

Estimating Signaling Games in International Relations: Problems and Solutions *

Casey Crisman-Cox[†] Michael Gibilisco[‡]

August 5, 2019

*Thanks to Rob Carroll, Kentaro Fukumoto, Jinhee Jo, Gleason Judd, Tasos Kalandrakis, James Lo, Gabriel Lopez-Moctezuma, Sergio Montero, Jacob Montgomery, Will Moore, Matt Shum, Wei Zhong, the editor Jude Hayes, and two anonymous referees for comments and suggestions. This paper benefited from audiences at Washington University in Saint Louis, the Southern California Methods Workshop, the International Methods Colloquium and the annual meetings of APSA, MPSA, and the Society for Political Methodology. Naturally, we are responsible for all errors.

[†]Corresponding author, Texas A&M University. Email: c.crisman-cox@tamu.edu

[‡]California Institute of Technology. Email: michael.gibilisco@caltech.edu

Abstract

Signaling games are central to political science but often have multiple equilibria, leading to no definitive prediction. We demonstrate that these indeterminacies create substantial problems when fitting theory to data: they lead to ill-defined and discontinuous likelihoods *even if* the game generating the data has a unique equilibrium. In our experiments, currently used techniques frequently fail to uncover the parameters of the canonical crisis-signaling game, regardless of sample size and number of equilibria in the data generating process. We propose three estimators that remedy these problems, outperforming current best practices. We fit the signaling model to data on economic sanctions. Our solutions find a novel *U*-shaped relationship between audience costs and the propensity for leaders to threaten sanctions, which current best practices fail to uncover.

1 Introduction

Political scientists use signaling games to analyze phenomena across practically all subfields. Scholars of international relations in particular use the models to address questions about economic sanctions (Lacy and Niu [2004](#); Lektzian and Sprecher [2007](#)), crisis bargaining (Tarar and Leventoğlu [2009](#)), escalation in interstate disputes (Schultz [1999](#)), and terrorism (Bueno de Mesquita and Dickson [2007](#); Lapan and Sandler [1993](#)). As a result of this ubiquity, scholars structurally estimate increasingly more complicated signaling models (Bas, Signorino and Whang [2014](#); Kurizaki and Whang [2015](#); Lewis and Schultz [2003](#); Wand [2006](#); Whang [2010a](#); Whang, McLean and Kuberski [2013](#)). Advocated by the movement for empirical implications of theoretical models, the structural approach allows researchers to account for strategic interdependence in the data generating process, estimate theoretical parameters of interest, and conduct counterfactual policy analysis in the absence of experimental conditions.

Despite these benefits, political scientists still face substantial theoretical and computational hurdles when estimating signaling games. In these games, each player knows her private information at the beginning of the interaction and behavior is characterized by perfect Bayesian equilibria. The most pressing problem is how to build a coherent empirical signaling model that smooths out issues arising from the multiplicity of equilibria common to these games. In this paper, we address this problem by adapting three techniques from the dynamic games and industrial organization literatures (e.g., De Paula [2013](#); Ellickson and Misra [2011](#)) to estimate the canonical crisis-signaling model in Lewis and Schultz ([2003](#)). We demonstrate that they outperform current best practices—in terms of statistical performance *and* computational feasibility. Through a series of experiments and applications, we argue

that these solutions are well suited for the simpler, but far more influential, models in political science.

Current best practices for estimating the crisis-signaling model use variants of the maximum likelihood (ML) routine proposed by Signorino (1999) to estimate the parameters of extensive-form games with quantal-response equilibria (QRE). In these best practices, a characterization of the game’s perfect Bayesian equilibria is used to derive a likelihood function for the observed data. Then a numerical optimizer maximizes this likelihood function by computing an equilibrium for every observation at every guess of the parameters. While straightforward, the procedure sidesteps a substantial problem in practice: an equilibrium is computed as if it is unique. Unlike the QRE models in McKelvey and Palfrey (1998) and Signorino (1999), multiple perfect Bayesian equilibria may exist in the crisis-signaling game under reasonable payoff parameters. This multiplicity creates an indeterminacy in the likelihood function, leading to inconsistent estimates (Jo 2011). Hereafter, we call the ML routines that ignore multiplicity “traditional” ML (tML), reflecting current practices (e.g., Bas, Signorino and Whang 2014; Kurizaki and Whang 2015; Whang, McLean and Kuberski 2013; Zeng 2019).

Past justifications for estimating crisis-signaling games with the tML routine rely on either using refinements to reduce the number of equilibria (e.g., Jo 2011) or verifying equilibrium uniqueness at the point estimates while ignoring multiplicity during estimation (e.g., Bas, Signorino and Whang 2014). We show that neither adequately solves the problem. Regarding the former, we prove formally that all equilibria of the crisis-signaling game almost always satisfy the regularity refinement, one of the most stringent in the literature (van Damme 1996). That is, equilibria are equally robust

to standard refinements.¹ Nonetheless, researchers could adopt an *ad hoc* selection rule (e.g., select the equilibrium maximizing the likelihood), but we show that this approach generates discontinuities in the tML’s likelihood function. Furthermore, the number of discontinuities can grow as the sample size increases. Thus, selection rules not only necessitate extraneous computations to identify the equilibrium of interest at every optimization step, but they also require maximization of discontinuous objective functions. These computational complexities dramatically reduce the tML’s already poor feasibility. Indeed, several scholars have abandoned the structural enterprise for reduced-form alternatives citing feasibility concerns (Gleditsch et al. 2018; Trager and Vavreck 2011).²

Regarding the latter, likelihood functions may be evaluated at parameter values under which multiple equilibria exist even if there is a unique equilibrium in the game generating the data. For example, optimization routines often take incorrect guesses at the parameters as they search for the ML estimates. As such, the routines may potentially evaluate the likelihood function at parameters under which multiple equilibria exist even if there is a unique equilibrium at the true parameters in the data generating process. This indeterminacy at incorrect parameter values allows the likelihood function to be evaluated incorrectly and leads to the same discontinuities discussed above, making it difficult to find the correct values. As such, we find that tML routines demonstrate consistently poor performance across a variety of experimental settings, regardless of sample size, use of global optimizers, or the number of equilibria in the game generating the data.

In contrast, we treat equilibrium selection as an empirical problem by allowing

¹Refinements on off-the-path beliefs such as Divinity or the Intuitive Criterion are inconsequential as all histories are reached with positive probability in all equilibria. We also investigate best-response stability, which does not guarantee uniqueness, but does eliminate some equilibria.

²Gleditsch et al. (2018) refer to the Lewis and Schultz (2003) model as “demanding” in justifying their alternative approach.

it to depend on observables. Indeed, having multiple equilibria allows the empirical model greater flexibility in matching real-world interactions. Furthermore, our solutions accommodate empirical selection in a manner that smooths out the issues created by multiple equilibria. Specifically, they rely on the observation that fixing the equilibrium beliefs to their true values when computing best responses removes the indeterminacies in the likelihood without generating discontinuities. Of course, these equilibrium quantities are unobserved, so our proposed solutions rely on estimating them. Specifically, we begin with the assumption that equilibrium strategies, and hence beliefs, can be inferred from observables, either because we observe several interactions from the same equilibrium or because dyads with similar covariates play similar equilibria. For an example in the international relations context, the latter suggests that countries with high levels of trade likely play the game similarly to each other but differently from non-trading countries. Estimating the equilibrium strategies in a first stage and using them in place of their true values in a second stage provides a feasible pseudo-likelihood (PL) solution to the problem of estimating the game's parameters of interest.

While relatively innocuous in principle, this approach requires accurate estimates of equilibrium quantities. We therefore introduce two additional methods to alleviate this reliance on first-stage estimates. The first is a nested-PL (NPL) approach that uses the PL estimates to update actors' beliefs which were estimated in the first stage, allowing the analyst to then update the payoff parameters. The process is iterated until convergence, making the final estimates less dependent on the initial guesses of the equilibrium strategies. The second approach is to estimate equilibrium beliefs as dyad-specific (game-specific) parameters in a single-stage constrained-ML estimator (CMLE) that maximizes a likelihood function while forcing the estimates to satisfy equilibrium conditions. While this approach does not require initial estimates of the

equilibrium beliefs, it does require panel-like data wherein we assume that each dyad plays from the same equilibrium every time it interacts. In other words, each time the U.S. and China interact, they play the same equilibrium—which may be very different than how the U.S. and Canada interact.³

All three of our proposed estimators outperform the tML by reducing variance and bias by orders of magnitude. Specifically, the CML is almost always the best performer, but it is also the most difficult to implement. The PL and NPL are very easy to implement and both work very well in a variety of settings. We also provide an R package to fit crisis-signaling models using the PL and the NPL.

By studying the widely used crisis-signaling model, this paper advances our understanding about the challenges that arise when connecting theory to data. More broadly, we demonstrate that theoretical issues such as equilibrium multiplicity, although often cast as a nuisance to be refined away, have important consequences when fitting models to data. Sidestepping these issues can result in mistaken substantive conclusions. While we focus on a specific game that holds a prominent place in international relations, identical problems arise in other games with multiple equilibria, e.g., games with simultaneous moves or infinitely repeated interactions. Our analysis should therefore encourage political scientists to structurally estimate a wider array of models.

Our empirical application uses the crisis signaling game to study the strategic incentives of sanction threats and impositions (as in Drezner 1998, 2003; Whang, McLean and Kuberski 2013). Past work has shown that domestic audiences affect sanction duration and effectiveness (Dorussen and Mo 2001; Krustev and Morgan 2011; Martin 1993; Whang 2010*b*) and that audience costs arise when leaders back

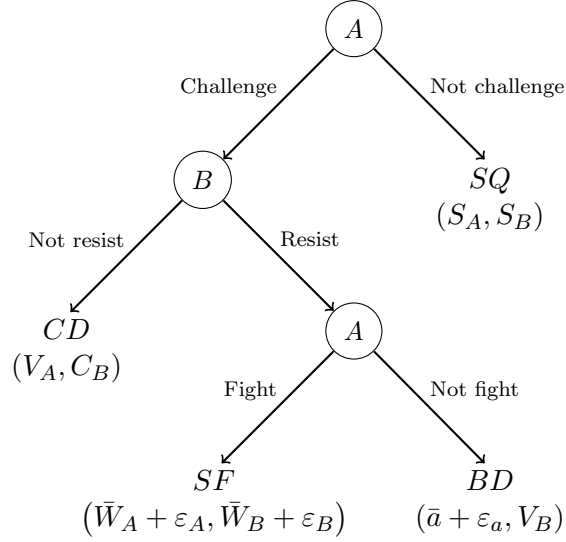
³This assumption can also be relaxed a bit by breaking dyads apart by time periods such as the dyad-decade, as we do below, or by using case-specific knowledge to determine changes.

down from sanction threats (Hart 2000; Thomson 2016). Yet scholars have not connected audience costs to the initiator’s decision to threaten sanctions.⁴ We fill this gap in the literature by fitting the crisis-signaling model to the Threat and Imposition of Sanctions (TIES) dataset. Our results indicate a novel *U*-shaped relationship where only leaders with large or small audience costs more freely threaten sanctions, as the former can credibly commit to such threats and the latter need not worry about the consequences of backing down. Such a result would be lost in traditional regressions that assume a monotonic relationship between audience cost measures and outcomes. Furthermore, the vast majority of observations are located on one side of the *U*-shaped curve: larger audience costs encourage leaders to threaten sanctions.

An important predecessor to this paper is Jo (2011) who demonstrates that multiple equilibria exist in the crisis-signaling game and that tML procedures ignoring multiplicity do not perform adequately. Indeed, this is the major problem we address in this paper, but we also build upon Jo’s endeavor in several ways. First, we explicate the computational issues that arise when researchers attempt to address multiplicity by either using refinements or verifying uniqueness post-estimation, including how multiple equilibria create discontinuous likelihood functions. Second, we provide three simple solutions to estimating the crisis-signaling games and benchmark their performances in a variety of experimental settings. Third, we apply the estimators to real-world data. While our solutions generally agree in sign and significance, they diverge considerably from tML routines. We also show that the tML is highly sensitive to implementation details. Our solutions remedy these issues.

⁴Exceptions to this include Peterson’s (2013) work on reputation costs and U.S. sanction threats and a brief aside in Whang, McLean and Kuberski (2013). Similarly, features of domestic audiences help to explain variation in the initiation of World Trade Organization disputes (Chaudoin 2014).

Figure 1: The canonical crisis-signaling game.



2 Model

States A and B compete over a good or a policy that is currently owned or controlled by B . At the beginning of the game, the states observe private information. State A observes $(\varepsilon_A, \varepsilon_a)$, where ε_A and ε_a are additively separable payoff shocks to A 's utility for war and backing down, respectively. Likewise, B observes ε_B which is an additively separable payoff shock to its war utility. All private information $(\varepsilon_A, \varepsilon_a$ and $\varepsilon_B)$ is independently drawn from a standard normal distribution.

Interaction proceeds according to Figure 1. First, A decides whether or not to challenge B for control over the good or policy, and if A does not challenge, then the game ends at node SQ with payoffs S_i for each state i . Second, after a challenge, B decides whether or not to resist A . If B does not resist, i.e., B concedes to A 's demands, then the game ends at node CD , and payoffs are V_A and C_B for states A and B , respectively. Finally, if B does resist, then A must decide whether to fight or not. When A fights or stands firm, the states receive $\bar{W}_i + \varepsilon_i$ at node SF . Similarly, when A backs down and does not fight, the game ends at node BD with A receiving

$\bar{a} + \varepsilon_a$ and B receiving V_B .

Perfect Bayesian equilibria (equilibria, hereafter) for the game can be represented as choice probabilities. Let p_C and p_F denote the probability that A challenges and fights (conditional on challenging) B , respectively, and let p_R denote the probability that B resists. Let $p = (p_C, p_R, p_F)$ denote a profile of choice probabilities. Furthermore, let θ denote the vector of payoffs, i.e., $\theta = (\bar{a}, C_B, (S_i, V_i, \bar{W}_i)_{i=A,B})$. The following result is due to Jo (2011) and characterizes the equilibria of the game in terms of a system of nonlinear equations.

Result 1 (Jo, 2011) *An equilibrium \tilde{p} exists, and \tilde{p} is an equilibrium if and only if it satisfies the following system of equations:*

$$\tilde{p}_C = 1 - \Phi\left(\frac{S_A - (1 - \tilde{p}_R)V_A}{\tilde{p}_R} - \bar{W}_A\right) \Phi\left(\frac{S_A - (1 - \tilde{p}_R)V_A}{\tilde{p}_R} - \bar{a}\right) \equiv g(\tilde{p}_R; \theta), \quad (1)$$

$$\tilde{p}_F = \Phi_2\left(\frac{\bar{W}_A - \bar{a}}{\sqrt{2}}, \bar{W}_A - \frac{S_A - (1 - \tilde{p}_R)V_A}{\tilde{p}_R}, \frac{1}{\sqrt{2}}\right) (g(\tilde{p}_R; \theta))^{-1} \equiv h(\tilde{p}_R; \theta), \quad (2)$$

and

$$\tilde{p}_R = \Phi\left(\frac{h(\tilde{p}_R; \theta)\bar{W}_B + (1 - h(\tilde{p}_R; \theta))V_B - C_B}{h(\tilde{p}_R; \theta)}\right) \equiv f \circ h(\tilde{p}_R; \theta), \quad (3)$$

where Φ is the CDF of the standard normal distribution and $\Phi_2(\cdot, \cdot, \rho)$ is the CDF of the standard bivariate normal distribution with correlation ρ .

In words, for a fixed θ , an equilibrium is completely pinned down by B 's probability of resisting. In addition, the functions f , g , and h are best-response functions. Specifically, the functions g and h compute A 's best response to B 's probability of resisting p_R , and function f denotes B 's best response to A 's probability of fighting. Furthermore, Jo (2011) illustrates that multiple equilibria exist in a nontrivial set of parameters, i.e., there exists several solutions to the equation $f \circ h(p_R; \theta) = p_R$.

Before proceeding two remarks about the model are in order. First, a key assumption in these models concerns how the analyst specifies the variance-covariance

matrix of the payoff shocks (Schultz and Lewis 2005; Whang 2010a). Since Lewis and Schultz’s (2003) first effort, several scholars have modified and generalized their assumptions governing private information (Bas, Signorino and Whang 2014; Wand 2006; Whang 2010a). Our goal is not to argue or determine which specification is better but to illustrate the computational difficulties that can arise from multiple equilibria and the structural estimation of these games.⁵ As such, we adopt the comparably sparse but original specification, where the shocks are independently and identically distributed according to the standard normal distribution. For completeness, in Appendix A we consider Whang’s (2010) specification, demonstrate that multiple equilibria can exist, and discuss how to modify our proposed procedures.

Second, researchers could change the model’s informational structure to ensure that a unique equilibrium exists. The most obvious possibility would be to adopt the quantal-response setup (McKelvey and Palfrey 1998; Signorino 1999). Bas, Signorino and Whang (2014) call this the “independent agent” approach. Essentially, actors are represented by different and independent agents at each of their assigned decision nodes. Each actor’s agents have the same mean utilities over terminal nodes but the utilities are subjected to independent payoff shocks, representing perceptual errors. For example, we currently assume that A observes ε_A and ε_a at the beginning of the interaction. Instead, A ’s agent at the beginning of the game would not observe the payoff shocks of A ’s agent at its last decision node. In this case, A ’s first agent is uncertain about whether A ’s second agent would stand firm or back down although the first could compute the probability of each action *ex ante*.

The QRE framework entails nontrivial substantive differences from the current model. Although the current information structure has the possibility of multiple perfect Bayesian equilibria, it also generates learning and signaling incentives. When

⁵Schultz and Lewis (2005) provide one justification for the specification used here.

State A decides to challenge B , she first observes the payoff difference between fighting and backing down. Thus, when State B observes A issuing a challenge, B updates her beliefs about whether A will fight or back down. Specifically, the amount of information communicated can be measured as the difference between the prior and the posterior probability that A fights:

$$\left| \underbrace{p_F}_{\text{posterior}} - \underbrace{\Phi\left(\frac{\bar{W}_A - \bar{a}}{\sqrt{2}}\right)}_{\text{prior}} \right|.$$

With the QRE information structure, the possibility of learning and signaling is null, however. If A only observes ε_A and ε_a conditional on arriving at the decision to stand firm or back down, then the state has no additional information to signal to state B when it issues a challenge. Furthermore, if A 's preferences are independent from B 's, then B has no opportunity to learn about the preferences of its opponent.⁶

⁶Signorino (2003, p. 343) concludes that the assumption underlying the standard QRE approach “is not innocuous. It greatly simplifies deriving the choice probabilities, but at a price: players cannot update in the model.”

3 Estimation: Problems and solutions

We consider D independent dyads or games. Each dyad is parameterized by covariates x_d and common payoff parameters β which determine the model's payoffs:

$$\theta(x_d, \beta) = \begin{bmatrix} S_{dA} \\ S_{dB} \\ V_{dA} \\ C_{dB} \\ \bar{W}_{dA} \\ \bar{W}_{dB} \\ \bar{a}_d \\ V_{dB} \end{bmatrix} = \begin{bmatrix} x_{dS_A} \cdot \beta_{S_A} \\ 0 \\ x_{dV_A} \cdot \beta_{V_A} \\ x_{dC_B} \cdot \beta_{C_B} \\ x_{d\bar{W}_A} \cdot \beta_{\bar{W}_A} \\ x_{d\bar{W}_B} \cdot \beta_{\bar{W}_B} \\ x_{d\bar{a}} \cdot \beta_{\bar{a}} \\ x_{dV_B} \cdot \beta_{V_B} \end{bmatrix}. \quad (4)$$

Each $x_{d(\cdot)}$ vector above contains zero or more explanatory variables.⁷ Hereafter, we are interested in the β parameters that are common across all games rather than $\theta(x_d, \beta)$.

Let β^* denote the parameters in the data generating process. Along with β^* , the covariate vector x_d determines the equilibrium $p^*(x_d, \beta^*) = (p_{dC}^*, p_{dF}^*, p_{dR}^*)$ that generates $T \geq 1$ outcomes $\{y_{dt}\}_{t=1}^T$, where y_{dt} is a terminal node in $\{SQ, CD, SF, BD\}$. Thus, $p_d^*(x_d, \beta^*)$ is a solution to the system of equations in Result 1, parametrized by payoffs $\theta(x_d, \beta^*)$. Thus the data are hierarchical: a complete observation is a dyad d with a single vector of exogenous traits x_d and a sequence of outcomes $\{y_{dt}\}_{t=1}^T$. While our structure reflects a common setup in the econometric literature on estimating games (e.g., Su 2014), it implies that T means something slightly different here than it does in common panel or time-series-cross-sectional (TSCS) data. With

⁷As in Lewis and Schultz (2003), identification depends on there being at least one variable (including the constant) for each player that does not appear in all of that player's utilities.

TSCS data, a complete observation d is composed of a covariate vector paired to a single outcome and “ T ” is the number of complete observations associated with some unit. In contrast, we envision an equilibrium of the game generating $T \geq 1$ outcomes. Of course, analysts can always set $T = 1$ in applications if they are concerned about this setup.⁸

There are two important assumptions implicit in our empirical setup. First, we assume that two states play from the same equilibrium conditional on x_d , rather than allowing the equilibrium to vary over within-dyad observations. Substantively, this assumption reflects the international system which has several forces incentivizing states to focus on a single equilibrium over time, including persistent international norms/institutions (Keohane 1984), a focal point specific to these two states (Schelling 1960), or other factors that emerge from their long time spans and repeated interaction. Technically, this is a standard assumption that is required in the recent empirical literature on estimating games with incomplete information (Bajari, Benkard and Levin 2007; Ellickson and Misra 2011). An alternative approach might assume that states play from the same equilibrium *across dyads* rather than within dyads. Such an assumption is more restrictive than ours, i.e., it requires the U.S. to interact with China the same way it interacts with Canada, which we do not require. Additionally, such an approach introduces an additional problem: because dyads are parameterized by different covariates, the number of equilibria may differ between two dyads even for a fixed θ , making it impossible to compare equilibria across observations.⁹

⁸One such concern is this structure may overstate the amount of information in the data due to the hierarchical setup. While increasing T provides more information about the choice probabilities, which makes the estimates of β more precise, the asymptotic standard errors provided in Appendix D.1 account for this structure.

⁹Consider the parameter values in Table 1, under the multiple equilibria setting with $D = 3$, where $x_1 = -\frac{2}{3}$, $x_2 = \frac{2}{3}$, and $x_3 = \frac{5}{3}$. According to Figure 2, there exist multiple equilibria in the second dyad but a unique equilibrium exists in the first and third. The third dyad indicates that the second should play from the largest equilibrium, but the first dyad indicates that the second should play from the smallest.

Table 1: Parameters for Monte Carlo experiments.

Utility	Multiple equilibria	Unique equilibrium
S_{dA}	0	0
S_{dB}	0	0
V_{dA}	1	1
C_{dB}	0	0
\bar{W}_{dA}	-1.9	-1.8
V_{dB}	1	1
\bar{W}_{dB}	$-2.9 + 0.1x_d$	$-2.45 + 0.1x_d$
\bar{a}_d	-1.2	-1.2

Second, when dyads play $T > 1$ rounds of play, we assume that a given equilibrium of the static game is played. Such an assumption can be justified if states play the game for a finite number of periods and private information is drawn independently over time.¹⁰ This assumption is a matter of convenience because our goal is to address the technical challenges that arise when estimating games with multiple equilibria *even in the most straightforward environments*. An added benefit of this simplicity is that we can easily enumerate the set of equilibria, allowing us to illustrate how equilibrium selection creates discontinuous likelihoods and the computational inefficiency of the tML in these situations. Nonetheless, scholars might be concerned about serial dependence in the private information shocks and long-term optimizing states. While we acknowledge the importance of considering dynamics in international relations, such concerns would require a dynamic theoretical model in order to address how repeated interaction affects states which is beyond the paper’s scope.¹¹

¹⁰The assumption can also be justified when the two states interact for an infinite number of periods if private information is not persistent and states use stationary strategies.

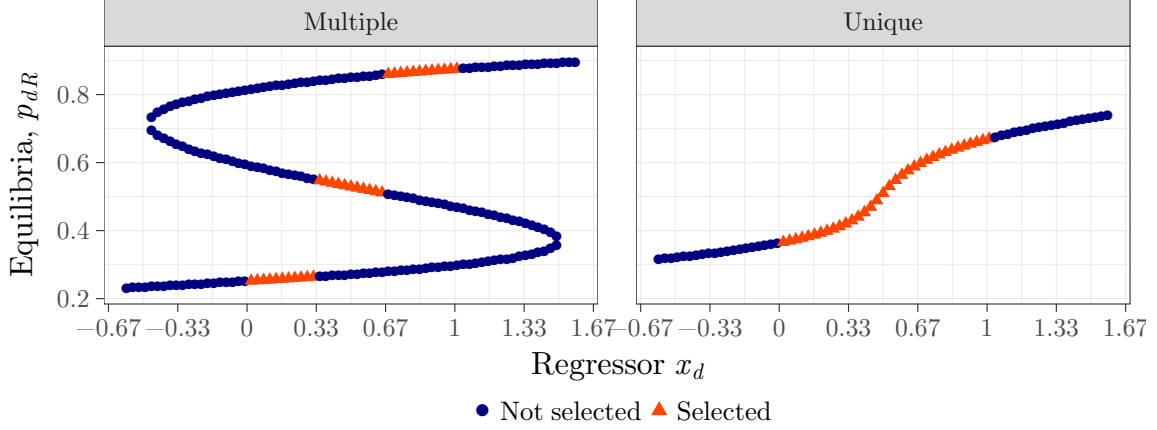
¹¹For an example of structurally estimating a dynamic game of crisis escalation see Crisman-Cox and Gibilisco (2018). A key property of their model is that states have no signaling incentives, i.e., private information is sufficiently transitory. With signaling incentives, a fully dynamic model becomes substantially more intractable.

Throughout we consider two numerical examples. Table 1 contains two sets of parameters that we use to demonstrate cases with a unique and with multiple equilibria. In both settings we include one regressor, $x_d \sim U[0, 1]$, which enters B 's war payoff.¹² There are a few things to note about the parameters in Table 1. First, we normalize the status-quo payoffs S_i and B 's concession payoff to zero, following standard identification assumptions (Jo 2011; Lewis and Schultz 2003). Second, the differences in the two columns are minor: by making small adjustments to only two parameters we can easily move into and out of situations where multiple equilibria exist. Third, these parameters reflect reasonable payoffs that satisfy the restrictions in Schultz and Lewis (2005). In addition, both war and backing down from threats are worse than the status quo, and actors only receive positive payoffs when their opponent backs down.

To illustrate the two settings, Figure 2 graphs the game's equilibrium correspondence with respect to x_d . In the left-hand panel of Figure 2, there are multiple equilibria for values of x_d between 0 and 1. Here, the orange triangles in the plots illustrate how we determine which equilibria generate the data in our Monte Carlo experiments. Specifically, when $x_d \in [0, \frac{1}{3})$, we use the smallest equilibrium probability of resisting p_R to generate the data for dyad d . When $x_d \in (\frac{2}{3}, 1]$, we use the largest. Finally, we use the moderate equilibrium in the remaining case. Notice that the equilibrium correspondence is smooth in the sense that it is upper hemicontinuous but selection creates discontinuities when modeling the probability of resistance p_{dR} as a function (not correspondence) of the covariate x_d . The right-hand side of Figure 2 graphs the equilibrium correspondence under parameters shown in the third column of Table 1, where there is a unique equilibrium for all values of x_d .

¹²A more realistic Monte Carlo experiment with multiple regressors can be found in Appendix C.4. Overall the results there confirm what we report here in the simpler setup.

Figure 2: The equilibrium correspondences for numerical examples



3.1 Problems with current practices

The current best practices in the literature closely follow the ML techniques discussed in Signorino (1999). For every β , an equilibrium to game d is computed by solving the system of equations in Result 1; call this solution $p(x_d, \beta)$. Note that this solution is not necessarily unique, and following standard practices, we do not search for all solutions.

Using $p(x_d, \beta)$, we define the probability of reaching each of the terminal nodes as

$$\Pr[y_{dt} \mid p(x_d, \beta)] = \begin{cases} (1 - p_{dC}) & \text{if } y_{dt} = SQ \\ p_{dC}(1 - p_{dR}) & \text{if } y_{dt} = CD \\ p_{dC}p_{dR}(1 - p_{dF}) & \text{if } y_{dt} = BD \\ p_{dC}p_{dR}p_{dF} & \text{if } y_{dt} = SF. \end{cases} \quad (5)$$

Under this setup, the log-likelihood takes the form

$$L(\beta \mid Y) = \sum_{d=1}^D \sum_{t=1}^T \log \Pr[y_{dt} \mid p(x_d, \beta)], \quad (6)$$

and the tML estimates attempt to maximize this log-likelihood.

As described in Jo (2011), the current approach evaluates the likelihood function as if a unique equilibrium exists. That is, for each guess of the parameters, we compute an equilibrium, $p(x_d, \beta)$, using a numeric equation solver. If there are multiple equilibria, then there is an indeterminacy in how analysts evaluate $p(x_d, \beta)$. If the equation solver of choice selects the wrong equilibrium, i.e., not the one in the data generating process, then the likelihood is computed incorrectly, resulting in mistaken inferences. To better see this problem, suppose there are D dyads, and fixing parameters β , suppose each dyad admits $n > 1$ equilibria. In this case, there are n^D possible values of the log-likelihood for just this one guess at the parameter vector. Standard equation solvers return just one of the n^D combinations. As D increases, it is increasingly implausible that the correct selection is made. An implication of this discussion is that two researchers can reach conflicting conclusions even when analyzing the same data if they implement the tML estimator with different equation solvers. We illustrate this problem in our empirical application.

Before proceeding, we first consider potential fixes to the commonly standard ML routine. To this end, we first ask: Can multiplicity in the crisis-signaling game be solved with traditional refinements? If so, tML techniques can be used so long as they are adjusted to always select the surviving equilibrium. Refinements based on off-the-path-of-play beliefs, such as the Intuitive Criterion or Divinity, are inconsequential here as all histories are reached with positive probability in every equilibrium. Because of this, an analyst may be tempted to use a refinement called regularity, which subsumes several other refinements such as perfection, essentialness, and strong stability (van Damme 1996).¹³ As we show in Appendix B, for almost all parameter values,

¹³For a formal definition, see Appendix B.

all equilibria of the crisis-signaling game satisfy regularity.¹⁴ Most importantly, the result demonstrates that multiplicity cannot be “refined away” using standard criteria, and the predictive indeterminacy that plagues traditional maximum likelihood methods still persists.¹⁵

With traditional refinements offering little headway, analysts may turn to *ad hoc* selection criteria such as selecting the equilibrium that maximizes a convex sum of A and B ’s payoffs. But determining the selection criterion forces an additional modeling choice onto the analyst. As we show in our empirical application, such a choice is consequential and can heavily influence the resulting estimates. Analysts could also consider empirical selection: for each dyad, select the equilibrium that maximizes the dyad’s contribution to the likelihood. This would also remove the indeterminacy in $p(x_d, \beta)$, but its implementation has several drawbacks. Researchers would need to reliably compute all equilibria for every dyad at every guess of the parameters, a computationally demanding task. In addition, imposing this (and other) selection criterion introduces discontinuities in the likelihood function as the number of equilibria and hence the solution to the criterion varies across different parameter values.¹⁶

To illustrate this, Figure 20 in Appendix F graphs the likelihood as a function of one parameter, while all remaining parameters are held at their true values. Note that the true value indeed maximizes the likelihood, but discontinuities arise as the number of equilibria change in different areas of the parameter space. Furthermore the number of discontinuities and equilibrium computations both increase in D . The

¹⁴We say that a property holds for almost all parameters θ , if it does not hold at most in a closed, Lebesgue-measure-zero subset of \mathbb{R}^8 .

¹⁵We also consider best-response stability. We prove formally that if multiple equilibria exist, then at least one is best-response unstable. Nonetheless, if multiple equilibria exist, then there are generally multiple best-response stable equilibria. For example, in the left-hand graph in Figure 2, the largest and smallest equilibria are best-response stable, while the middle is unstable.

¹⁶Technically, this problem arises because the equilibrium correspondence, and hence likelihood correspondence, is upper, but not lower, hemicontinuous.

problems are further exacerbated if the remaining parameters are not known and need to be estimated as well.

3.2 Pseudo-likelihoods

Our first proposal involves a two-step estimator based on Hotz and Miller (1993) that essentially removes the indeterminacy associated with multiple equilibria by using the observed data to select appropriate equilibrium beliefs. In the first step, we produce consistent (in T or D) estimates of the equilibrium choice probabilities p_{dR}^* and p_{dF}^* , for $d = 1, \dots, D$. We label these estimates $\hat{\mathbf{p}}_{\mathbf{R}} = (\hat{p}_{1R}, \dots, \hat{p}_{DR})$ and $\hat{\mathbf{p}}_{\mathbf{F}} = (\hat{p}_{1F}, \dots, \hat{p}_{DF})$. While in theory we are agnostic about how an analyst obtains the first-stage estimates, in practice we have found that random forests tend to work very well across a variety of sample sizes and settings.

Next, consider how actors best respond to these first-stage estimates. By Result 1, the best-responses take the form:

$$\hat{p}(\hat{p}_{dR}, \hat{p}_{dF}; x_d, \beta) = \begin{bmatrix} g(\hat{p}_{dR}; x_d, \beta) \\ h(\hat{p}_{dR}; x_d, \beta) \\ f(\hat{p}_{dF}; x_d, \beta) \end{bmatrix}. \quad (7)$$

In other words, if actors play the game *as if* they believed their opponents use strategies estimated in the first stage, \hat{p}_{dR} and \hat{p}_{dF} , then \hat{p} are their best responses. These best-responses approach their true values as the first-stage estimates become more accurate. Using the first-stage estimates and the associated best-responses, we build the pseudo-log-likelihood function as

$$PL(\beta \mid \hat{\mathbf{p}}_{\mathbf{R}}, \hat{\mathbf{p}}_{\mathbf{F}}, Y, X) = \sum_{d=1}^D \sum_{t=1}^T \log \Pr[y_{dt} \mid \hat{p}(\hat{p}_{dR}, \hat{p}_{dF}; x_d, \beta)]. \quad (8)$$

What is the intuition behind the estimator? If we know the equilibrium choice probabilities, i.e., $\hat{p}_{dR} = p_{dR}^*$ and $\hat{p}_{dF} = p_{dF}^*$ for all dyads d , then the pseudo-likelihood is the likelihood in Equation 6 with the correct equilibrium selection. In addition, it is a continuous function of the parameters β . The equilibrium choice probabilities are unobserved variables, however. Thus, we estimate them from the data, which is possible given our assumptions that two states play from the same equilibrium conditional on x_d . For example, because the states in dyad d are playing from one equilibrium, when T is large, we can estimate p_{dR}^* and p_{dF}^* using frequency estimators:

$$\hat{p}_{dR}^* = \frac{\sum_{t=1}^T \mathbb{I}[y_{dt} \in \{SF, BD\}]}{\sum_{t=1}^T \mathbb{I}[y_{dt} \in \{SF, BD, CD\}]} \quad \text{and} \quad \hat{p}_{dF}^* = \frac{\sum_{t=1}^T \mathbb{I}[y_{dt} = SF]}{\sum_{t=1}^T \mathbb{I}[y_{dt} \in \{SF, BD\}]},$$

where \mathbb{I} is the indicator function. As we observe more draws from the same equilibrium, i.e., T goes to infinity, the frequency estimates converge to their true values because the equilibrium p_d^* puts positive probability on all histories. Substituting the frequency estimates into Equation 8 demonstrates that the pseudo-likelihood converges to the true likelihood as T increases, and under standard regularity conditions the PL estimates converge to the true ML estimates. Thus, by estimating equilibrium beliefs from the data in a first-stage, we can select the appropriate equilibrium in a continuous manner when estimating payoff parameters during the second stage.

In finite samples, frequency estimators may be impractical. One alternative is to pool information across dyads and estimate the choice probabilities as functions of covariates x_d , albeit with highly flexible methods—hence our assumption that two observationally equivalent dyads play from the same equilibrium. As mentioned, we have found that random forests work particularly well in both our simulations and applications. Nonetheless, the PL estimator may perform poorly if the first stage is misspecified or imprecise. The two methods we discuss below attempt to overcome

this issue.

3.2.1 Nested pseudo-likelihood

The NPL approach, proposed by Aguirregabiria and Mira (2007), builds on the PL by using best responses to update the first-stage choice probabilities upon knowing the PL estimates. This process is repeated until convergence. More precisely, the NPL algorithm begins with the PL estimates,

$$(\hat{\beta}_0^{NPL}, \hat{\mathbf{P}}_{\mathbf{R},0}, \hat{\mathbf{P}}_{\mathbf{F},0}) = (\hat{\beta}^{PL}, \hat{\mathbf{P}}_{\mathbf{R}}, \hat{\mathbf{P}}_{\mathbf{F}}),$$

and for the k th iteration, set

$$\begin{aligned}\hat{p}_{dF,k} &= h(\hat{p}_{dR,k-1}; x_d, \beta_{k-1}) \\ \hat{p}_{dR,k} &= f(\hat{p}_{dF,k-1}; x_d, \beta_{k-1}) \\ \hat{\beta}_k^{NPL} &= \operatorname{argmax}_{\beta} PL(\beta \mid \hat{\mathbf{P}}_{\mathbf{R},k}, \hat{\mathbf{P}}_{\mathbf{F},k}, Y, X).\end{aligned}$$

The algorithm is repeated until the parameters and choice probabilities cease changing. The intuition is to decrease the analyst's reliance on correct first-stage estimates by updating the choice probabilities with the new information captured in the estimated payoff parameters.

Without a particular stability condition on the data generating process, the NPL algorithm may fail to converge (Pesendorfer and Schmidt-Dengler 2010). Specifically, if the data generating equilibrium is best-response stable, the above iteration will converge to the correct equilibrium as long as the starting value is not too far away. In contrast, if the data generating equilibrium is unstable, the above iteration may not converge to the true equilibrium. In Appendix C.3, we consider how sensitive the

PL and NPL are to unstable equilibria. Overall, we find that both the PL and the NPL still outperform the tML even if best-response unstable equilibria dominate in the data.

3.3 Constrained MLE

An alternative approach is to use a full-information CMLE, as proposed by (Su and Judd 2012). Applied to this problem, we maximize the likelihood in Equation 6 subject to the equilibrium constraints in Result 1. Define

$$\bar{p}(p_{dR}; x_d, \beta) = \begin{bmatrix} g(p_{dR}; x_d, \beta) \\ h(p_{dR}; x_d, \beta) \\ p_{dR} \end{bmatrix}, \quad (9)$$

then the CMLE solves the following problem:

$$\begin{aligned} \max_{\beta, \mathbf{PR}} \quad & \sum_{d=1}^D \sum_{t=1}^T \log \Pr[y_{dt} \mid \bar{p}(p_{dR}; x_d, \beta)], \\ \text{s.t.} \quad & f \circ h(p_{dR}; x_d, \beta) = p_{dR}, \quad d = 1, \dots, D. \end{aligned} \quad (10)$$

Su and Judd (2012) demonstrate that the CMLE is equivalent to the *true* MLE procedure in which equilibria are selected to maximize each dyad’s contribution to the likelihood. Thus, the estimator is essentially using the data to select equilibria, which is similar to the PL procedure where data were used to estimate equilibrium beliefs. As mentioned above, modifying the tML to compute every equilibrium at every guess of the parameters and to select the ones that maximize the likelihood dramatically reduces its feasibility because it requires repeated equilibrium computations and introduces discontinuities.

The CMLE avoids these problems. By not requiring that \mathbf{p}_R satisfy the equilibrium condition at every step in the constrained optimization, the CMLE avoids any equilibrium computation while ensuring that the objective function is well-behaved and continuous. As such, the true maximum likelihood estimates are discovered with a much lower computational burden than the corrected-tML discussed above. Additionally, the CMLE improves on the pseudo-likelihood procedures by eliminating the need to rely on first-stage estimates, resulting in both bias and efficiency gains.

Despite these improvements, the CMLE has two drawbacks. First, the full-information constrained optimization approach introduces D auxiliary parameters in the form of $\hat{\mathbf{p}}_R$; as such we need $T > 1$ in order use this estimator. In contrast, the pseudo-likelihood approaches cover the $T = 1$ case. However, our Monte Carlo experiments demonstrate that the CMLE performs well even with a small number of within-game observations. Second, solving this constrained optimization problem requires specialized —but still open source— software; Appendix D contains complete implementation details.

4 Performance

We now evaluate the performance of the estimators in two settings: when there are multiple equilibria in the data generating process and when there is a unique equilibrium. We continue to use the parameter values from Table 1, where x_d is distributed standard uniform.¹⁷ Throughout, we use the ordinary implementation of the tML as our baseline for comparison, which uses arbitrary equilibrium selection

¹⁷The results we present here are unchanged when we a more realistic Monte Carlo experiment with multiple covariates in Appendix C.4.

and Nelder-Mead’s simplex method to find the estimates. These implementation choices match current practices as found in replication archives.¹⁸

To estimate equilibrium choice probabilities in the PL and NPL methods we use random forests. There are two models in the first-stage, where the dependent variables are the nonparametric frequency estimates of the probability that B resists (for $\hat{\mathbf{p}}_{\mathbf{R}}$) and A fights (for $\hat{\mathbf{p}}_{\mathbf{F}}$). We fit the former only with observations in which A challenges, and we fit the latter only with observations in which B resists. For predictors, we include the one regressor x_d .

We vary the number of dyads, D , between 25 and 200 and the number of within-game observation, T , between 5 and 200 to create simulated datasets of various sizes. For each combination of D and T , we draw x_d from the standard uniform distribution and then select the appropriate equilibrium that generates the data for the corresponding dyad as shown in Figure 2. Finally, we use the simulated data to estimate the game using all four estimators. Starting values for the PL and tML are drawn from a standard uniform distribution, and the same values are used within each Monte Carlo iteration. The CMLE and NPL use the PL estimates as starting values.¹⁹ We repeat this process 1,000 times for each pair of D and T and for each of the parameter settings in Table 1.²⁰

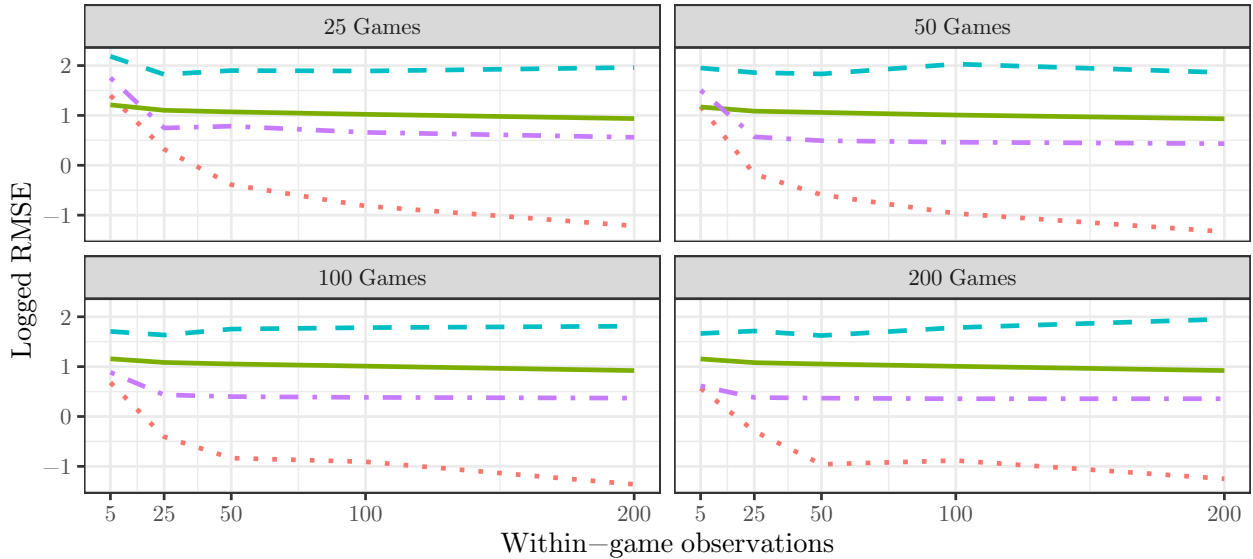
The main results of the experiment are summarized in Figures 3 and 4, which

¹⁸As the tML’s objective function contains discontinuities, gradient-free methods, such as Nelder-Mead, are a common choice for avoiding expensive global optimization. We also considered global and quasi-Newton methods our conclusions are unchanged. In contrast to the tML, our proposals have continuous log-likelihood functions, and so we use the gradient based Newton-Raphson method for the PL and NPL and a Newton-based interior point method for the CMLE.

¹⁹The choice of random starting values for the tML and PL reflect the fact that they are competing methods in this experiment. In contrast, the CMLE and NPL are natural extensions of the PL approach and use the PL to inform them. We return to this in Appendix E where we check the tML’s performance when it is also started from the PL estimates.

²⁰Within each Monte Carlo iteration, results are considered converged and recorded only if a successful convergence code is returned by the optimizer in question *and* all the point estimates are between -50 and 50 .

Figure 3: RMSE in signaling estimators with multiple equilibria.



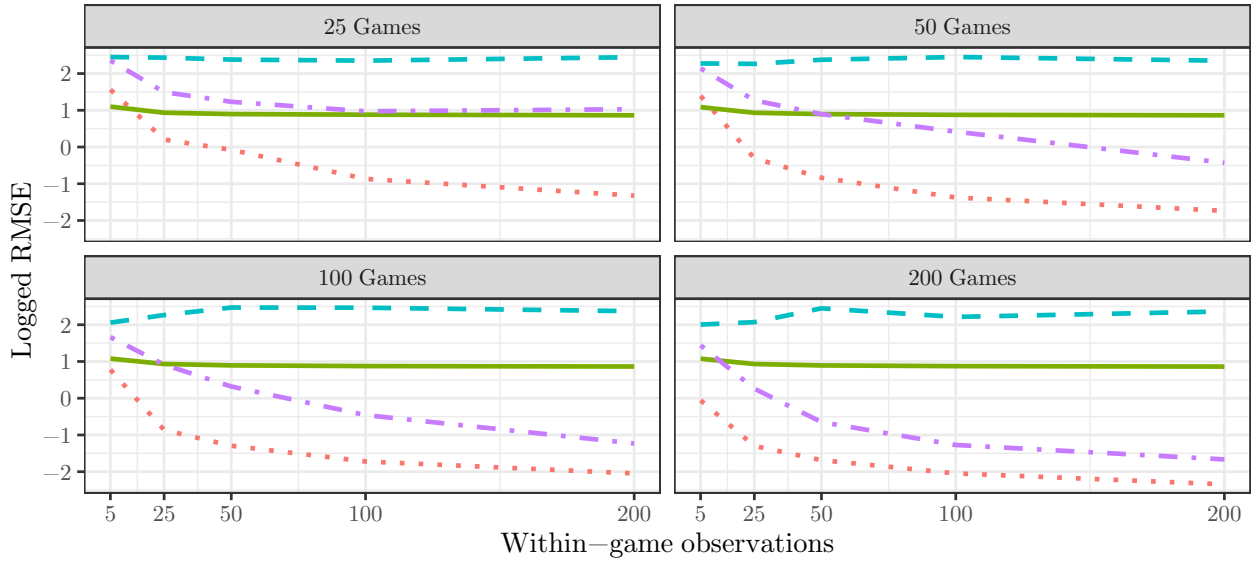
Estimator ⋯ CMLE — PL - - - tML · - · NPL

compare the logged root-mean-square error (RMSE) of the estimators. The first thing we note is that the tML (blue, dashed line) performs consistently bad and shows no improvement as the amount of data increases in either D or T . In many cases, its performance worsens as T increases.²¹

Contrast these results to those from the other estimators, which generally all improve with more data. The PL (green, solid line) tends to be best performing estimator when both T and D are small. Additional analysis in Appendix C shows that the estimator tends to have more bias than the others and that its strong performance is driven by low variance. The NPL (pink, dot-dashed line) greatly improves the bias associated with the PL method without adding too much variance, and as a result, we see that it performs very well in most settings, particularly as the amount of data increases. Overall, the CMLE (red, dotted line) tends to be the best. How-

²¹Appendix C contains additional Monte Carlo results relating to the bias, variance, convergence rates, and computational time.

Figure 4: RMSE in signaling estimators with a unique equilibrium



Estimator CMLE ——— PL - - - - tML - · - · NPL

ever, this great performance often comes at the cost of decreased convergence rates and non-standard software choices (see Appendix D for more information).

Comparing Figures 3 and 4 reveals that the tML has uniformly poor performance regardless of the number of equilibria that exist in the signaling game generating the data.²² What explains the poor performance of the tML in the unique equilibrium experiment? Even in this setting the tML’s likelihood function is often evaluated at incorrect parameter values. For example, we pick starting values that are drawn uniformly over the interval $[0, 1]$. These are obviously incorrect, and the optimizer will need to search over the parameter space, evaluating the likelihood function at incorrect parameter values. In some instances, dyads parameterized (incorrectly) by these values will have multiple equilibria, and the objective function will need to select an equilibrium in an *ad hoc* manner. This selection will lead to discontinuities and

²²Figures 12 and 13 in Appendix C.2 compare the estimators’ bias and variance in the unique equilibrium setting and illustrate that the tML has the worst performance on both measures.

creates the possibility that an incorrect equilibrium is selected, i.e., an equilibrium that has little relation to the one generating the data. These issues can lead even more robust optimizers astray.

We also find that the tML appears to face numerical challenges during the optimization process. Even in cases where we verify that the tML only considers candidate parameter vectors that are associated with a unique equilibrium, we find that the optimizer frequently converges to a wrong answer. These issues do not go away (and often get worse) when we consider alternative optimization routines to the simplex method, such as global or quasi-Newton methods. Additionally, with our empirical example we find that very small implementation differences, including simply changing software versions, result in wildly different tML estimates. Overall, this level of sensitivity indicates that the equilibrium computation in the tML’s likelihood creates a highly nonlinear optimization problem that is difficult to solve. We do not observe these kinds of stability issues with any of the other methods.

With the above theoretic and numeric concerns in mind, it is worth considering how sensitive the tML’s performance is to starting values; we investigate this in Appendix E. After all, using informative starting values may smooth over some of these issues and provide a path forward for the tML. Indeed, we find that the tML’s performance improves if (a) there is a unique equilibrium at the true parameters and (b) the tML has starting values that are either the true parameters or the PL estimates. However, even when initialized with the PL estimates, the tML rarely improves much on, and sometimes worsens, the PL’s performance, and it is almost always worse than the NPL or CMLE. Overall, relying on informed starting values and equilibrium uniqueness in the data generating process is perilous for applied researchers because neither can be verified before estimation. Furthermore, the PL, NPL, and CMLE perform at least as well as tML, and often much better, across all

the experiment conditions we consider.

Before turning our attention to economic sanctions, we report the following conclusions.

1. The tML routine performs the worst in both multiple and unique settings, regardless of the number of observations.
2. The NPL and PL methods consistently perform well, but the PL outperforms the NPL when the number of within-game observations is small, and *vice versa* when the number of within-game observations is large. In every experiment, the NPL is less biased than the PL.
3. The CMLE is almost always the best, but it is the most difficult to implement as the highly nonlinear nature of the constraint requires advanced software.

5 Application to economic sanctions

Our application is motivated by Whang, McLean and Kuberski (2013, WMK, hereafter) who also use the empirical crisis-signaling game to study the implementation of economic sanctions. They test the hypotheses that greater economic dependence decreases the probability that state B resists and increases the amount of belief updating, finding substantial support for the former but not the latter. The game is reproduced in Figure 5. The outcomes are status quo, concede to the threat, impose sanctions, and back down, which are denoted SQ , CD , SF , and BD , respectively.

An observation in WMK is a politically relevant directed dyad-decade. In their study, a directed dyad is politically relevant if there exists at least one sanction threat issued from State A to State B in the TIES dataset during the 1971–2000 period. Within each directed dyad, WMK aggregate the dependent variable to be the most extreme outcome within a directed dyad-decade, dividing the time frame into three groups 1971–80, 1981–90, and 1991–2000.

Like WMK we aggregate covariates x_d to the decade level, but unlike WMK, we dis-aggregate the outcomes y_{dt} to the monthly level ($T = 120$). We treat the observed y_{dt} within each directed dyad-decade as if they are repeated draws from the same equilibrium.²³ Effectively this means that each game d consists of a decade-level covariate vector x_d that is thought to produce each directed dyad’s monthly interaction over the course of the decade. In terms of our setup, each game is a politically relevant, directed dyad-decade, and we observe $T = 120$ observations from each game.²⁴

For our purposes, this approach has two important advantages. First, the CMLE procedure requires within-game multiple observations for identification. Without this setup, we could not illustrate this estimator even though it performed quite well in the Monte Carlo experiments. Second, we do not ignore variation within each decade: a directed dyad with only one threat issued in a decade may be substantially different than one with several threats in the same period. Thus, our application does not replicate previous analyses but rather highlights the differences between tML routines and those that we propose.

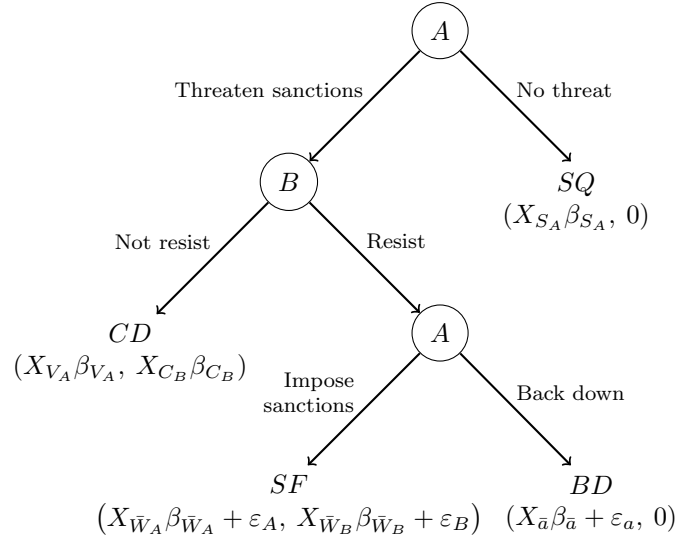
Following WMK, we use the *Final Outcome* variable to record the dependent variable, which denotes how sanction-threat episodes end.²⁵ When there is no episode in a month, we record the status quo. When *Final Outcome* records either “acquiescence” by the target or a negotiated settlement, we record the outcome as B giving

²³WMK specify payoff shocks following Whang (2010), discussed in Appendix A, where they also estimate covariance parameters. WMK not only report that these covariance estimates are below 0.07 in magnitude, but also fail to reject the null hypothesis that the covariances are equal to zero at the $p < 0.08$ level.

²⁴This is stricter than WMK’s threshold for political relevance, but using their less restrictive inclusion criteria does not affect our substantive conclusions on audience costs as we show in Appendix I.

²⁵Note all the action is coded as occurring in the month when the episode starts. If a play of the game actually unfolds over a back-and-forth, we might be overstating the number of status quo observations. To address this, we also consider a robustness check in our supplementary information where we redo our analysis at the quarterly level, i.e., $T = 40$.

Figure 5: Economic sanctions as a signaling interaction.



into A 's threat (node CD). Likewise, whenever the *Final Outcome* variable notes that actual sanctions are imposed, we list A as standing firm on its threat (node SF). Finally, when *Final Outcome* denotes that A either “capitulates” or the situation is unresolved, we list A as backing down (node BD). After dropping irrelevant dyad-decades, i.e., those with no recorded threats or sanctions, we are left with 418 games, each with 120 within game observations that span one of the three time frames, 1971–80, 1981–90, and 1991–2000.²⁶

The independent variables, their sources, and how they enter the actors' payoffs are listed in Table 2, following the specifications in WMK. All variables are measured on the dyad-decade level as discussed above.²⁷

²⁶Some countries enter/exit the data in the 1990s so there are 15 dyads where T is between 72 and 96.

²⁷While there are legitimate concerns associated with aggregating any set of variables to the decade level, we use it in our main analysis to follow WMK. Additionally, we show in Appendix H that there is actually very little variation in covariates within each dyad-decade. Furthermore, in Appendix I we also check the analysis with five years $T = 60$ and one year $T = 12$. When $T = 12$ our covariates are measured on the standard dyad-year level. Finally, we also consider a situation where both x_d and y_d are measured at the dyad-year level ($T = 1$). Our coefficients on audience costs remain stable in sign, significance, and magnitude.

Table 2: Variables in the economic sanctions model

Variable	Utilities	Description	Source
Fixed to 0	S_B, V_B	Identification restriction	–
Constant	$V_A, C_B, \bar{W}_i, \bar{a}$	Omitted from S_A for identification	–
Econ. Dep _A	S_A, \bar{W}_A	A’s economic dependence on B	COW
Dem _A	S_A, \bar{a}, \bar{W}_A	A’s Polity2 score	Polity IV
Contiguity	S_A, C_B	Contiguity between the states (1-6)	COW
Alliance _B	V_A, C_B	Alliance between the states (0/1)	COW
Costs _A	V_A	Anticipated costs to A	TIES
Econ. Dep _B	C_B	B’s economic dependence on A	COW
Costs _B	C_B	Anticipated costs to B	TIES
Cap. Ratio	\bar{W}_i	(log) ratio of A’s capabilities to B	COW
Dem _B	\bar{W}_B	B’s Polity2 score	Polity IV

5.1 Point estimates

Table 3 displays our main results. Each column contains parameter estimates and standard errors using the different estimators.²⁸ There are several notable patterns. First, the techniques derived from the dynamic games literature produce estimates that agree in direction, magnitude, and significance. Models 2–4 match signs for 14 out of 21 coefficients, and when we reject a null hypothesis using one estimator, we generally do the same for one of the others. Second, the tML returns estimates that diverge wildly from the other three. The problem appears particularly bad for coefficients that enter the target state’s concession payoffs, C_B .

Third, not only does the tML routine return different point estimates, it also produces substantive implications that diverge from the other three estimations. In Model 1, for example, economic dependence is negative in B’s concession payoff, C_B . This suggests that greater dependence decreases the target state’s preference

²⁸We also attempt to use the corrected-tML approach using a global optimizer and selecting the equilibrium that maximizes the likelihood of each dyad. After 21 days, convergence was not achieved and the log-likelihood value was still worse than the PL, NPL, and CML estimators. As a result, we conclude that the corrected-tML, while theoretically sound, is too computationally intense for applied research.

Table 3: Economic sanctions application

	tML Model 1	Pseudo-Likelihood Model 2	Nested Pseudo Likelihood Model 3	CMLE Model 4
S_A : Econ. Dep $_A$	0.05 (0.29)	-0.32 (0.67)	-0.22 (1.11)	-0.52 (0.55)
S_A : Dem $_A$	-0.00 (0.00)	0.00 (0.07)	0.03 (0.07)	0.01 (0.03)
S_A : Contiguity	0.27* (0.10)	0.01 (0.01)	0.05 (0.04)	0.04* (0.01)
S_A : Alliance	-0.06 (0.08)	-0.18* (0.03)	-0.17 (0.10)	-0.16* (0.07)
V_A : Const.	-0.06 (0.08)	-0.33 (0.83)	1.60 (1.65)	1.31* (0.12)
V_A : Costs $_A$	-0.04 (0.03)	0.36 (0.24)	-0.05 (0.27)	-0.19 (0.17)
C_B : Const.	0.81 (0.91)	-1.12* (0.31)	-2.14* (0.80)	-4.53* (2.24)
C_B : Econ. Dep $_B$	-0.21 (0.16)	2.32* (0.68)	2.34* (1.11)	2.83* (0.63)
C_B : Costs $_B$	-0.08* (0.03)	0.08 (0.06)	0.12 (0.06)	0.19* (0.05)
C_B : Contiguity	-0.25* (0.02)	0.13* (0.04)	0.12* (0.04)	0.10* (0.03)
C_B : Alliance	0.10 (0.09)	-0.05 (0.13)	-0.03 (0.12)	-0.02 (0.11)
\bar{W}_A : Const.	-0.15 (0.78)	-2.43* (0.10)	-2.42* (0.13)	-2.46* (0.08)
\bar{W}_A : Econ. Dep $_A$	0.07 (0.75)	0.39 (0.90)	0.01 (1.14)	-0.05 (0.18)
\bar{W}_A : Dem $_A$	0.01 (0.01)	0.01 (0.08)	0.04 (0.07)	-0.00 (0.03)
\bar{W}_A : Cap. Ratio	-0.01 (0.01)	0.02 (0.01)	0.03 (0.03)	0.04* (0.01)
\bar{W}_B : Const.	-0.38 (1.13)	0.48* (0.22)	-0.91 (1.11)	-4.42 (2.92)
\bar{W}_B : Dem $_B$	0.01* (0.00)	-0.00 (0.01)	-0.00 (0.01)	-0.01* (0.01)
\bar{W}_B : Cap. Ratio	0.01 (0.01)	0.11* (0.04)	0.12* (0.05)	0.29* (0.09)
\bar{a} : Const.	-0.56 (0.77)	-2.63* (0.09)	-2.64* (0.13)	-2.71* (0.10)
\bar{a} : Dem $_A$	-0.00 (0.01)	-0.00 (0.07)	0.02 (0.07)	-0.00 (0.03)
Log L	-4102.76	-3964.03	-3932.45	-3927.91
$D \times T$	418 \times 120	418 \times 120	418 \times 120	418 \times 120

Notes: * $p < 0.05$

Asymptotic standard errors in parenthesis, see Appendix D.1 for details.

for peaceful resolutions over escalation. Yet, the remaining estimators suggest the opposite. Here, more dependence implies B should be more likely to prefer peace to escalation, a finding which closely matches previous work (Hafner-Burton and Montgomery 2008; Lektzian and Souva 2003; McLean and Whang 2010).²⁹ In a similar vein, Schultz and Lewis (2005, p. 123–4) discuss “sensible” payoff restrictions that imply bargaining concessions are valuable. They expect state A to prefer immediate concessions from B to the status quo, i.e., $V_A > S_A$. Similarly, state B should prefer A backing down to immediately conceding, i.e., $V_B > C_B$. We determine whether these inequalities hold given the parameter estimates in Table 3. Using Models 3–4, all observations satisfy these restrictions. In Model 2, only 20% satisfy the restriction, and no observation satisfies these restrictions with the estimates from Model 1.

For another example, consider audience costs, i.e., the initiating state’s payoff from backing down, \bar{a} . Notice that the relevant constant term is negative, significant, and large in magnitude in all three models that accommodate multiple equilibria. This suggests that states or leaders are indeed punished for backing down after issuing threats.³⁰ In fact, in Models 2–4, we reject the null hypothesis that $\bar{a} \geq 0$ at the $p < 0.05$ level in every observation. In contrast, we cannot reject the null hypothesis that $\bar{a} \geq 0$ at the $p < 0.1$ level in the tML model for any observation. Our analysis suggests that researchers may underestimate audience or belligerence costs if estimation techniques do not accommodate the multiplicity of equilibria.

Overall, our results demonstrate that tML routines can produce point estimates and substantive implications that diverge from our proposed methods. To better

²⁹Although our analysis is not a replication exercise, we confirm WMK’s central finding that greater economic dependence for State B increases its preferences for the concession outcome, C_{dB} , and decreases its probability of resisting, p_{dR}^* .

³⁰Note that the estimate captures both belligerence and audience costs (as in Kertzer and Brutger 2016). We dig deeper into this in the subsequent section by conducting counterfactuals that isolate the substantive effects of audience costs while fixing belligerence costs.

illustrate that the differences are due to equilibrium selection and the computational problems addressed above, we conduct two additional analyses in Appendix F. First, we fit the sanctions model using a tML routine that is identical to what we use in Model 1 except for how it computes equilibria. Most surprisingly, the two tML results diverge in both sign (for 9/20 estimates) and significance (for 13/20 estimates). Second, we show the importance of starting values. Perhaps unsurprisingly, when we use the PL estimates as starting values, the tML improves, but is still worse than the NPL and CMLE in terms of log-likelihood. Thus, two researchers can reach substantively diverging conclusions with different software choices even when analyzing identical data sets.

5.2 Audience costs and substantive effects

How do audience costs affect the likelihood of leaders threatening sanctions? In the previous section, we demonstrated that tML routines can produce point estimates that diverge wildly from our solutions. In this section, we analyze the substantive effects of audience costs on the equilibrium probability of threatening sanctions, p_C , illustrating that the tML routines can fail to uncover important comparative statics. We focus on audience costs because of their importance to the economic sanctions literature (Dorussen and Mo 2001; Drezner 1998, 2003; Martin 1993; Whang, McLean and Kuberski 2013). In addition, previous work has not connected audiences to the likelihood that leaders threaten sanctions.

Regarding audience costs and the incentives to initiate threats, there exist *a priori* expectations from the conflict literature.³¹ On the one hand, larger audience costs may discourage a leader from initiating disputes if the leader never expects subse-

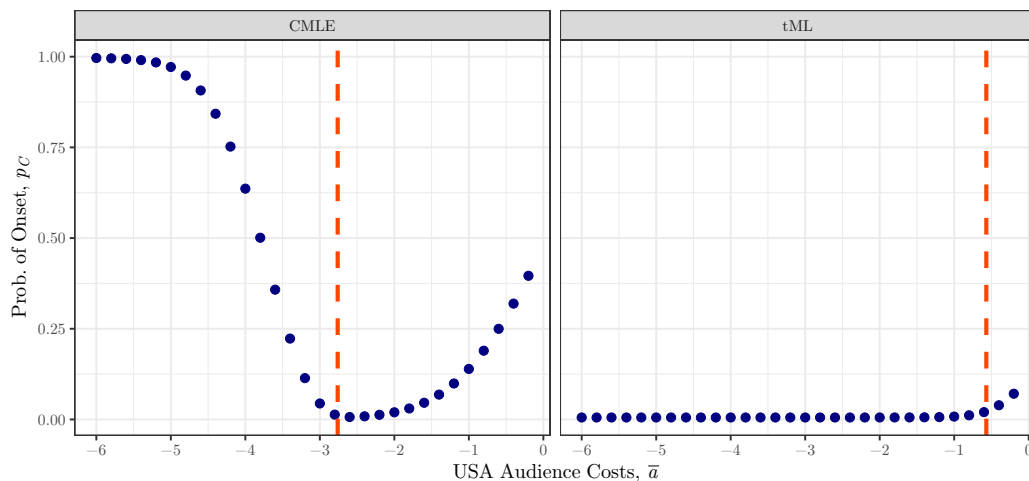
³¹While the decisions around conflict and sanction initiation are obviously different, the conflict literature provides a good jumping-off point for theorizing about the effects of audience costs on threatening sanctions.

quent concessions; Weeks (2012) and Prins (2003) find such a relationship using data from militarized interstate disputes. On the other hand, larger audience costs create bargaining advantages within disputes (Dorussen and Mo 2001; Schultz 1999), and Martin (1993) shows that leaders with larger costs expect additional cooperation from the international system when imposing sanctions. As such larger costs may encourage leaders to initiate disputes, and Crisman-Cox and Gibilisco (2018) find that this effect dominates in interstate crises between 1993–2007. Finally, even though survey experiments demonstrate that audience costs are highly salient for economic crises (Thomson 2016), they may be unrelated to the likelihood of sanction threats given our estimates in the previous section.

To see which, if either, effect appears here, we consider the directed dyad in which the U.S. is the initiating state A and China is the target state B between 1991–2000, the most recent decade in the sample. We vary the U.S.’s audience cost, \bar{a} , from -6 to 0 while fixing the remaining payoffs estimated using the tML and CMLE from Table 3. For every value of \bar{a} , we compute all equilibria using a line-search method. Then we plot the associated equilibrium probabilities of the U.S. initiating a conflict, p_C , in Figure 6. For all values of \bar{a} considered, there is a unique equilibrium, pictured with the blue circles. The vertical line denotes the estimated value of U.S. audience costs, around -2.7 for the CMLE and -0.6 for the tML. Throughout, we fix the other payoffs at their estimated values, thereby implicitly controlling for the other (belligerent) costs leaders face when choosing to start a crisis. Hence, our analysis allows us to isolate the effects of audience costs from belligerence costs, a traditionally difficult objective when using experiments or reduced form analyses (Kertzer and Brutger 2016).

The figure illustrates three notable results. First, there is substantial difference between the substantive effects from the CMLE and tML. That is, even with the same

Figure 6: Substantive effects of audience costs in U.S. and China dyad, 1991–2000.



theoretical model and data, the choice of estimation procedure matters. Second, given the CMLE, audience costs have a large substantive effect on the probability of threat initiation, covering the entire range between zero and one. These large effects are lost when using the tML estimates. Third, there is a *U*-shaped relationship between audience costs and threat initiation. Leaders only initiate threats when audience costs are very small or quite large. In the former case, leaders do not pay a cost for backing down and do so with impunity. In the latter case, their threats are quite credible, coercing rivals to concede with higher probability.³² With intermediate audience costs, however, leaders almost never threaten rivals with sanctions, as their threats are not credible and backing down entails nontrivial costs.

Notice that if we were to increase the U.S.’s audience costs beginning from the value estimated in the data, then the model predicts an increase in sanction threats toward China. That is, the true value of audience costs tend to fall on the left-hand-side of the *U*-shaped curve, where larger (more negative) audience costs increase the likelihood of interstate threats. This pattern generalizes to other observations in the

³²Appendix G illustrates these additional comparative statics.

data. We compute the marginal effect of making audience costs, \bar{a} , more negative on the equilibrium probability of issuing threats. Conclusively, larger (more negative) audience costs increase the likelihood of states threatening their rivals with sanctions. This result holds in 97% of observations.

6 Conclusion

In this paper, we analyze problems that emerge when fitting games with multiple equilibria to data in international relations. We demonstrate that frequently used maximum likelihood routines perform poorly when estimating the parameters of the canonical crisis-signaling game not only if there are multiple equilibria in the signaling game generating the data but also if the equilibrium is unique. In the former case, without further information, the likelihood function may select the wrong equilibrium when evaluating different parameter guesses, leading to estimates that do not increase in accuracy with more observations. In the latter case, the likelihood function will often be evaluated at parameter guesses under which multiple equilibria exist, leading to similar problems. Imposing a selection rule does not fix these problems, rather, it makes the estimation problem more difficult because it introduces discontinuities into the likelihood. Our analysis should give researchers pause before using these techniques in international relations.

For solutions, we adapt several estimators from the dynamic games literature and show that they are particularly useful in the crisis bargaining context. In a series of experiments and applications, we show that all three perform better than the currently used tML routines, but the CMLE and NPL are consistently good choices. Although the CMLE is far and away the best choice, it requires repeated within-game observations, which may not be appropriate in all situations. Additionally, it requires

specialized constrained optimization software. In general, we propose the following advice when estimating crisis-signaling games.

1. Estimate the game with the PL method, using a flexible first-stage estimator. In our experience, random forests work well.
2. To verify whether bias in the first-stage estimates has affected the second stage, estimate the game with either the NPL or CMLE approach. If these converge, then they should be prioritized. If these do not converge, then the PL results should be prioritized.
3. The tML routine should not be used; it generally performs worse than the other procedures.

We provide R implementations of the PL and NPL estimators in our computational appendix and in the `sigInt` package. This accessibility should help researchers to uncover theoretically informed parameters rather than engaging in more reduced-form analyses.

Finally, the paper raises an important avenue for future research into the empirical crisis-signaling model. Throughout, we have assumed that within each dyad or game, states play the same equilibrium for all within unit observations $t \in \{1, \dots, T\}$. However, it could be the case that the dyad switches equilibria over time; that is, y_{dt} and $y_{dt'}$ were generated from two different equilibria for some $t \neq t'$. If this equilibrium selection rule probabilistically depends on equilibrium properties such as Pareto efficiency or the amount of updating, then scholars could attempt to model the equilibrium selection as in Bajari, Hong and Ryan (2010). If this equilibrium selection rule is arbitrary, then scholars could employ techniques that estimate identification bounds around the coefficients of interest as in Ciliberto and Tamer (2009). Either

approach would relax an assumption in our analysis. A major difficulty in this area is that current work considers games of complete information, as these techniques require repeatedly enumerating the entire set of equilibria. With incomplete information and signaling incentives, this task becomes substantially more complicated.

References

- Aguirregabiria, Victor and Pedro Mira. 2007. "Sequential Estimation of Dynamic Discrete Games." *Econometrica* 75(1):1–53.
- Bajari, Patrick, C. Lanier Benkard and Jonathan Levin. 2007. "Estimating dynamic models of imperfect competition." *Econometrica* 75(5):1331–1370.
- Bajari, Patrick, Han Hong and Stephen P. Ryan. 2010. "Identification and Estimation of a Discrete Game of Complete Information." *Econometrica* 78(5):1529–1568.
- 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.
- Bueno de Mesquita, Ethan and Eric S. Dickson. 2007. "The Propaganda of the Deed: Terrorism, Counterterrorism, and Mobilization." *American Journal of Political Science* 51(2):364–381.
- Chaudoin, Stephen. 2014. "Audience features and the Strategic Timing of Trade Disputes." *International Organization* 68(4):877–911.
- Ciliberto, Federico and Elie Tamer. 2009. "Market Structure and Multiple Equilibria in Airline Markets." *Econometrica* 77(6):1791–1828.
- Crisman-Cox, Casey and Michael Gibilisco. 2018. "Audience Costs and the Dynamics of War and Peace." *American Journal of Political Science* 62(3):566–580.
- De Paula, Aureo. 2013. "Econometric Analysis of Games with Multiple Equilibria." *Annual Review Economics* 5(1):107–131.
- 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.
- Drezner, Daniel W. 1998. "Conflict Expectations and Paradox of Economic Coercion." *International Studies Quarterly* 42:709–731.

- Drezner, Daniel W. 2003. "The hidden hand of economic coercion." *International Organization* 57(3):643–659.
- Ellickson, Paul B. and Sanjog Misra. 2011. "Structural Workshop Paper Estimating Discrete Games." *Marketing Science* 30(6):997–1010.
- Gleditsch, Kristian Skrede, Simon Hug, Livia Isabella Schubiger and Julian Wucherpfennig. 2018. "International Conventions and Nonstate Actors." *Journal of Conflict Resolution* 62(2):346–380. Online First.
- Hafner-Burton, Emilie M. and Alexander H. Montgomery. 2008. "Power or Plenty: How Do International Trade Institutions Affect Economic Sanctions?" *Journal of Conflict Resolution* 52(2):213–242.
- Hart, Robert A., Jr. 2000. "Democracy and the Successful Use of Economic Sanctions." *Political Research Quarterly* 53(2):267–284.
- Hotz, V. Joseph and Robert A. Miller. 1993. "Conditional Choice Probabilities and the Estimation of Dynamic Models." *The Review of Economic Studies* 60(3):497–529.
- Jo, Jinhee. 2011. "Nonuniqueness of the Equilibrium in Lewis and Schultz's Model." *Political Analysis* 19(3):351–362.
- Keohane, Robert O. 1984. *After Hegemony*. Princeton: Princeton University Press.
- Kertzer, Joshua D. and Ryan Brutger. 2016. "Decomposing Audience Costs: Bringing the Audience Back into Audience Cost Theory." *American Journal of Political Science* 60(1):234–249.
- Krustev, Valentin L. and T. Clifton Morgan. 2011. "Ending Economic Coercion: Domestic Politics and International Bargaining." *Conflict Management and Peace Science* 28(4):351–376.
- Kurizaki, Shuhei and Taehee Whang. 2015. "Detecting Audience Costs in International Crises." *International Organization* 69(4):949–980.
- Lacy, Dean and Emerson MS Niou. 2004. "A theory of economic sanctions and issue linkage: The roles of preferences, information, and threats." *Journal of Politics* 66(1):25–42.
- Lapan, Harvey E. and Todd Sandler. 1993. "Terrorism and Signaling." *European Journal of Political Economy* 9:383–397.
- Lektzian, David J. and Christopher M. Sprecher. 2007. "Sanctions, Signals, and Militarized Conflict." *American Journal of Political Science* 51(2):415–431.

- Lektzian, David and Mark Souva. 2003. "The Economic Peace Between Democracies: Economic Sanctions and Domestic Institutions." *Journal of Peace Research* 40(6):641–660.
- 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.
- Martin, Lisa L. 1993. "Credibility, Costs, and Institutions: Cooperation on Economic Sanctions." *World Politics* 45(3):406–432.
- McKelvey, Richard D. and Thomas R. Palfrey. 1998. "Quantal Response Equilibria for Extensive Form Games." *Experimental Economics* 1(1):9–41.
- McLean, Elena V. and Taehee Whang. 2010. "Friends or Foes? Major Trading Partners and the Success of Economic Sanctions." *International Studies Quarterly* 54(2):427–447.
- Pesendorfer, Martin and Philipp Schmidt-Dengler. 2010. "Sequential Estimation of Dynamic Discrete Games: A Comment." *Econometrica* 78:833–842.
- Peterson, Timothy M. 2013. "Sending a Message: The Reputation Effect of US Sanction Threat Behavior." *International Studies Quarterly* 57:670–682.
- Prins, Brandon C. 2003. "Institutional Instability and the Credibility of Audience Costs: Political Participation and Interstate Crisis Bargaining, 1816-1992." *Journal of Peace Research* 40(1):67–84.
- Schelling, Thomas C. 1960. *The Strategy of Conflict*. New York: Oxford University Press.
- Schultz, Kenneth A. 1999. "Do Democratic Institutions Constrain or Inform?" *International Organization* 53(2):233–266.
- Schultz, Kenneth A. and Jeffrey B. Lewis. 2005. "Learning About Learning." *Political Analysis* 14(1):121–129.
- Signorino, Curtis S. 1999. "Strategic Interaction and the Statistical Analysis of International Conflict." *American Political Science Review* 93(2):279–297.
- Signorino, Curtis S. 2003. "Structure and Uncertainty in Discrete Choice Models." *Political Analysis* 11(4):316–344.
- Su, Che-Lin. 2014. "Estimating Discrete-Choice Games of Incomplete Information: Simple static examples." *Quantitative Marketing and Economics* 12(2):167–207.

- Su, Che-Lin and Kenneth L. Judd. 2012. "Constrained Optimization Approaches to Estimation of Structural Models." *Econometrica* 80(5):2213–2230.
- Tarar, Ahmer and Bahar Leventoğlu. 2009. "Public Commitment in Crisis Bargaining." *International Studies Quarterly* 53(3):2009.
- Thomson, Catarina P. 2016. "Public Support for Economic and Military Coercion and Audience Costs." *British Journal of Politics and International Relations* 18(2):407–421.
- Trager, Robert F. and Lynn Vavreck. 2011. "The Political Costs of Crisis Bargaining: Presidential Rhetoric and the Role of Party." *American Journal of Political Science* 55(3):526–545.
- van Damme, Eric. 1996. *Stability and Perfection of Nash Equilibria*. 2nd ed. Berlin: Springer.
- Wand, Jonathan. 2006. "Comparing Models of Strategic Choice: The Role of Uncertainty and Signaling." *Political Analysis* 14(1):101–120.
- Weeks, Jessica L. 2012. "Strongmen and Straw Men." *American Political Science Review* 106(2):326–347.
- Wang, Taehee. 2010a. "Empirical Implications of Signaling Models." *Political Analysis* 18(3):381–402.
- Wang, Taehee. 2010b. "Structural Estimation of Economic Sanctions: From Initiation to Outcomes." *Journal of Peace Research* 47(5):561–573.
- Wang, 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.
- Zeng, Yuleng. 2019. "Bluff to Peace: How Economic Dependence Promotes Peace Despite Increasing Deception and Uncertainty." *Conflict Management and Peace Science* pp. 1–22. Online First.