal and incorporate the possibility that state- and action-specific payoffs depend on observable variables. Doing so allows for testing hypotheses concerning how country- or dyad-specific characteristics affect the prevalence of conflict while accounting for the long-term strategic behavior of countries and their leaders.

Consider two countries. We use i to denote an arbitrary country and $j \neq i$ is its rival. Time is discrete and indexed by $t = 1, ..., \infty$. In each period t, country i first observes a common state variable $s^t \in \{1, 2, 3\}$ and a private state variable ε_i^t, which represents action-specific shocks and is unknown to its rival. Here s^t denotes the current level of hostility, where $s^t = 1$ denotes the countries are in a state of peace, $s^t = 2$ a state of crisis, and $s^t = 3$ a state of war.[4] Each country then simultaneously chooses a level of hostility against its competitor. Let $a_i^t \in \{1, 2, 3\}$ denote country i's action in period t and a^t a profile of actions, i.e., $a^t = (a_i, a_j)$. Here, a_i^t takes the values 1, 2, and 3 which indicate peaceful, crisis-level (threat/demand), and war-level (attack/invasion) actions, respectively.

The common state variable s^t evolves according to past actions, and we assume escalation is deterministic and unilateral, that is, $s^t = \max\{a_i^{t-1}, a_j^{t-1}\}$.[5] Thus, the model captures situations in which a country declares war ($a_i^t = 3$) on its rival, and the next period begins with the two countries in a state of war ($s_i^{t+1} = 3$). We denote country i's shock to action a_i in period t as $\varepsilon_i^t(a_i^t)$ where $\varepsilon_i^t(a_i^t)$ are independently and identically distributed type I extreme value across actions, players, and states, which are standard independence assumptions in these types of games.

Let θ denote a vector of relevant structural parameters. Country i's per period payoff against country j is given as $u_{ij}\left(a^t, s^t; \theta\right) + \varepsilon_i^t(a_i)$. In other words, the per-period payoff consists of a deterministic and a stochastic component. Given a sequence of action profiles, states, and action-specific shocks $\{(a^t, s^t, \varepsilon_i^t)\}_{t=1}^{\infty}$, country i's total payoff is the discounted sum of per-period utilities:

$$\sum_{t=1}^{\infty} \delta^{t-1} \left[u_{ij}\left(a^t, s^t; \theta\right) + \varepsilon_i^t(a_i^t)\right],$$

where $\delta \in [0, 1)$ denotes a common discount factor.[6]

We endow u_{ij} with the following functional form:

$$u_{ij}(a, s; \theta) = x_{ij} \cdot \beta^s + z_i \cdot \kappa^{a_i} + \mathbb{I}[a_j \geq s > a_i]\alpha_i + \mathbb{I}[a_i > 1]\mathbb{I}[a_j > 1]\gamma^s, \tag{1}$$

which means country i's utility consists of four components. First, it receives a state-specific payoff from being in a state s with country j. This takes the form $x_{ij} \cdot \beta^s$, where x_{ij} is a vector of dyad-specific variables and β^s a vector of associated coefficients. Dyad variables could be directed, e.g., military capability ratio between i and j, or undirected, e.g., minimum democracy between countries i and j. Furthermore, the state payoff $x_{ij}^s \cdot \beta^s$

[4]Hereafter, a *state* denotes the commonly observed level of hostility s^t, and we refer to the actors in the game as countries.

[5]Unilateral escalation is common in the crisis and conflict literature (Bueno De Mesquita, Morrow and Zorick, 1997; Fearon, 1994; Kurizaki and Whang, 2014; Schultz, 2001).

[6]Notice we do not let δ vary according to the country. Without further structure, the discount factor is highly colinear with state and action payoffs. Specifically, as δ increases, the maginitude of the current state and action payoffs can increase and produce similar expected utilities. In other words, it is difficult to identify δ with the data and model we employ. Instead, we fix $\delta = 0.9$ throughout the paper as is common when estimating dynamic games.

could also represent an expected utility of country i in some lottery or finite-period game that varies with s as long as the outcomes of this lottery or game do not affect the state transition probabilities or payoffs of the dynamic game detailed here. For example, consider the classic war lottery. If p_{ij} represents the probability i wins in a war against j, π_{ij} the benefit of winning, and c_{ij} the destructiveness of war, we could assume the war payoff takes the form $x_{ij} \cdot \beta^3 = p_{ij}\pi_{ij} - c_{ij}$.

Second, regardless of the state, if country i chooses action a_i, i pays some costs $z_i \cdot \kappa^{a_i}$, where z_i is a vector of country-specific variables and κ^{a_i} a vector of associated coefficients. These costs of escalation capture important transaction costs from declaring war, formally threatening a rival or maneuvering a military in order to later fight a war, although they are often not explicitly considered in previous models.[7] Notice that i's cost of action a_i does not depend on the characteristics of j. This is an important identification assumption, but paired with the state-specific payoff, this leads to a natural interpretation. Together, these two payoffs imply, for example, that while the US pays the same cost from declaring war on Afghanistan and on Russia, it can still possess a preference for being at war with Afghanistan over being at war with Russia. Without loss of generality, we adopt the normalization that $\beta^1 = \kappa^1 = 0$.

The remaining two components in i's utility depend on j's actions. The parameter α_i measures country i's audience costs. That is, if i and j are engaged in a conflict ($s > 1$), j continues or escalates the conflict ($a_j \geq s$) and i backs down ($a_i < s$), then i incurs cost α_i. Intuitively, this means that countries never realize audience costs when in peace or if their rivals do not escalate/maintain the current conflict. In addition, countries can receive audience costs regardless of who initially began the conflict as in Fearon (1994). Notice that we do not impose a functional form on α_i, such as a function of country specific characteristics (e.g. democracy levels). Thus the model uncovers audience costs without further parametric assumptions. This allows us to test hypotheses from theories that relate observed characteristics of countries to audience costs, without baking them into the estimated quantities. In addition, the presence of action-specific costs means we also control for the possibility that some countries are more or less likely to engage in conflict, regardless of their rival's expected actions.

The parameter γ^s measures how i's cost of escalation varies with j's actions in state s. When $\gamma^s > 0$, i's cost of escalation ($a_i > 1$) decreases when j escalates. Similarly, when $\gamma^s < 0$, i's cost of escalation increases when j escalates. For example, if countries receive a large benefit in state s when their rivals attack first, such as support from an

[7] This may appear to be a non-standard modeling choice, but this version subsumes the case in which the coefficients κ^{a_i} are zero.

Figure 1: A Diagram of Standard Crisis-Escalation Models

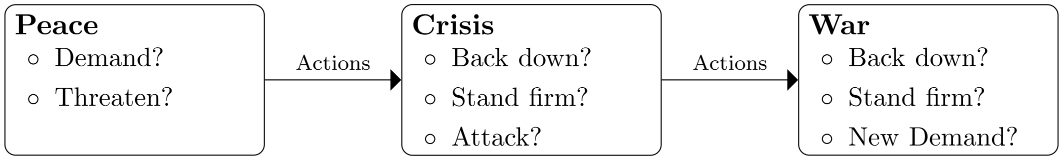

international community, then γ^s would be strongly negative.[8] Theoretically, γ^s represents other strategic incentives as to why a country does not escalate a conflict independently of audiences costs including a host of factors that contribute to potential second-strike advantages. Substantively, the sign of γ^s relates to Braumoeller (2008)'s empirical analysis of spiral and deterrent theories of systemic conflict. If γ^s is positive, conflict spirals in state s because a country's cost of conflict decreases when it expects its rival to attack. In contrast, conflict deters when γ^s is negative.

Before describing strategies and equilibria, we briefly discuss how the model presented here relates to those in the previous literature. In a large class of crisis escalation games, the strategic interaction is sequential, and each decision node belongs to one of the three states of peace, crisis, and war. Figure 1 summarizes the extensive-form representation of these games (Bueno De Mesquita and Lalman, 1992; Bueno De Mesquita, Morrow and Zorick, 1997; Fearon, 1994; Kurizaki and Whang, 2014; Schultz, 2001; Slantchev, 2005).[9] Often in these models, two countries begin in a state of peace and determine whether to make demands or threaten their rivals. If either country makes a demand, the pair transitions to a crisis state in which they then decide whether to stand firm with their demands, back down, or engage in war. Here, standing firm and attacking can be equivalent, and escalation to war can be either bilateral (e.g. Bueno De Mesquita and Lalman, 1992) or unilateral (e.g. Bueno De Mesquita, Morrow and Zorick, 1997; Fearon, 1994). Finally, depending on the actions and these transitions, the countries either remain in the crisis state or transition to the other two states, which are absorbing at this point in the game. Our model relaxes the linearity seen in Figure 1, by allowing countries to rotate into and out of the three states. Furthermore, none of these states are absorbing in our model, which emphasizes the dynamic (in both the colloquial and technical sense) nature of the strategic interaction.

[8] It is possible these mutual payoffs γ^s are actually realized in the next period rather than the current period. If this is the case, we can always weight the estimated parameters by δ^{-1}, and the estimates reported below would only be attenuated toward zero.

[9] Obviously, Figure 1 ignores some important differences, which include the possibility that the countries begin in crisis (e.g. Fearon, 1994) or the endogenous creation of military power (e.g. Slantchev, 2005). Note that Senese (1997) independently drew a nearly identical figure to characterize theories of conflict and escalation.

3 Strategies and Equilibria

We focus on stationary Markovian strategies as is standard in these dynamic games. In addition, equilibria with these types of strategies potentially serve as focal points because of their simplicity. Accordingly, we drop references to the period hereafter. As in Aguirregabiria and Mira (2007), Bayesian-Nash equilibria in stationary Markovian strategies can be represented as vectors of expected utilities. Formally, $v_i(a_i, s)$ denotes 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

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

which 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

$$P(a_i, s; v_i) = \frac{\exp(v_i(a_i, s))}{\sum_{a_i'} \exp(v_i(a_i', s))} . \tag{2}$$

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

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

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). This closed-form solution to $G(s, v_i)$ greatly simplifies the computational burden in the subsequent analysis by removing the need to perform numerical integration, and this is the major reason for assuming the distribution of private information described above.

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 is given by function $\Phi_{ij}(a_i, s, v; \theta)$, where

$$\Phi_{ij}(a_i, s, v; \theta) = \sum_{a_j'} P(a_j', s; v_j)\left[u_i(a_i, a_j', s; \theta) + \delta G\left(\max\{a_i, a_j'\}, v_i\right)\right] . \tag{3}$$

In words, Equation 3 takes a profile of values v, supposes countries play according to the associated choice probabilities in Equation 2, and then computes new expected values of each action in each state. If v is an equilibrium profile, it must be a fixed point of these 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 profile v such that $\Phi(v;\theta) = v$.[10]

4 Constrained Maximum Likelihood

For our empirical strategy, we use a full-information constrained maximum likelihood estimator (CMLE) as advocated by Su and Judd (2012). This process involves maximizing a constrained multinomial likelihood function derived from the choice probabilities in Equation 2 while simultaneously imposing the equilibrium constraints in Equation 3.[11]

We consider $N \geq 2$ countries and D dyads from these N countries. Each $k \in \{1, ..., D\}$ denotes a game described in the previous section, and we index the dyad with a superscript k hereafter. We use data that can be summarized as a list $\{X, Z, Y\}$. Here X and Z are matrices of ordered-dyad and country-specific variables, respectively, which enter the stage utilities through Equation 1. In addition, Y is a collection of time-series matrices, which detail observed state and action profiles for each dyad. Specifically, $Y^k = \left(s^{kt}, a^{kt}_{i^k}, a^{kt}_{j^k}\right)_{t=1}^{T}$. Let $\bar{\theta}$ denote the true vector of parameters. For each dyad k, we assume the data Y^k were generated from a *single* equilibrium, $\bar{v}^k$. i.e., $\Phi^k(\bar{v}^k; \bar{\theta}) = \bar{v}^k$. While multiple equilibria potentially exist in the game between the countries i^k and j^k, the procedure requires that Y^k comes from only one of these. The goal then is to now estimate $\bar{\theta}$ and v^k.[12]

Let $\mathbf{v} = (v^1, ..., v^D)$ denote the vector of all profiles of expected utilities. The log-likelihood takes the following form:

$$\mathcal{L}(\mathbf{v} \mid Y) = \sum_{k=1}^{D} \sum_{t=1}^{T} \left[\log P\left(a^{kt}_{i^k}, s^{kt}; v^k_{i^k}\right) + \log P\left(a^{kt}_{j^k}, s^{kt}; v^k_{j^k}\right) \right], \tag{4}$$

which is the ordinary multinomial logit log-likelihood summed over dyads, time periods, and players. To ensure that the choice probabilities describe an equilibrium, we require v^k satisfy equilibrium constraints. With a slight abuse of notation, the CMLE estimates,

[10] An equilibrium exists. The function $\Phi_{ij}(a_i, s, v; \theta)$ is weighted sums 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.

[11] Egesdal, Lai and Su (2013) compare the CMLE to other methods when estimating dynamic games and find that the CMLE has better (less biased and more efficient) parameter estimates and convergence diagnostics. Because this is a relatively new method and our data include games corresponding to country-dyads, this section describes the estimator in more detail.

[12] We do *not* require that Y^k and $Y^{k'}$ come from the same equilibrium. In fact a strength of our approach is that we estimate a separate equilibrium for each dyad in our data.

$(\hat{\mathbf{v}}; \hat{\theta})$, solve the following optimization problem:

$$\begin{aligned} &\max_{(\mathbf{v};\theta)} \frac{1}{T} \mathcal{L}(\mathbf{v} \mid Y) \\ &\text{subject to } \Phi^k(v^k; \theta \mid X, Z) = v^k, \; k = 1, ..., D. \end{aligned} \tag{5}$$

The consistency of the estimator as the number of observed periods, T, increases follows from the standard results concerning the Lagrange Multiplier Test in Silvey (1959). The key idea behind consistency in T originates from the nature of the equilibria we consider. Specifically, with the model's assumptions, we can learn the equilibrium choice probabilities precisely as we observe longer and longer segments from the path of play. Under some regularity and identification conditions, we can then accurately estimate the parameters of interest, θ.[13] Note that consistency in the number of games or dyads is not guaranteed.[14] Indeed, there is a obvious incidental parameters problem because adding another dyad requires an additional 18 equilibrium constraints. Nonetheless, we gain leverage by pooling information across dyads when T is sufficiently large, and this pooling is necessary for identification with a larger number of variables.[15] Furthermore, Monte Carlo experiments suggest that, with a reasonable number of time periods, moderate increases in the number of dyads reduce the estimator's bias and variance, and these numbers correspond to those in our data. (See Appendix B.)

The CMLE requires estimating a potentially large number of auxiliary parameters. Our data contains 180 dyads and 119 countries, and to solve Equation 5 we estimate more than $3,347$ parameters, of which $3,222$ are expected utility constraints. The high-dimensionality of the problem raises some questions about feasibility. To efficiently estimate the model, we use a large-scale optimizer created to solve problems with hundreds of thousands of variables (Wächter and Biegler, 2006) and algorithmic differentiation for derivatives (ADOL-C). This approach allows us to efficiently simulate standard errors using a parametric bootstrap. Further details are relegated to Appendix A. Additionally, Appendix B contains a Monte Carlo experiment analyzing the estimator's performance and the identification of the model.

[13] Here, the key additional assumptions, besides those required for the traditional MLE results, are the continuity of the constraint functions and that the true vector $(\bar{\mathbf{v}}, \bar{\theta})$ resides in the interior of the larger, compact and convex parameter space. Silvey (1959) also characterizes the asymptotic normality of the estimator.

[14] Such a result would require that all dyads or games play the same equilibrium, an impossibility with dyad-specific variables. In this case, increasing the number of games is identical to increasing the number of time periods.

[15] See Pesendorfer and Schmidt-Dengler (2008) for a upper-bound on the number of parameters one can estimate using information from a single game.

Table 1: Distribution of transitions in the data.

Transition	Percent of data	Percent within each state
Peace → Peace	92.5%	97.0%
Peace → Crisis	1.81%	1.90%
Peace → War	1.09%	1.10%
Crisis → Peace	1.82%	72.2%
Crisis → Crisis	0.49%	19.6%
Crisis → War	0.21%	8.24%
War → Peace	1.09%	53.4%
War → Crisis	0.22%	10.6%
War → War	0.74%	36.0%

Caption: The middle column displays the probability distribution over the nine types of possible transitions, and the far-right column presents the conditional distribution in each state.

5 Data

To estimate the model we use the MIDs 4.0 data (Ghosn, Palmer and Bremer, 2004) to define the observed path of play, Y^k, for each dyad. In particular we make use of the incident level data known as MID-IP 4.01 (Kenwick et al., 2013), which records each interstate incident between 1993 and 2010. Incidents record the actions taken by individual states within a dispute. Dispute numbers can be used to determine what state or states the actions were taken against. The actions recorded by the MID-IP are on the standard 22-point MID scale ranging from no action to joining an interstate war. We use this scale to form the three levels of hostility countries can take against each other: peace, crisis, and war. The peaceful action is recorded for country i against j in period t if the MID-IP records no militarized action (MID-IP action 0) from i to j. We code the action as 'crisis-level' if the country commits an action that is between a threat and an attack (MIP-IP actions 1-15). Lastly, a 'war-level' action is recorded if the state attacks or takes a more hostile action (MIP-IP actions 16-21). These actions are then used to create the state transitions. Because MIDs data records the day of each incident, we can make our time periods a calender month. Ultimately, we are left with 179 dyads with 180 observations each for the final analysis. With this classification in hand, approximately 93% of states in the data are peace, and 97% of actions are peace-level. In addition, Table 1 records the nine different types of possible transitions and provides preliminary evidence that countries condition their behavior on the state variable of interest. That is, the conditional distribution of transitions changes substantially across our observed states.

Note that we limit ourselves to the study of countries and dyads that have experienced a MID, which we do for two reasons. First, we cannot include countries that have never entered a state of conflict or experienced a militarized dispute with any other country. Including these countries introduces substantial separation issues into the maximization routine because the likelihood of the observations including these countries will be strictly increasing as their audience cost parameters tend toward negative infinity. Because a major goal is to estimate these audience costs, we focus on dyads that have a recorded MID, removing separation issues.[16] Second, the computational cost of implementing our estimator with more than 300 dyads is significant. Currently, we require approximately 500 megabytes of memory to compute and store the Hessian of the Lagrangian in a sparse format. More than $1,000$ dyads requires over 30 gigabytes of RAM.[17]

For the dyad-level independent variables, those associated with β^s above, we use the minimum democracy level in the dyad, the logged capability ratio, and the square-root of the trade interdependence. As is standard practice we measure the minimum level of democracy by using the dyadic minimum polity2 score from the Polity IV database. Capability ratios are computed as the ratio of CINC scores from the Correlates of War (COW) National Material Capabilities 3.0 (NMC) dataset (Singer, Bremer and Stuckey, 1972). Trade interdependence is measured in the usual fashion (Gartzke, 1998; Oneal and Russett, 1997) where state i's interdependence on state j is the sum of exports and imports between i and j divided by i's GDP. Trade data comes from the COW dyadic trade data (Barbieri, Keshk and Pollins, 2009), supplemented by data from Gleditsch (2002). GDP data is take from the Penn World Table (PWT) 8.0 (Feenstra, Inklaar and Timmer, N.d.) and supplemented with data from the World Bank. For the variables associated with country specific costs, we include logged GDP per capita and logged military personnel per capita. GDP is again taken from the PWT and population and military personnel are taken from the NMC.

We use the mean value of all of these variables over the course of the time period we study (1993-2007) to produce the x_{ij} and z_i. While there is a legitimate concern that some these variables are endogenous to the conflict process itself, the variables in the analysis show little change over the time frame considered here. Even when these variables do change, we find that their is no correlation between these changes and the observed states and actions. See Appendix C for details.

6 Results

Using the data described above, we performed the constrained optimization problem in Equation 5, and the results are reported in Table 2. Note that Table 2 is not reporting five models, but rather reports the output from one model across five columns. The first

[16] These dyads would be removed in any fixed-effects analysis (Beck and Katz, 2001).

[17] We are working on a sampling scheme in which we sample from peaceful dyads consisting of those countries in the MIDs dataset and optimize a weighted likelihood function subject to equilibrium constraints.

Table 2: Estimation results omitting α_i

	β^{Crisis}	β^{War}	κ^{Crisis}	κ^{War}	γ
Joint Democracy	0.00	-0.03^{***}			
	(0.01)	(0.01)			
Cap. Ratio	0.01	$-0.04^{\dagger}$			
	(0.01)	(0.02)			
Trade Depend.	0.33	-2.43^{*}			
	(0.45)	(0.97)			
GDP pc			0.12^{***}	-0.15^{***}	
			(0.02)	(0.03)	
Mil. Per. pc			10.18^{**}	$-10.29^{\dagger}$	
			(3.39)	(5.75)	
γ^{Peace}					-53.32^{***}
					(3.13)
γ^{Crisis}					9.52^{***}
					(0.58)
γ^{War}					13.98^{***}
					(0.58)
Constant	19.78^{***}	19.26^{***}	-21.72^{***}	-21.17^{***}	
	(1.08)	(1.52)	(0.99)	(1.37)	
Log L			-45.50		

Notes: $^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$; $^{\dagger}p < 0.1$
Bootstrapped Standard Errors in Parenthesis

two columns estimate state-specific payoffs that country i gets from its current relationship 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 has the estimates of the structural parameters γ^s which describe whether or not aggression is mutually advantageous or destructive in state s. The 125 audience costs parameters are omitted from the table, and we defer their discussion until the following section.

First, we observe in the model is that we find support for a liberal peace in war but not crisis. Our results suggest that as joint democracy levels (measured by the minimum polity2 score in the dyad) increase the utility from being in the conflict state decreases. This is in line with classic works on the democratic peace like Oneal and Russett (1997) and Gartzke (1998). What is interesting is that we find support for the proposition that while democracies do not like to be in conflict with each other they are just as likely to engage in crisis as any other pair of countries. While this is not a new finding in the conflict literature, it has not receive a lot of attention since being observed by Senese (1997). What it tells us is that while democracies tend not to prefer going to war with each other, they want to saber rattle and shout just as often as any other regime type. This suggests evidence for a democratic peace where democracies can use crises to credibly size each other up and then reach a settlement without resorting to war, which is consistent with theories by Bueno de Mesquita et al. (1999); Schelling (1960); and Schultz (1998).

In addition we also find evidence for a pacifying effect of trade on war but not crisis. Increased trade dependency appears to result in a lower utility for being in the war state, which makes sense as war can upset trade relations between states. Being highly trade dependent then reduces the utility of being at war which is the second half of the liberal peace. What is interesting here is that, like democracy, trade appears to have no, but possibly positive, effect on the utility a country gets from being in crisis. This suggests that trading partners, like pairs of democracies, are content to have minor disputes and scuffles, but they want to avoid the detrimental effects of war more than non-trading dyads. The result here speaks to a debate on trade's effect on MIDs with some scholars finding a pacifying effect (e.g. Gartzke, 1998; Oneal and Russett, 1997) and others finding that trade increases hostility (Barbieri, 1996). Our findings suggest that the strategic context and the level of conflict matter in determining a country's preferences over the two types of disputes in relation to peace.

The last dyadic variable we consider is capability ratios. What we can see here is that as country i gains power relative to j its utility for being in crisis may increase. Presumably this is because it can do better in crisis because it does better under the status quo. Our finding of a negative coefficient on capability ratios in the conflict state likely means that increasing a country's relative strength means it prefers to solve its disputes without going to war, because it already prefers the status quo.

Turning our attention to the country-specific costs for each action we first note the large negative constants. This is reassuring and provides some face validity to our results. The negative constants tell us that the baseline cost of starting a crisis or conflict is utility reducing, which just means that we find support for the standard supposition that war is costly. Increasing either GDP per capita or military personnel per capita appears to decrease the cost of playing the crisis action. This tells us that wealthy and militarily powerful states can better afford crisis and they pay less for these actions, an intuitively satisfying result.

However, playing the war action seems to have very large costs which increase (become more negative) in both per capita GDP and military personnel. This is an interesting result, but it is also believable. Entering war requires spending more money and resources than being at peace. This funnels money away from consumption, and is (at least in the short-run) bad for growth and economically destructive. In other words, wealthier states have larger opportunity costs than their counterparts with comparatively smaller economies. Likewise, having a larger military might mean that countries spend more to mobilize, supply, coordinate, and transport it. Playing the crisis action does not require large military states to do any of these things.

Finally, we examine the last column in Table 2 which are the parameters that capture the additional payoff that actors receive if they both play a non-peaceful action in a particular state of the world, i.e., γ^s. When the dyad is at peace, we see that mutual hostility is very destructive. In other words, when country j attacks country i, i's cost of escalation increases (becomes more negative). Intuitively, disputes are costly situations and countries only begin one if they expect the other country to roll over. In other words a negative gamma in peace suggests that countries are playing something like chicken, and the expectation of conflict deters. In contrast to this cautiously optimistic finding that there is a durability of peace, we find a persistence of war and crisis. The positive coefficients on γ^2 and γ^3 mean that once countries are in a dispute, country's i's cost of continuing the conflict *decreases* when its rival j continues as well. In these states, conflicts spiral. This means that concerns about the security dilemma are particularly relevant when there is an active dispute between two countries. This tragic persistence of conflict may in turn generate commitment problems and pose barriers to settlement.

7 Audience Costs

We now turn our attention to the 125 country-specific audience costs parameters that we estimated. The point estimates of these parameters are presented in Figure 2.[18] Countries are sorted by the size of their audience cost from the most negative (highest costs) to the most positive (lowest costs).

One thing we can notice that although most of the top ten are democracies, it seems that there are plenty of exceptions to the basic idea that democracy is synonymous with high audience costs. On the surface it seems that autocracies and anocracies are well represented throughout the range of audience costs. This offers some *prima facie* evidence in favor of arguments that suggest that autocrats leaders in weak states can face removal threats that create large audience costs (Chiozza and Goemans, 2011; Weeks, 2008).[19]

Further visual inspection of the estimates reveals a few points of interest, one of which is the similarity of audience costs between North and South Korea. South Korea's audience cost estimate is very low (less negative) compared to other advanced democracies. One

[18] Point estimates in table format can be found in Appendix E.

[19] Despite this, difference-in-means tests reveal that democracies have larger audience costs than both autocracies and anocracies, but found no difference between the later two.

reason this may be the case is the fact that South Korea exists in a state of perpetual siege which might mean that its voters are more willing to give their leader free range to do whatever he or she thinks is best to avoid war with North Korea which may include backing down when necessary. A similar story may explain why Israel and India both have very low audience costs for democracies.

The skeptical reader may, at this point, be concerned that these estimates actually reflect other costs/benefits associated with backing down. One such alternative story could be the existence of a second-strike disadvantages/advantages. However, as mentioned above, benefits or drawbacks of being the second mover is actually captured by the γ parameters. Likewise, the audience cost parameters do not include a general preference for peace as such an incentive is measured by constants in κ^{a_i}. Indeed, we have a number of reasons to be confident that we are in fact measuring audience costs. Chief among these is that the utility function characterizes audience costs in a manner consistent with other major works (Fearon, 1994; Kurizaki and Whang, 2014; Lewis and Schultz, 2003).

The second reason is that all audience cost estimates are negative suggesting that leaders face punishment for backing down in either crisis or war. Thus, we have estimated actual penalties even though the parameter estimates were not constrained to be negative. Even more impressive is that we uncover believable looking estimates without the strong functional form assumptions imposed by other studies (e.g. Kurizaki and Whang, 2014).

For one more additional assurance that we are producing estimates that really capture the effect of domestic audiences on conflict behavior we turn to Weeks (2012). In this paper, she argues that different types of autocratic regimes should be affected differently by their domestic audiences. To test this, she classifies autocrats into four regime types: Machine (nonpersonalist civilian regime), Junta (nonpersonalist military regime), Boss (personalist civilian regime), and Strongman (personalist military regimes).[20] Her theory suggests machine regimes will have audience costs very similiar to democracies, juntas will have smaller domestic costs, and bosses and strongmen will be the least constrained. We can use these classifications to provide some validity that we have a estimated a measure of audience costs. Table 3 shows the estimated relationships between these regime types and audience costs using weighted least squares (WLS), where we weight by the number of dyads in which country i is involved. In the regression, democratic regimes are the excluded category and so the estimated coefficients represent the average difference in audience costs between that type of autocratic regime and a democratic regime. Larger audience costs are more negative and so a negative slope means that on average a regime has higher audience costs than a democracy.

Note that we have uncovered the exact same relationships that Weeks hypothesized

[20] For more discussion on these classifications of authoritarian regimes see Weeks (2012) or Slater (2003).

Table 3: Week's Regime Classification and Audience Costs

	Audience Costs WLS
Machine	-1.43
	(1.63)
Junta	1.93
	(2.11)
Boss	3.94^{*}
	(1.58)
Strongman	1.15
	(2.00)
Other Nondemocracy	2.13^{*}
	(1.01)
Population	0.36
	(0.28)
Constant	-18.32^{***}
	(2.85)
adj. R^2	0.04
N	125

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

regarding the strength of domestic audiences in these various regime types. Machines have, if anything, larger audience costs than democracies, while the rest have lower audience costs. What is most remarkable is that boss leaders have, on average, *much* smaller audience costs than democrats which is Week's exact theoretical expectations. This regression provides reasonable evidence that the values we have estimated from our model are in fact capturing domestic constraints on the use of force.

With these assurances we can now turn our attention to using these measures to explore the guts of audience costs. Something else we may be interested in is how well standard measures of audience costs map into our estimated parameters. Note that this exercise cannot be done with past efforts to uncover audience costs, such as those by Kurizaki and Whang (2014) as they use a linear in the parameters specification to bake in variables like democracy. The structural approach we employ does not assume any relationships between our audience costs measures and standard observable proxies, which means that we can use them to examine how well standard proxy variables do at predicting audience costs. To do this, we regress our estimated results on a series of candidate proxies for audiences; we again use WLS with the same weights as above. We use five different proxies for audience costs to see which are good predictors. The proxies we use are polity2 and executive constraints (exconst) from the polity IV dataset, the winning coalition (W) from (Bueno de Mesquita et al., 2005), a freedom of the press dummy that is taken from Freedom House and Li (2005), and from the Institutions and Elections Project (IEP) we use dummies for whether country directly elects its chief executive and whether the executive can be removed by the legislature (Regan, Frank and Clark, 2009). For additional controls we include dummies for the electoral system within the country and logged population. The results of the WLS are presented in Table 4.

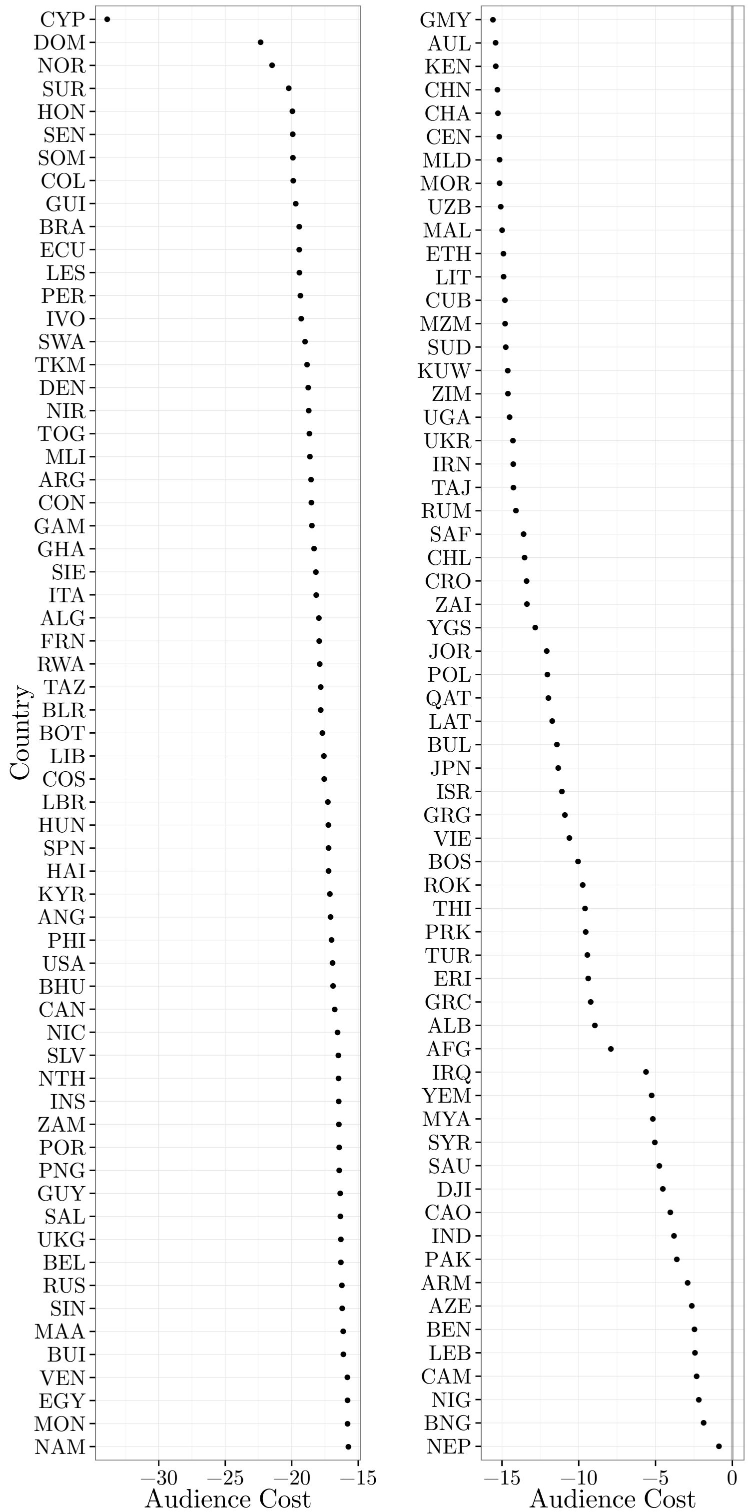

Figure 2: Country-specific audience costs

Table 4: Correlates of Audience Costs

	Audience Costs					
	WLS					
	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
polity2	-0.23^{**}					
	(0.08)					
Executive Constraints		$-0.42^{\dagger}$				
		(0.24)				
W			-3.85^{*}			
			(1.90)			
Free Press				-2.26^{*}		
				(1.04)		
Elected Executive					-1.26	
					(0.98)	
Removal of Executive						-1.87
						(1.21)
Plurality	0.25	0.18	0.28	0.13	-0.03	0.14
	(0.28)	(0.29)	(0.30)	(0.28)	(0.29)	(0.30)
Majority	-1.00	-1.63	-1.91	-2.40	$-2.56^{\dagger}$	-2.35
	(1.60)	(1.62)	(1.56)	(1.51)	(1.55)	(1.63)
PR	-2.06	$-3.04^{\dagger}$	-2.24	-3.11^{*}	-3.51^{*}	-3.24^{*}
	(1.58)	(1.57)	(1.70)	(1.50)	(1.57)	(1.60)
Mixed	-0.60	-2.12	-1.82	$-2.84^{\dagger}$	-3.19^{*}	-2.62
	(1.80)	(1.76)	(1.72)	(1.54)	(1.59)	(1.70)
Population	-0.19	-1.64	-1.07	$-2.39^{\dagger}$	$-2.56^{\dagger}$	-2.00
	(1.63)	(1.54)	(1.63)	(1.37)	(1.43)	(1.54)
Constant	-15.02^{***}	-11.87^{***}	-12.72^{***}	-12.18^{***}	-10.14^{**}	-11.58^{***}
	(3.21)	(3.08)	(3.09)	(3.05)	(3.21)	(3.18)
adj. R^2	0.08	0.05	0.05	0.06	0.04	0.05
N	125	122	125	125	124	121

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

Standard Errors in Parenthesis

What we see in Table 4 is that many of these common predictors are associated with larger (more negative) audience costs, which is consistent with common conceptions that democracies have larger audience costs. Polity2, polity's measure of executive constraints, Bueno de Mesquita, et al.'s W, freedom of the press are all significantly associated with audience costs in ways that have face validity. Likewise, the two measures of executive constraints/independence that we took from IEP are in the expected directions although they are not significant.

Another seemingly good predictor is the presence of a PR or mixed system. These systems tend to have higher audience costs than countries without electoral systems (the excluded category). They are the only democratic systems to exhibit this property in more than one of the regression models and may suggest a higher level of responsiveness to public outrage at decisions to back down than other systems. PR systems, in particular, tend to have a large negative coefficient suggesting that on average these regimes have the largest domestic costs. This is consistent with some research that demonstrates that PR systems are very responsive to public demands (Powell, 2000).

We may be interested in how these candidate predictors of audience costs do in a head-to-head competition. To see this we employ the machine learning technique of a regression tree (Venables and Ripley, 2003). The process selects variables that best predict breaks within the data and discards those candidate proxies that do not have sufficient predictive oomph. An advantage of this approach is it allows us to uncover interactive relationships between the candidate predictors that we may not otherwise find. In this process we remove executive constraints from the list of candidates because it is a component of and correlates highly with polity. The output of the regression tree is shown in Figure 3, where the values at each terminal node refer to the average audience cost amongst classified countries and the percentage of observations at the node.

According to this method the first best predictor of higher audience costs is if a state has an average polity2 score greater than or equal to –5. As a first cut this interesting as it seems to lump democracies and anocracies together as having larger audience costs than autocracies. This pattern is supported by a glance at the Figure 2, as we see the states with the 14 largest audience costs are in this range of the polity scale. What is interesting beyond this is that tree tells us that size of the winning coalition, W, is a good predictor conditional on the state having a polity score greater than –5. On average, states with the highest audience costs (most negative) are those with polity greater than –5 and a winning coalition that is larger than 0.84. The model estimates that states in this category have, on average, audience costs equal to about –16, which is just above the median estimated audience cost. Countries in this category are mostly advanced democracies like the United States, Great Britain, France, and Norway.

Regression Tree: Predictors of Audience Costs

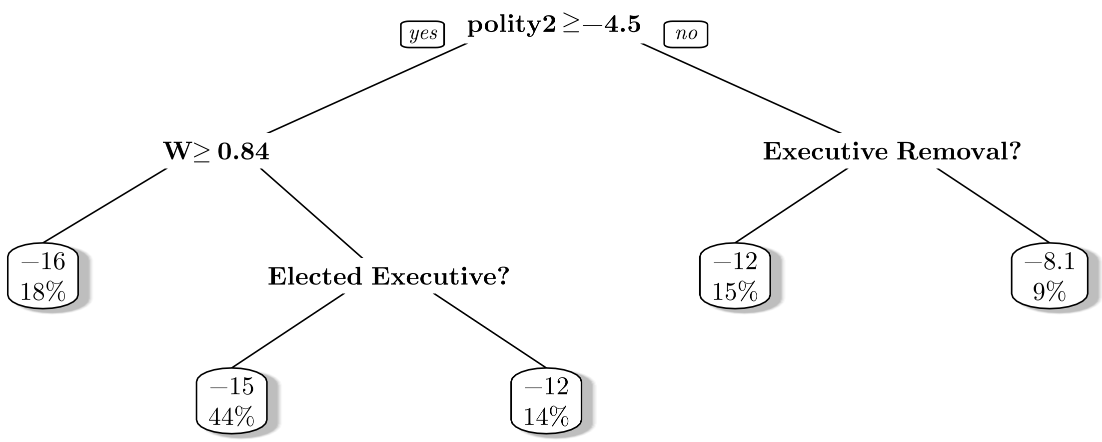

Figure 3: Predictors of Audience Costs

States with a polity greater than –5 buy with a W less than 0.84 are further split into two groups depending on whether or not the chief executive is directly elected. In countries where there is direct election the leaders seem to face larger audience costs with an average of –15 compared to –12. Countries in the former group include a grab-bag of democracies, anocracies, and autocracies, such as South Korea, Israel, Georgia, Iran, Turkey, and Yemen. The later group is formed almost entirely of weak democracies, developing countries, and a few autocrats. Included here are India and Cambodia.

Turning our attention to the truest of autocrats, polity score less than –5, we see that they are split into two groups depending on whether or not there are provisions for legislative removal of the executive. Autocracies that do not have these institutions for executive removal are the countries with the lowest average audience costs while the remaining regimes look more like their anocratic and democratic counterparts. In the set of autocrats with legislative removal we find an average audience cost of –12 and states like Syria, China, Vietnam, and Egypt. Finally, among the truest of autocrats there is an average audience cost of only –8 and we find states like Libya (under Qaddafi), Saudi Arabia, and Swaziland.

The regression tree tells a clear and interesting story where democratic characteristics like W and polity2 are important proxies for audience costs, but even in countries that are not democratic there are still institutions that constrain executives. These institutions appear to have an effect on audience costs which can explain why we see both autocrats and democrats mixing in Figure 2. Conversely, this means there is also evidence that democracies with less inclusive institutions are difficult to distinguish from constrained

Table 5: Predicted aggregate transitions

Transition	Predicted transitions	Predicted within state
Peace → Peace	90.1%	95.9%
Peace → Crisis	2.12%	2.26%
Peace → War	1.69%	1.80%
Crisis → Peace	2.27%	79.3%
Crisis → Crisis	0.45%	15.8%
Crisis → War	0.14%	4.94%
War → Peace	1.54%	47.0%
War → Crisis	0.29%	8.96%
War → War	1.44%	44.0%

Caption: The middle column displays the probability distribution over simulated transitions and the far-right column presents the conditional distribution in each state.

autocrats, which supports theories proposed by Chiozza and Goemans (2011) and Weeks (2008, 2012). On average, the tree model tells us that democrats may have higher audience costs than autocrats, but once we condition on regime type, the individual institutions within the country matter in where exactly countries place on the audience cost scale. Future research should be done to find out what the true drivers of audience costs are and what variable or combinations of variables make for good proxies that empirical researchers can use.

8 Model Fit

To assess overall model fit, this section considers the degree to which our model reproduces moments observed in the data but not used in estimation. Although the model is relatively parsimonious, the analysis suggests that it fits the data quite well. This is certainly remarkable given the small numbers of variables included in the baseline model.

When assessing model fit, we examine patterns in aggregate and at the dyadic level. Aggregating across all dyads, our model predicts 93.9% of states should be peace and 2.87% of states should be crisis.[21] This compares to rates of 95.4% and 2.52% observed in the data, respectively. Likewise, Table 5 illustrates our predictions concerning the nine different types of possible transitions and should be compared to Table 1, which reports transitions in the data. Again, the model's predictions fit real-world observations. A notable exception occurs in the war state, and the model over predicts the number of war transitions conditional on the dyads being in war by 8%. Likewise, we over predict peace transitions in

[21] These are based on simulating one path of play of length $100,000$ for each dyad observed in the data.

Table 6: Model fit and comparison at the dyad level

Model	Observed states		Conditional transitions	
	Nulls rejected	Per. incorrect	Nulls rejected	Per. incorrect
Fitted Game	4	4.1%	3	13.0%
Pooled-Average	6	5.0%	7	89.6%

the crisis state by 7%. Nonetheless, the model matches other conditional transitions, and the unconditional distribution essentially mimics the observed one.

We also examine these same statistics at the dyadic level. Specifically, for each dyad, we compare the states observed in the data to the invariant (stationary) distribution that arises from equilibrium play. A dyad's invariant distribution describes the long-run proportion of time spent in each state. To make the comparison, we use a Chi-squared goodness-of-fit test and count the number of dyads in which we reject the null hypothesis that the estimated invariant distribution generates the observed data.[22] In addition, we also compute a percent incorrectly predicted measure and compare our model to a pooled-average model, i.e., one that averages the number of states observed in the data across all dyads.[23] Such a model is equivalent to a naive, all-constants multinomial model over the three states. Second, we repeat the same process with transitions in the observed data, and we examine expected transitions conditional on state from the estimated model and a naive one using similar measures.[24]

Table 6 reports the results, and smaller numbers in all measures indicate better fit. As before, the analysis suggests our model appears to fit the data quite well: We reject the hypothesis that our estimated equilibria generate observed states in roughly 2% of dyads, i.e. 4/179. Likewise, we only reject the null hypothesis that our estimated equilibria generate observed transitions in less than 1% of dyad-state pairs, i.e. $3/(3 \cdot 179)$. Furthermore, there does not appear to be an obvious pattern as to why our model performs poorly on certain dyads. For example, the model under predicts peace in the US-Iraq dyad, and over predicts it in the Israel-Lebanon dyad. Finally, our model out performs the pooled-average model on all four performance criteria. Nonetheless, the pooled-average model is particularly bad at predicting the observed transitions within the data. This poor performance is due the

[22] We use the conventional level of significance and reject the null when $p < .05$.

[23] Consider π a probability distribution over the three states, where $\pi(s)$ is the probability state $s \in \{1, 2, 3\}$. For a dyad k, the percent incorrectly predicted given data Y^k is

$$\frac{1}{T} \sum_s \max \left\{ \sum_{t=1}^{T} \mathbb{I}[s^{kt} = s] - T\pi(s), 0 \right\}.$$

[24] We condition on state because the small number of observed non-peaceful states means all transitions besides Peace $\rightarrow$ Peace are predicted to occur in less than 5/180 observations.

diversity of conditional transitions especially in the crisis and conflict states observed in the data. For example, in some dyads conflict persist, and once states enter the conflict and crisis states they transition to the same state (e.g. Greece and Turkey). In contrast, once some dyads enter the conflict and crisis states, they transitions to peace (e.g. US and Canada).

9 Comparative Statics

In this section, we use the estimated coefficients and equilibria to conduct five counterfactual experiments. These counterfactual experiments illustrate substantive effects (e.g. predicted probability plots) standard for nonlinear models. More precisely, we want to know the effect of changing the data from $\{X, Z\}$ to $\{\tilde{X}, \tilde{Z}\}$ or the parameters from $\hat{\theta}$ to $\tilde{\theta}$ on equilibrium-play between two countries. For example, how does the equilibrium between Lebanon and Israel change as we make Lebanon more democratic or how does the equilibrium between North and South Korea change as we increase the intensity of North Korea's audience costs? Because these counterfactuals involve equilibrium analysis, we encounter two difficulties when computing substantive effects.

First, there exist multiple equilibria. To ensure that our comparative statics describe the estimated equilibrium, we implement a method from Aguirregabiria (2012) that traces equilibria as continuous functions of data or parameters. The key assumptions underlying this procedure and our specific implementation are discussed in the Appendix D. Throughout we focus on specific dyads as we cannot consider the average dyad without information on what equilibrium such a dyad would play.

Second, even after tracing an equilibrium, changes in expected utility do not lend themselves to substantive interpretation.[25] For example, if making Lebanon more democratic encourages Israel to choose peace and Lebanon to choose war, the overall substantive effect on the probability of peace is unclear. To better summarize our equilibria, we use invariant distributions described in the previous section.

Throughout this section, we focus on four theoretically interesting dyads: the United States and Iran, Cyprus and Turkey, Lebanon and Israel, and North and South Korea. In addition, we run two general classes of experiments. The first class consists of three experiments and proceeds as follows. We vary the dyad-specific variables x_{ij} to $\tilde{x}_{ij}$. Next, we trace the estimated equilibria $\hat{v}^k$ to $\tilde{v}^k$ and compute the respective invariant distributions π^{vk} and $\pi^{\tilde{v}k}$. Finally, to determine whether or not the two distributions are statistically different, we use Chi-squared goodness-of-fit tests to determine whether we can reject the null hypothesis that the two distributions produce the same data in expectation. In Ex-

[25] Even conditional choice probabilities in Equation 2 may not better summarize equilibrium play because they reside in a high-dimensional space.

periment 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 the domain of the independent variables.

Table 7: Counterfactual Experiments 1-3.

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

Caption: In Experiment 1, we increase every country's trade dependence by one-half of a standard deviation. In Experiment 2, we increase the dyad's minimum polity score to 10. In Experiment 3, we set each capability ratio to 1. The p-values describe the results of a Chi-squared goodness-of-fit test comparing the estimated invariant distribution to the counterfactual distribution with 180 observations.

Table 7 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.[26] 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 and Israel dyad increases peace by more than 10%, but a similar change with North and South Korea decreases the probability of peace by 6%. To better illustrate the effect of minimum democracy in the Israel-Lebanon dyad, Figure 4 graphs the probability of peace and war as we vary the dyad's minimum

[26] In the data, trade dependence from Lebanon on Israel is 0, and we increase this dependence by approximately 0.01.

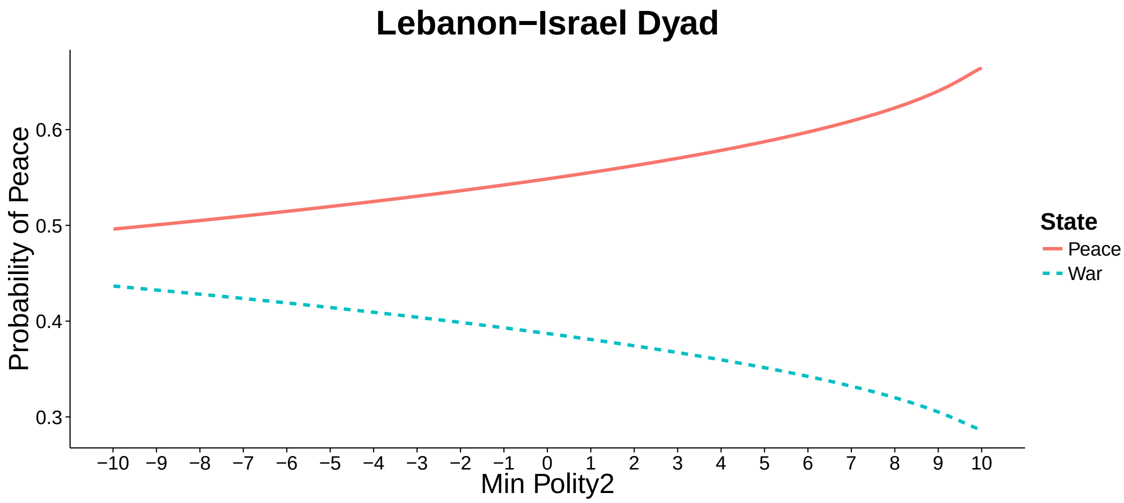

Figure 4: The effect of minimum democracy in the Lebanon-Isreal dyad

polity score from the smallest to the largest possible values.[27] As minimum polity varies from -10 to 10, the probability of peace increases from 50% to 80%, approximately, and the probability of conflict decreases from 45% to 25%, approximately. In contrast, there are no such noticeable changes when looking at the ratio of military capabilities.

The second class of experiments consists of two additional counterfactuals in which we vary the parameters in θ. These experiments resemble familiar comparative-static exercises in theoretical papers where researchers examine how equilibria change as functions of exogenous parameters. Here, we consider the model parametrized by the estimated coefficients $\hat{\theta}$ and data X and Z,and conduct comparative statics on the estimated equilibria $\hat{\mathbf{v}}$. That is, we change θ from our estimated parameter vector, $\hat{\theta}$, to a new vector $\tilde{\theta}$. In the fourth experiment, we reduce the degree to which conflict is mutually reinforcing in the crisis and war states, i.e., we decrease γ^2 and γ^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. In the fifth experiment, we increase every country's audience costs (move α_i toward $-\infty$) by 1.0. In the regression from Table 4, this is similar to increasing a country's executive constraint, as measured by the polity2 score, by two points.

Table 8 reports the results from Experiments 4 and 5. 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 to this we observe that despite making mutual continuance of conflict less profitable the probability of peace

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

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

Table 8: Counterfactual Experiments 4-5.

	Dyad	Invariant Dist. In Data	Invariant Dist. Counterfactual
Experiment 4 (γ^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)**
Experiment 5 (α_i)	United States and Iran	(94.75, 3.81, 1.44)	(94.95, 3.62, 1.43)
	Cyprus and Turkey	(88.69, 9.19, 2.12)	NA
	Lebanon and Israel	(54.86, 6.44, 38.71)	(35.13, 8.32, 56.55)**
	North and South Korea	(86.08, 9.33, 4.59)	(91.70, 5.93, 2.37)*

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

Caption: In Experiment 4, we decrease γ^{WAR} and γ^{CRISIS} by 1.0. In Experiment 5, we decrease α_i by 1.0 for both countries simultaneously in the first two dyads and by 2.0 in the second two dyads. In the first experiment, we could not continuously trace the equilibrium between Lebanon and Israel, and the same problem arose with Cyprus and Turkey in the second.

Turning our attention to audience costs we first note that moving Lebanon and Israel's audience costs α_i toward $-\infty$ by 1.0 leads to a nearly 10% drop in the probability of peace between the two countries. In contrast, a similar shift occurs in the Korean dyad where the probability of peace increases by about 10%. The differences in these results illustrate the two mechanisms through which audience costs influence a country's decision in crisis escalation (Chiozza and Goemans, 2011; Downs and Rocke, 1995). In particular in the Koreas, higher audience costs discourage hostile behavior (either crisis or conflict) because neither country wants to bear the future costs associated with backing down, and in this sense audience costs have a pacifying effect on the dyad. The dark side of audience costs is that they encourage gambling for resurrection among leader who stay in conflicts and continue the destruction because they do not want to back down and pay this cost (LBJ in Vietnam). This could be what we see in the Lebanon – Israel case where we observe increases in the probability of being in crisis and war. These states, not wanting to pay the public cost of backing down, spend more time in conflict.

Finally, Figure 5 graphs the probability of peace in two dyads as functions of the estimated audience costs. In both graphs, the x-axis denotes changes to the relevant countries' audience costs at the same rate. The axis varies across figures because we can only consider values in a non-empty open subset around the original estimated coefficients. In the North

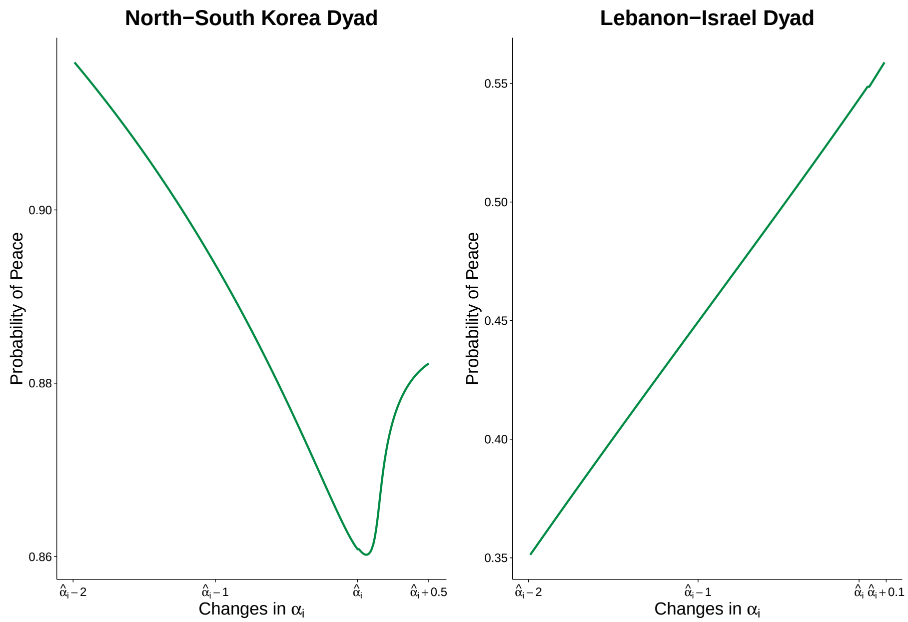

Figure 5: The effect of audience costs in two dyads

and South Korean dyad, audience costs exhibit a non-monotonic relationship with the probability of peace when tracing its equilibrium for a larger set values than in Experiment 5. This non-monotonicity is lost in just comparing starting and final values in the experiments and illustrates that both effects of audience can appear in a single dyad. In contrast, the probability of peace increases as audience decrease (move closer to positive infinity) in the Lebanon-Israel dyad, and this is consistent with Experiment 5.

10 Conclusion

This paper examines the dynamic interaction between pairs of countries as they navigate the interstate system. We explicitly model how countries move into and out of war using a dynamic game and develop a constrained maximum likelihood estimator (CMLE) that accounts for the forward-looking nature of the actors we considered. This approach finds new evidence for a host of theories important to peace scholars. In particular, we uncover new evidence for the pacifying effects of democracy and trade. These two cornerstones of the Kantian peace discourage dyads from entering war, but an interesting caveat is that this pacifying effect does not discourage crises or other low-level disputes, a nuance that has been noted in some past research.

Additionally, we measure how a country's escalation costs change depending on the current state of affairs and its rival's escalation decision. We find evidence for spiraling models of conflict only in cases where countries are already engaged in hostilities, while in peace a country's cost of escalation increases when it expects its opponent to do the same. This explains why peace and hostility tend to be self-enforcing, and why we tend to observe temporal clustering of conflicts.

Furthermore, our approach takes a substantial leap forward in the search for and analysis of audience costs in international politics. First, we uncover believable estimates without imposing strict or questionable functional form assumptions. Second, we analyze their competing effects on peace; for example, audience costs encourage peace in the Korean dyad but discourage peace between Lebanon and Israel. Third, the flexibility of our approach allows us to analyze proxies for audience costs. Democracy explains only part of their variation, and other factors like institutions that constrain chief executives capture previously unidentified components of audience costs. Overall, evidence suggests that the relationship between democracy and audience costs that we take for granted is not necessarily as straightforward as previously thought. Future research should contribute a more nuanced view of audience costs, their inner working, and best proxies.

Our theoretically and empirically unified approach to analyzing international conflict dynamics offers important avenues for future research. First, the estimation procedure can be easily applied to other structural models in political science, and this side-steps issues with multiple equilibria as in crisis-signaling models. Second, the model and estimator are flexible enough to incorporate other actors such as an international system and other states such as nuclear capabilities and diplomatic connections. Third, we provide estimates of audience costs that can be used in future empirical studies as key variables of interest.

Finally and most substantively, our results highlight the complicated relationship between democracy and peace. While a large minimum democracy score in a dyad encourages countries to prefer peace over war, it is unclear as to the partial effect of democracy through audience costs. This paints a darker picture of democracy and its relationship with peace. Future work is needed to measure the overall effect of democracy on peace across dyads, especially as policy practitioners justify the spread of democracy in the name of peace.

APPENDIX

A Implementation

In this Appendix, we detail our implementation of the CMLE because it not only involves a constrained optimization but a potentially a very large number of parameters. For more detailed discussion, see Su and Judd (2012).

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). Interior point methods are particularly useful for two reasons. First, the algorithms find a local solution to optimization problem in Eq. 5. Thus, our estimates are similar to the classical unconstrained MLE results at convergence. Second, this method does not require that the constraints are satisfied in every iteration, just at convergence. This side steps the need to repeatedly compute equilibria as in Rust (1987). Third, the algorithms will account for the possibility that multiple equilibria or solutions exist at $\bar{\theta}$. In time trials, IPOPT had better convergence and performance properties than other optimizers such as KNITRO and a version of the Augmented Lagrangian Method.

A potential drawback of Interior-point methods is that they require accurate representations of the Hessian of the Lagrangian for the problem in Eq. 5. In our experiments, numerical approximations using finite differences substantially inflate standard errors. To work around this, we compute this Hessian and 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 AD routines produce the derivatives by repeatedly applying the chain rule to the supplied code. The result are derivatives within numerical precision and with a similar complexity of the original code. We implement the estimator using Python 2.7 on Xubuntu using the pyipopt software developed by Xu (2014) to use IPOPT within Python and the pyadolc package developed by Walter (2014) to use the AD routines discussed above.

Finally, using the above tools we can efficiently solve the constrained optimization problem, which means we can then simulate standard errors. We use a version of the model-based bootstrap in which one simulates new paths of play using the estimated expected utilities, $\hat{\mathbf{v}}$, and then re-estimates the model. This is identical to what Davison and Hinkley (1997) present in their discussion of time-series data, and they call this a "model-based bootstrap." More specifically, to compute standard errors we employ the following procedure:

1. Estimate $(\hat{\theta}, \hat{v})$ using X, Z, and Y.

2. Using $\hat{v}$, generate a new synthetic time series (path of play) for each dyad in the data; call this data Y^*.

3. Re-estimate the model using X, Z, and Y^*; store $(\hat{\theta}^*, \hat{v}^*)$.

4. Repeat steps 2 and 3 n times.

The output from this procedure produces n simulated draws from the joint sampling distribution of $(\hat{\theta}, \hat{v})$, and this output can be used to construct an estimated variance-covariance matrix of the parameters as well as confidence intervals on substantive effects.

Table 9: Coefficients used in the first Monte Carlo experiment analysing the performance of the CMLE as a function of N and T.

Coefficient	β^2	β^3	κ^2	κ^3	γ^1	γ^2	γ^3	α_i
Value	$(-1, 1)$	$(-2, 2)$	-0.5	-1	-0.5	0	0.5	$-2 + \frac{2(i-1)}{N-1}$

The parametric nature of this process is obvious given that it relies on producing consistent estimates of v in the original model and then resampling data from the assumed Markov process. The advantages of this approach is that it accounts for the time series nature of each game which makes ordinary nonparametric techniques either inappropriate or unclear as to how they should be applied. It's worth noting that there is nothing limiting this process to our type of strategic estimator. In fact we can easily apply them to the strategic logits and probits proposed by Signorino (1999), or indeed any non-strategic estimator where the distribution of the dependent variable is assumed. Some preliminary Monte Carlo evidence with Signorino's strategic probit suggests that this procedure returns results roughly similar to those obtained by a nonparametric bootstrap, and in many cases it produces slightly more conservative estimates than the nonparametric bootstrap.

B 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 data set below.

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 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 9. 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. 3. Next, we generate T periods of data using the computed equilibrium, the associated conditional choice probabilities in Eq. 2, 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 burnin period, we combine the generated data and estimate the parameters of the model by solving the constrained optimization problem in Eq. 5 using the tools described in the previous section. The procedure is repeated 50 times for each value of N and T.

Table 10: Summary of Monte Carlo Experiment

N	T	$E[\hat{\beta}_1^2]$	$E[\hat{\beta}_2^2]$	$E[\hat{\beta}_1^3]$	$E[\hat{\beta}_2^3]$	$E[\hat{\kappa}_1^2]$	$E[\hat{\kappa}_1^3]$	$E[\hat{\gamma}^1]$	$E[\hat{\gamma}^2]$	$E[\hat{\gamma}^3]$
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 10 reports the means and standard errors in parentheses for the parameters β^s,

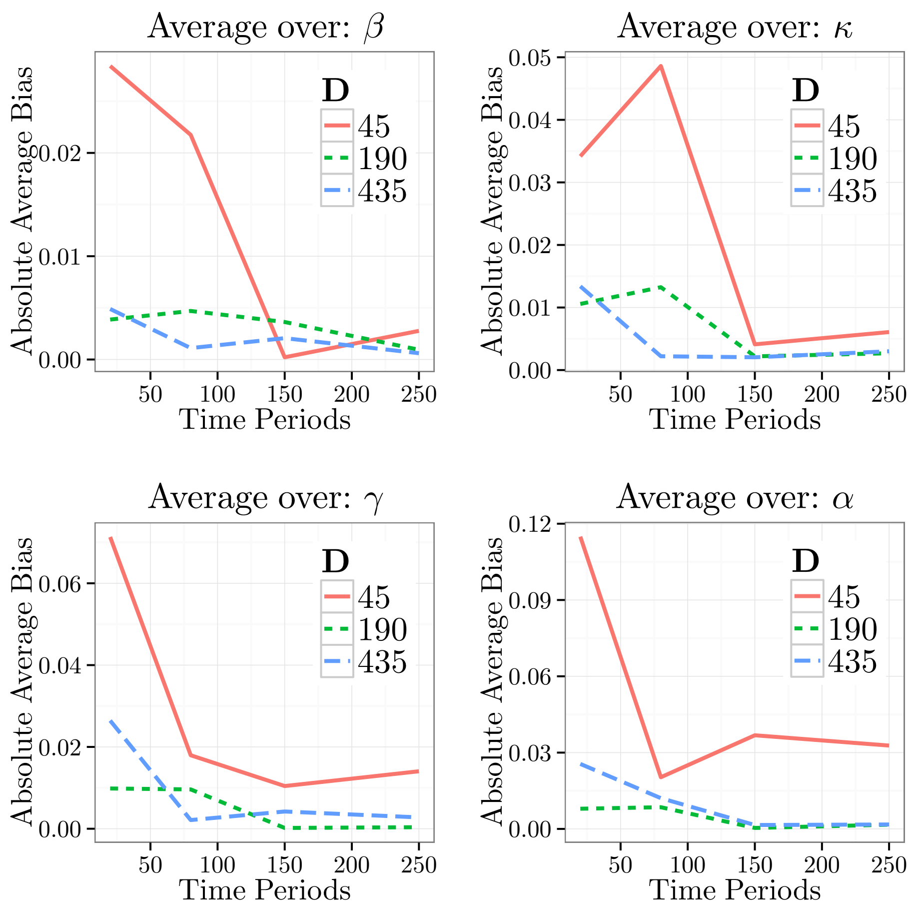

Figure 6: The average bias of the constrained ML estimator by four different types of coefficients as functions of the number of countries N and time periods T. In the analysis, $D = \binom{N}{2}$. Note that the average bias over β is $\frac{1}{4}\sum_{s=2}^{3}\sum_{j=1}^{2}|\beta_j^s - \mathrm{E}[\hat{\beta}_j^s]|$.

κ^{a_i}, and γ^s. Due to space concerns, we do not report the audience cost parameters. In addition, Figures 6 and 7 summarize the results, Here we graph the CMLE's bias averaged over the four different sets of coefficients. More specifically, to produce the upper-left graph of Figure 6, we first compute the expected bias of $\hat{\beta}_l^s$ for each $s = 1, 2$ and each $l = 1, 2$, and then we averaged these values for each specification of N and T. The upper-left graph of Figure 7 averages the variance of $\hat{\beta}_l^s$. The remaining graphs are produced in a similar manner, and Appendix B contains more detailed results with coefficient estimates and associated standard errors. Most importantly, the bias and the variance of the constrained ML estimator decreases as we increase N and T. This monotonic relationship is especially pronounced with the estimator's variance. Even though increasing N means estimating an additional audience cost parameter, 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 period in the observation may increase the estimator's bias, especially concerning the action-specific cost parameters, κ^s. 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. In contrast, when $T > 20$, the convergence rate is 100%. In addition, Figure 8 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 parameters.) Nonetheless, even with 435 dyads the average estimation time in Monte Carlos is less than 3 hours.

B.1 Comparing Random Utility Models

In this section, we conduct an additional Monte Carlo experiment in which we compare our estimation procedure to a standard logit and multinomial logit. We chose these models because they form similar choice probabilities and are common in the literature (e.g. Clare, 2007; Huth and Allee, 2002). This section's results indicate that specifying the estimation procedure without incorporating the long-term dynamics of crisis escalation can generate incorrect inferences concerning the size and direction of the coefficients of interest.

Specifically, we fix $N = 10$, $\mathcal{D} = \{\{i, j\} \mid i, j \in \{1, ..., N\}, i \neq j\}$, and $T = 250$. We chose parameter values in Table 11, and use an identical process to generate the data as in the previous section. The major difference between Tables 9 and 11 is that in the latter we do not vary the audience cost parameters, α_i, and now the structural parameters, e.g. γ^s, are larger in magnitude. Besides the CMLE, we attempt to estimate the coefficients using a multinomial logit, in which we assume, albeit incorrectly, that the probability of country

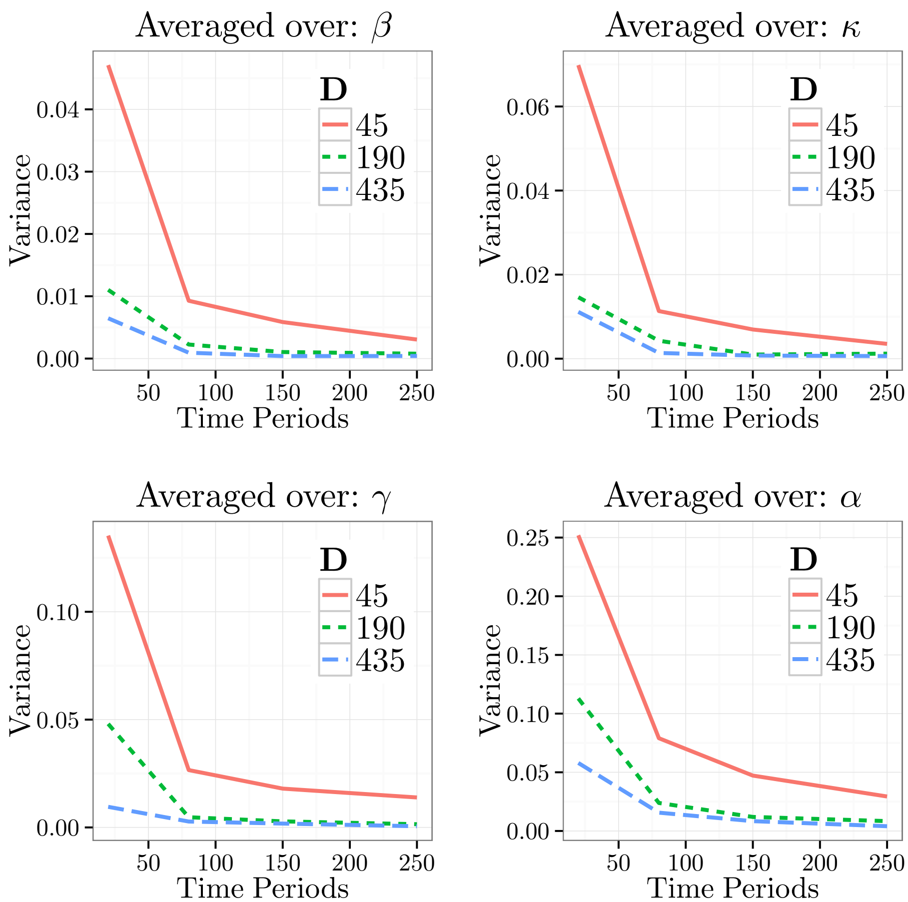

Figure 7: The average variance of the constrained ML estimator by four different types of coefficients. In the analysis, $D = \binom{N}{2}$. Note that the average bias over β is $\frac{1}{4}\sum_{s=2}^{3}\sum_{j=1}^{2} \mathrm{Var}[\hat{\beta}_j^s]$.

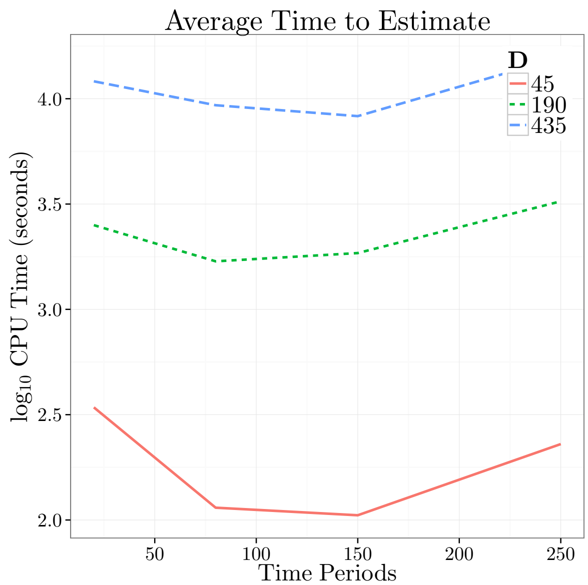

Figure 8: CPU time of the CMLE.

Table 11: Coefficients used in the second Monte Carlo experiment analysing the performance of the CMLE compared to a logit and multinomial logit.

Coefficient	β^2	β^3	κ^2	κ^3	γ^1	γ^2	γ^3	α_i
Value	$(-3, 2)$	$(-6, -2)$	-1	-0.5	-2	2	4	$-5 + \frac{10(i-1)}{9}$

i chooses action a_i against country j is

$$P_{ij}(a_i) = \frac{\exp\left(\beta^{a_i} \cdot x_{ij} + \kappa^{a_i} \cdot z_i\right)}{\sum_{a_i'} \exp\left(\beta^{a_i'} \cdot x_{ij} + \kappa^{a_i'} \cdot z_i\right)}.$$

We also estimate the model using a logit in which we assume that the probability the i chooses action $a_i > 1$ against country j is

$$P_{ij}(a_i > 1) = \frac{\exp\left(\beta^2 \cdot x_{ij} + \kappa^2 \cdot z_i\right)}{1 + \exp\left(\beta^2 \cdot x_{ij} + \kappa^2 \cdot z_i\right)}.$$

Notice that using a multinomial logit permits us to estimate $L(W-1)$ coefficients, but using a logit, we only estimate L coefficients.

Table 12 reports the results of this analysis. The most important thing to note is that, unlike the CMLE, the other regression models cannot adequately uncover the data generating process. This leads to significant incorrect inferences: the non-strategic models cannot estimate the correct directions of some coefficients. For example, both logits estimate

β_2^2 with the opposite sign as the data generating process, which leads to large root Mean-squared errors when compared to the CMLE. While it is quite easy to generate data with one model and find that other models cannot as readily uncover the original parameters, the model misspecification bias can become large enough that even the signs of the coefficients are not correctly recovered. Furthermore, notice there exists a relationship between the CMLE and multinomial logit models in special cases. When $\delta = 0$, $\gamma^s = 0$ for all states s, and $\alpha_i = 0$ for all countries i, that is, where there is no present or future strategic interactions, a multinomial logit will correctly estimate the κ^{a_i} parameters.

Table 12: Comparison of three random utility models. The first three columns report coefficient estimates and their associated standard errors in parentheses. The last three columns report the root mean-squared error.

	Estimates			Root Mean-Squared Error		
Coefficients	Constrained MLE	Multinomial Logit	Logit	Constrained MLE	Multinomial Logit	Logit
β_1^2	-3.06 (0.07)	-0.65 (0.62)	0.10 (0.45)	0.09	2.43	3.13
β_2^2	1.99 (0.04)	-1.21 (0.75)	-0.44 (0.21)	0.05	4.84	2.45
β_1^3	-5.93 (0.13)	0.38 (0.27)		0.14	1.64	
β_2^3	-1.98 (0.09)	-1.41 (0.24)		0.09	0.64	
κ_1^2	-0.93 (0.08)	-0.79 (0.59)	-0.56 (0.35)	0.44	0.63	0.56
κ_1^3	-0.36 (0.09)	-0.36 (0.62)		0.28	0.63	

C 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 these 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.

Figure 9: Country-Year Deviations from Mean Values of Independent Variables

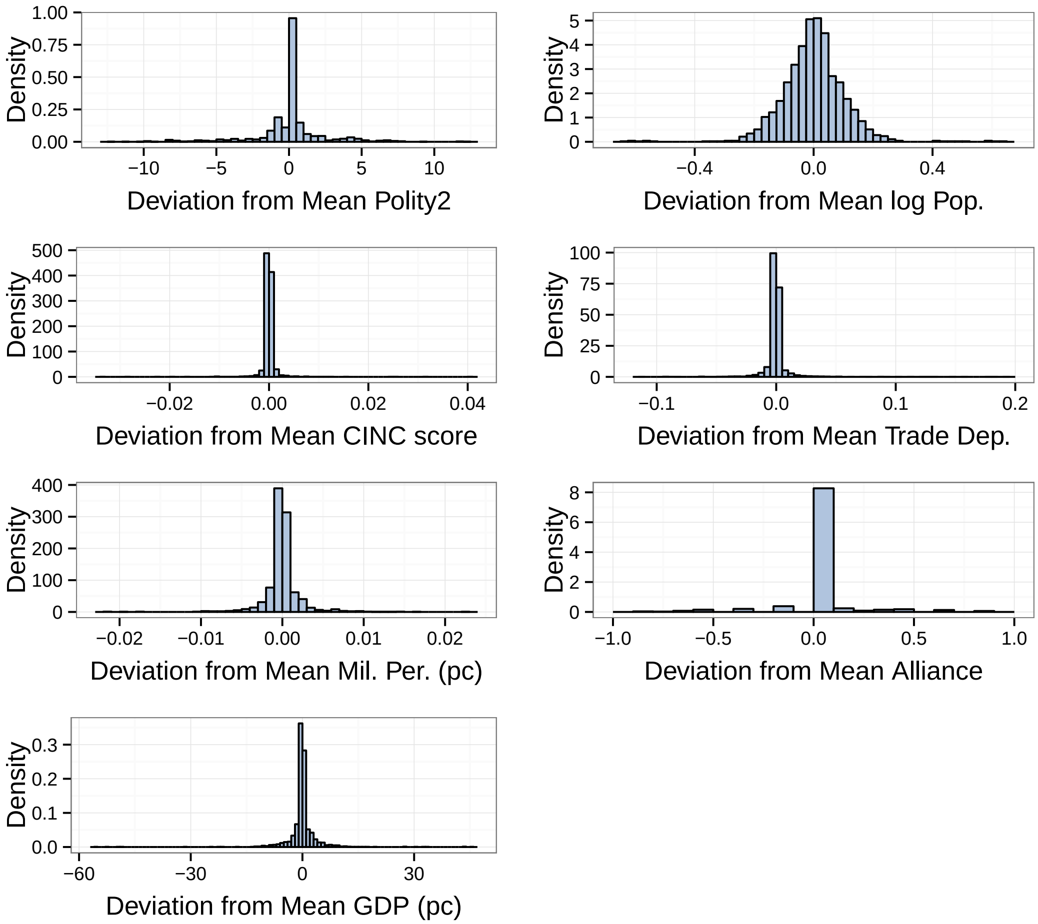

First, we first examine whether our independent variables change over time. Our independent variables come in two types, and the first type varies by country. These variables include Polity2, Military Personnel per capita, GDP per capita, CINC score and population. For each variable, we compute its means 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 9 displays histograms of these deviations for each variable and illustrates that observed deviations from the mean are relatively small across our data set.

Second, we then attempt to explain our independent variables using observed states and actions in a panel data analysis. Specifically, we regress these variables on the number of conflictual 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 13 display regression results with country-year observations, and Models 5-6 displays a similar regressions with directed dyad-year observations.[28] The main take-away 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.

[28]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 13: 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†	0.00†	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

D Comparative Statics

In this Appendix, we detail how to compute comparative statics in the model using a simple homotopy predictor-corrector method. 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 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. 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 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. 3 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 get a continuous representation of the equilibrium, which we use to graphically verify that the output has indeed continuously traced the equilibrium.

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 $\tilde{v}^k$

E Audience Cost Parameters

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

Table 14: Audience Cost Estimates

Country	Audience Cost	St. Err.	p
Nepal	-0.87	8.54	0.92
Bangladesh	-1.86	1.29	0.15
Nigeria	-2.18	2.23	0.33
Cambodia	-2.32	7.19	0.75
Lebanon	-2.43	0.91	< 0.01
Benin	-2.46	9.45	0.79

Azerbaijan	-2.64	0.99	< 0.01
Armenia	-2.90	1.15	0.01
Pakistan	-3.61	0.94	< 0.01
India	-3.80	0.95	< 0.01
Cameroon	-4.03	2.70	0.13
Djibouti	-4.52	9.87	0.65
Saudi Arabia	-4.75	2.42	0.05
Syria	-5.04	5.06	0.32
Myanmar (Burma)	-5.17	2.42	0.03
Yemen	-5.25	2.89	0.07
Iraq	-5.62	0.89	< 0.01
Afghanistan	-7.91	2.09	< 0.01
Albania	-8.95	1.56	< 0.01
Greece	-9.22	0.95	< 0.01
Eritrea	-9.37	5.70	0.10
Turkey	-9.43	0.89	< 0.01
Korea North	-9.54	1.46	< 0.01
Thailand	-9.59	2.80	< 0.01
Korea South	-9.74	1.63	< 0.01
Bosnia	-10.04	2.85	< 0.01
Vietnam	-10.61	6.62	0.11
Georgia	-10.90	4.27	0.01
Israel	-11.09	1.93	< 0.01
Japan	-11.33	3.26	< 0.01
Bulgaria	-11.42	6.60	0.08
Latvia	-11.72	7.70	0.13
Qatar	-11.97	6.81	0.08
Poland	-12.04	7.73	0.12
Jordan	-12.08	6.00	0.04
Yugoslavia	-12.83	1.23	< 0.01
Congo Kinshasa	-13.37	2.01	< 0.01
Croatia	-13.39	2.18	< 0.01
Chile	-13.52	6.40	0.03
South Africa	-13.59	14.48	0.35
Romania	-14.08	3.84	< 0.01
Tajikistan	-14.24	5.21	< 0.01
Iran	-14.26	2.44	< 0.01

Ukraine	-14.29	2.76	< 0.01
Uganda	-14.50	1.85	< 0.01
Zimbabwe	-14.60	7.14	0.04
Kuwait	-14.62	2.59	< 0.01
Sudan	-14.75	1.53	< 0.01
Mozambique	-14.79	9.95	0.14
Cuba	-14.80	4.92	< 0.01
Lithuania	-14.88	4.19	< 0.01
Ethiopia	-14.90	2.44	< 0.01
Malaysia	-14.99	2.48	< 0.01
Uzbekistan	-15.07	4.62	< 0.01
Morocco	-15.14	5.56	< 0.01
Moldova	-15.15	1.81	< 0.01
Central African Republic	-15.17	3.55	< 0.01
Chad	-15.26	1.93	< 0.01
China	-15.29	1.90	< 0.01
Kenya	-15.40	3.56	< 0.01
Australia	-15.41	6.21	0.01
Germany	-15.58	2.06	< 0.01
Namibia	-15.73	16.46	0.34
Mongolia	-15.79	2.54	< 0.01
Egypt	-15.79	2.50	< 0.01
Venezuela	-15.80	1.93	< 0.01
Burundi	-16.10	3.10	< 0.01
Mauritania	-16.12	6.15	< 0.01
Singapore	-16.20	5.45	< 0.01
Russia	-16.22	1.18	< 0.01
Belgium	-16.29	4.99	< 0.01
United Kingdom	-16.30	2.03	< 0.01
El Salvador	-16.33	10.03	0.10
Guyana	-16.35	4.64	< 0.01
Papua New Guinea	-16.42	2.41	< 0.01
Portugal	-16.43	4.68	< 0.01
Zambia	-16.45	2.30	< 0.01
Indonesia	-16.46	1.48	< 0.01
Netherlands	-16.47	4.73	< 0.01
Slovenia	-16.48	2.26	< 0.01

Nicaragua	-16.55	3.97	< 0.01
Canada	-16.76	2.65	< 0.01
Bhutan	-16.89	5.48	< 0.01
United States	-16.92	1.13	< 0.01
Philippines	-17.00	3.88	< 0.01
Angola	-17.07	1.81	< 0.01
Kyrgyzstan	-17.13	4.96	< 0.01
Haiti	-17.22	1.85	< 0.01
Spain	-17.23	3.34	< 0.01
Hungary	-17.23	2.69	< 0.01
Liberia	-17.27	1.88	< 0.01
Costa Rica	-17.55	5.90	< 0.01
Libya	-17.58	4.20	< 0.01
Botswana	-17.69	15.83	0.26
Belarus	-17.81	4.63	< 0.01
Tanzania	-17.81	1.97	< 0.01
Rwanda	-17.89	2.46	< 0.01
France	-17.92	1.14	< 0.01
Algeria	-17.94	4.01	< 0.01
Italy	-18.15	1.94	< 0.01
Sierra Leone	-18.18	2.61	< 0.01
Ghana	-18.31	15.06	0.22
Gambia	-18.47	11.18	0.10
Congo Brazzaville	-18.52	4.47	< 0.01
Argentina	-18.54	1.72	< 0.01
Mali	-18.63	2.03	< 0.01
Togo	-18.66	11.58	0.11
Niger	-18.71	1.96	< 0.01
Denmark	-18.75	3.28	< 0.01
Turkmenistan	-18.84	4.64	< 0.01
Swaziland	-18.99	9.95	0.06
Ivory Coast	-19.28	1.77	< 0.01
Peru	-19.34	4.52	< 0.01
Lesotho	-19.41	14.18	0.17
Ecuador	-19.42	2.68	< 0.01
Brazil	-19.43	3.74	< 0.01
Guinea	-19.69	1.75	< 0.01

Colombia	-19.88	3.35	< 0.01
Somalia	-19.91	6.37	< 0.01
Senegal	-19.92	8.39	0.02
Honduras	-19.94	3.89	< 0.01
Suriname	-20.22	4.78	< 0.01
Norway	-21.47	2.34	< 0.01
Dominican Rep	-22.34	3.38	< 0.01
Cyprus	-33.87	4.33	< 0.01

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.

Axelrod, Robert. 1984. *The Evolution of Cooperation.* New York: Basic Books.

Barbieri, Katherine. 1996. "Economic Interdependence: A Path to Peace or a Source of Interstate Conflict." *Journal of Peace Research* 33(1):29–49.

Barbieri, Katherine, Omar M.G. Keshk and Brian Pollins. 2009. "TRADING DATA: Evaluating our Assumptions and Coding Rules." *Conflict Management and Peace Science* 26(5):471–491.

Bas, Muhammet A., Curtis S. Signorino and Taehee Whang. 2014. "Knowing Ones Future Preferences: A Correlated Agent Model with Bayesian Updating." *Journal of Theoretical Politics* 26(1):3–34.

Beck, Nathaniel, Jonathan Katz and Richard Tucker. 1998. "Taking Time Seriously: Time-Series-Cross-Sectional Analysis with a Binary Dependent Variable." *American Journal of Political Science* 42(4):1260–1288.

Beck, Nathaniel and Jonathan N. Katz. 2001. "Throwing Out the Baby with the Bath Water: A Comment on Green, Kim, and Yoon." *International Organization* 55(2):487–495.

Braumoeller, Bear F. 2008. "Systemic Politics and the Origins of Great Power Conflict." *American Political Science Review* 102(01):77–93.

Bueno de Mesquita, Bruce, Alastair Smith, Randolph M. Siverson and James D. Morrow. 2005. *The Logic of Political Survival.* Cambridge, MA: MIT Press.

Bueno De Mesquita, Bruce and David Lalman. 1992. *War and reason: Domestic and International Imperatives*. New Haven: Yale University Press.

Bueno De Mesquita, Bruce, James D. Morrow and Ethan R. Zorick. 1997. "Capabilities, perception, and escalation." *American Political Science Review* 91(1):15–27.

Bueno de Mesquita, Bruce, James D. Morrow, Randolph M. Siverson and Alastair Smith. 1999. "An Institutional Explanation of the Democratic Peace." *American Political Science Review* 93(4):791–807.

Carter, David B. 2010. "The Strategy of Territorial Conflict." *American Journal of Political Science* 54(4):969–987.

Chiozza, Giacomo and H.E. Goemans. 2011. *Leaders and International Conflict*. New York: Cambridge University Press.

Clare, Joe. 2007. "Domestic Audiences and Strategic Interests." *Journal of Politics* 69(3):732–745.

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.

Debs, Alexandre and Nuno P. Monteiro. 2014. "Known Unknowns: Power Shifts, Uncertainty, and War." *International Organization* 68(01):1–31.

Dorussen, Han and Jongryn Mo. 2001. "Ending Economic Sanctions Audience Costs and Rent-Seeking as Commitment Strategies." *Journal of Conflict Resolution* 45(4):395–426.

Downs, George W. and David M. Rocke. 1995. *Optimal Imperfection?* Princeton University Press.

Egesdal, Michael, Zhenyu Lai and Che-Lin Su. 2013. "Estimating Dynamic Discrete-Choice Games of Incomplete Information." Unpublished manuscript.

Esarey, Justin, Bumba Mukherjee and Will H. Moore. 2008. "Strategic Interaction and Interstate Crises: A Bayesian Quantal Response Estimator for Incomplete Information Games." *Political Analysis* 16(3):250–273.

Eyerman, Joe and Robert A. Hart. 1996. "An Empirical Test of The Audience Cost Proposition Democracy Speaks Louder than Words." *Journal of Conflict Resolution* 40(4):597–616.

Fearon, James D. 1994. "Domestic Political Audiences and the Escalation of International Disputes." *American Political Science Review* 88(3):577–592.

Fearon, James D. 2004. "Why Do Some Civil Wars Last So Much Longer than Others?" *Journal of Peace Research* 41(3):275–301.

Fearon, James D. 2013. "Fighting Rather than Bargaining." Unpublished manuscript.

Feenstra, Robert C., Robert Inklaar and Marcel P. Timmer. N.d. "The Next Generation of the Penn World Table." available for download at www.ggdc.net/pwt.

Gartzke, Erik. 1998. "Kant We All Just Get Along? Opportunity, Willingness, and the Origins of the Democratic Peace." *American Journal of Political Science* 42(1):1–27.

Ghosn, Faten, Glenn Palmer and Stuart Bremer. 2004. "The MID3 Data Set, 1993-2001: Procedures, Coding Rules, and Description." *Conflict Management and Peace Science* 21(2):133–154.

Gleditsch, Kristian S. 2002. "Expanded Trade and GDP Data." *Journal of Conflict Resolution* 46:712–724.

Gleditsch, Kristian S. and Michael D. Ward. 2000. "War and Peace in Space and Time: The Role of Democratization." *International Studies Quarterly* 44(1):1–29.

Goemans, H.E. 2000. *War and Punishment: The Causes of War Termination and the First World War*. Princeton 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.

Jackman, Simon. 2000. "In and Out of War and Peace: Transitional Models of International Conflict." Unpublished manuscript.

Jackson, Matthew O. and Massimo Morelli. 2009. "Strategic Militarization, Deterrence and Wars." *Quarterly Journal of Political Science* 4(4):279–313.

Jensen, Nathan M. 2003. "Democratic Governance and Multinational Corporations: Political Regimes and Inflows of Foreign Direct Investment." *International organization* 57(03):587–616.

Jo, Jinhee. 2011. "Nonuniqueness of the Equilibrium in Lewis and Schultz's Model." *Political Analysis* 19(3):351–362.

Kenwick, Michael R., Matthew Lane, Benjamin Ostick and Glenn Palmer. 2013. "Codebook for the Militarized Interstate Incident Data, version 4.0." Unpublished manuscript.

Kurizaki, Shuhei and Taehee Whang. 2014. "Detecting Audience Costs in International Crises." Unpublished manuscript.

Leemann, Lucas. 2014. "Strategy and Sample Selection: A Strategic Selection Estimator." *Political Analysis* 22(3):374–397.

Lewis, Jeffrey B. and Kenneth A. Schultz. 2003. "Revealing Preferences: Empirical Estimation of a Crisis Bargaining Game with Incomplete Information." *Political Analysis* 11(4):345–367.

Li, Quan. 2005. "Does Democracy Promote or Reduce Transnational Terrorist Incidents?" *Journal of Conflict Resolution* 49(2):278–297.

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.

McKelvey, Richard D. and Thomas R. Palfrey. 1995. "Quantal Response Equilibria for Normal Form Games." *Games and Economic Behavior* 10(1):6–38.

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

Partell, Peter J. and Glenn Palmer. 1999. "Audience Costs and Interstate Crises: An Emprical Assessment of Fearon's Model of Dispute Outcomes." *International Studies Quarterly* 43(2):389–405.

Pesendorfer, Martin and Philipp Schmidt-Dengler. 2008. "Asymptotic Least Squares Estimators for Dynamic Games." *The Review of Economic Studies* 75(3):901–928.

Powell, G. Bingham. 2000. *Elections as Instruments of Democracy*. New Haven: Yale University Press.

Powell, Robert. 1999. *In the Shadow of Power*. Princeton: Princeton University Press.

Prins, Brandon C. 2003. "Institutional Instability and the Credibility of Audience Costs: Political Participation and Interastate Crisis Bargaining, 1816-1992." *Journal of Peace Research* 40(1):67–84.

Regan, Patrick M., Richard W. Frank and David H. Clark. 2009. "New Datasets on Political Institutions and Elections, 1972–2005." *Conflict Management and Peace Science* 26(3):286–304.

Rust, John. 1987. "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher." *Econometrica* 55(5):999–1033.

Schelling, Thomas C. 1960. *The Strategy of Conflict*. New York: Oxford University Press.

Schultz, Kenneth A. 1998. "Domestic Opposition and Signaling in International Crises." *American Political Science Review* 92(4):829–844.

Schultz, Kenneth A. 2001. "Looking for Audience Costs." *Journal of Conflict Resolution* 45(1):32–60.

Senese, Paul D. 1997. "Between Dispute and War: The Effect of Joint Democracy on Interstate Conflict Escalation." *The Journal of Politics* 59(1):1–27.

Signorino, Curtis S. 1999. "Strategic Interaction and the Statistical Analysis of International Conflict." *American Political Science Review* 93(2):279–297.

Signorino, Curtis S and Ahmer Tarar. 2006. "A Unified Theory and Test of Extended Immediate Deterrence." *American Journal of Political Science* 2(3):586–605.

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

Singer, J. David, Stuart A. Bremer and John Stuckey. 1972. Capability Distribution, Uncertainty, and Major Power War, 1820-1965. In *Peace, War, and Numbers*, ed. Bruce Russett. Beverly Hills: Sage Publications.

Slantchev, Branislav L. 2003. "The Principle of Convergence in Wartime Negotiations." *American Political Science Review* 97(4):621–632.

Slantchev, Branislav L. 2005. "Military Coercion in Interstate Crises." *American Political Science Review* 99(4):533–547.

Slater, Dan. 2003. "Iron Cage in an Iron Fist: Authoritarian Institutions and the Personalization of Power in Malaysia." *Comparative Politics* 36(1):81–101.

Spaniel, William. 2014. "Arms Treaties and the Credibility of Preventive War." Unpublished manuscript.

Su, Che-Lin and Kenneth L. Judd. 2012. "Constrained Optimization Approaches to Estimation of Structural Models." *Econometrica* 80(5):2213–2230.

Venables, W. N. and B. D. Ripley. 2003. *Modern Applied Statistics with S.* 4th ed. New York: Springer.

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(01):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.

Whang, Taehee, Elena V. McLean and Douglas W Kuberski. 2013. "Coercion, Information, and the Success of Sanction Threats." *American Journal of Political Science* 57(1):65–81.

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