

ORIGINAL ARTICLE

Estimating signaling games in international relations: problems and solutions

Casey Crisman-Cox^{1*}  and Michael Gibilisco²

¹Department of Political Science, Texas A&M University, College Station, TX, USA and ²Division of Humanities and Social Sciences, California Institute of Technology, Pasadena, CA, USA

*Corresponding author. Email: c.crisman-cox@tamu.edu

(Received 11 January 2019; revised 11 June 2019; accepted 30 July 2019; first published online 9 December 2019)

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.

Keywords: Maximum likelihood estimation (MLE); Structural estimation; Crisis-signaling; Economic sanctions

Political scientists use signaling games across practically all subfields. Scholars of international relations in particular use the models to address questions about economic sanctions, crisis bargaining, escalation in interstate disputes, and terrorism. As a result of this ubiquity, scholars structurally estimate increasingly more complicated signaling models (Lewis and Schultz, 2003; Wand, 2006; Whang, 2010; Whang *et al.*, 2013; Bas *et al.*, 2014; Kurizaki and Whang, 2015). 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., Ellickson and Misra, 2011; De Paula, 2013) 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., Whang *et al.*, 2013; Bas *et al.*, 2014; Kurizaki and Whang, 2015).

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 *et al.*, 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. 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 (Trager and Vavreck, 2011; Gleditsch *et al.*, 2018).¹

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 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

¹Gleditsch *et al.* (2018) refer to the Lewis and Schultz (2003) model as "demanding" in justifying their alternative approach.

provides a feasible pseudo-likelihood (PL) solution to the problem of estimating the game's parameters.

While relatively innocuous in principle, this approach requires accurate estimates of equilibrium quantities. We therefore introduce two additional methods to alleviate the 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 strategies as dyad-specific (game-specific) parameters in a single-stage constrained-ML estimator (CMLE). 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.

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, 2003; Whang *et al.*, 2013). Past work has shown that domestic audiences affect sanction duration and effectiveness (Martin, 1993; Dorussen and Mo, 2001; Krustev and Morgan, 2011) and that audience costs arise when leaders back down from sanction threats (Hart, 2000; Thomson, 2016). Yet scholars have not connected audience costs to the initiator's decision to threaten sanctions.² 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 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 data on economic sanctions.

²Exceptions to this include Peterson's (2013) work on reputation costs and US sanction threats and a brief aside in Whang *et al.* (2013). Similarly, features of domestic audiences help to explain variation in the initiation of World Trade Organization disputes (Chaudoin, 2014).

1. 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 (ϵ_A, ϵ_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 $(\epsilon_A, \epsilon_a, \text{ and } \epsilon_B)$ 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, 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 $\bar{W}_i + \epsilon_i$ at node *SF*. Similarly, when *A* backs down and does not fight, the games end at node *BD* with *A* receiving $\bar{a} + \epsilon_a$ 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:

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

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

and

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

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 essentially best-response functions. Specifically, the functions g and h compute how *A* best responds to *B*'s probability of resisting p_R , and function f captures how *B* best responds to *A*. 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 it is worth noting that there are different ways to specify how private information is introduced in the model. We discuss how the problems we consider appear under some of the most common information structures in Appendix A.




Figure 1. The canonical crisis-signaling game.

2. 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:

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

Each $x_{d(\cdot)}$ vector above contains zero or more explanatory variables.³ 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 $\theta(x_d, \beta^*)$. Additionally, 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 $\{y_{dt}\}_{t=1}^T$.

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 repeated interaction. Technically, this is a standard assumption that is required in the recent empirical literature on estimating games with

³As 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.

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
W_{dA}	-1.9	-1.8
V_{dB}	1	1
W_{dB}	$-2.9 + 0.1x_d$	$-2.45 + 0.1x_d$
$\bar{\sigma}_d$	-1.2	-1.2

incomplete information (Bajari *et al.*, 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, and it 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.

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. 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. While we acknowledge the importance of considering dynamic interactions, this would require a different theoretical model, which is beyond the paper's scope.⁴

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.⁵ There are a few additional things to note about the parameters. First, we normalize the status-quo payoffs S_i and B 's concession payoff to zero, following standard identification assumptions (Lewis and Schultz, 2003; Jo, 2011). 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). 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 gray 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 .

⁴For 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 as private information is transitory. With signaling incentives, a fully dynamic model becomes substantially more intractable.

⁵A 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.

2.1. Problems with current practices

Current best practices 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

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

Under this setup, the log-likelihood takes the form

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

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.

Before proceeding, we first consider potential fixes to the standard ML routine. 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).⁶ As we show in Appendix B, for almost all parameter values, all equilibria of the crisis-signaling game satisfy regularity.⁷ 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.⁸

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.⁹ We return to this point in Appendix G.

2.2. 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 $\hat{\mathbf{p}}_R = (\hat{p}_{1R}, \dots, \hat{p}_{DR})$ and $\hat{\mathbf{p}}_F = (\hat{p}_{1F}, \dots, \hat{p}_{DF})$. 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:

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

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

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

⁶For a formal definition, see Appendix B.

⁷We 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 .

⁸We 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.

⁹Technically, this problem arises because the equilibrium correspondence, and hence likelihood correspondence, is upper, but not lower, hemicontinuous.

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.

2.2.1. *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,

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

and for the k th iteration, set

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

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.

2.3. 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:

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

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 tML with empirical selection 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 \mathbf{p}_R ; as such we need $T > 1$ in order to 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 software; Appendix D contains complete implementation details.

3. 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.¹⁰ Throughout, we use the ordinary implementation of the tML as our baseline for comparison, which uses arbitrary

¹⁰The results we present here are unchanged when we use a more realistic Monte Carlo experiment with multiple covariates in Appendix C.4.

equilibrium selection and Nelder–Mead’s simplex method to find the estimates. These implementation choices match current practices as found in replication archives.¹¹

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{p}_R) and *A* fights (for \hat{p}_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.¹² We repeat this process 1000 times for each pair of *D* and *T* and for each of the parameter settings in Table 1.¹³

The main results of the experiment are summarized in Figures 3 and 4, which compare the logged root-mean-square error (RMSE) of the estimators. The first thing we note is that the tML (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.¹⁴

Contrast these results to those from the other estimators, which generally all improve with more data. The PL (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 (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 (dotted line) tends to be the best. However, this great performance often comes at the cost of decreased convergence rates and non-standard software choices.

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.¹⁵ 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 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.

¹¹As 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, and our conclusions were unchanged. In contrast to the tML, our proposals have continuous log-likelihood functions, so we use the gradient-based Newton–Raphson method for the PL and NPL and a Newton-based interior point method for the CMLE.

¹²The 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 explore how the tML’s performance varies with starting values in Appendix E.

¹³Within 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 .

¹⁴Appendix C contains additional Monte Carlo results relating to the bias, variance, convergence rates, and computational time.

¹⁵Figures 11 and 12 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.

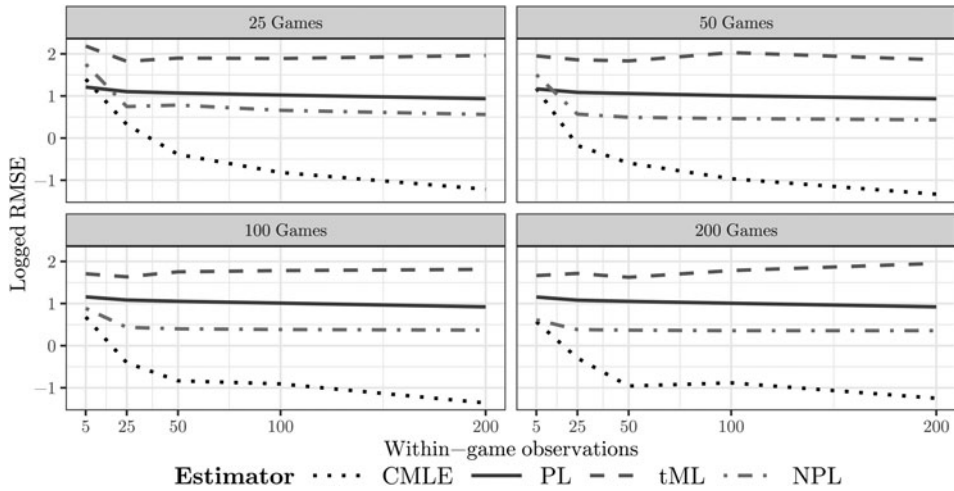


Figure 3. RMSE in signaling estimators with multiple equilibria.

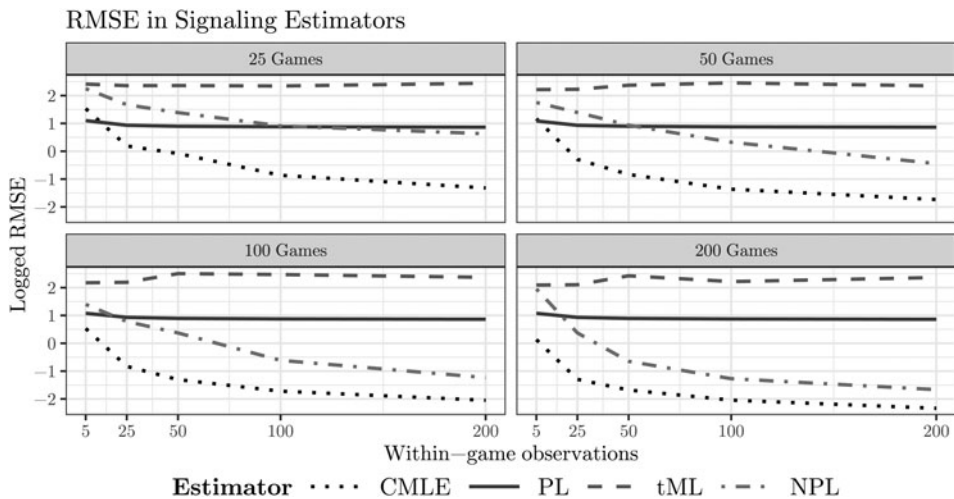


Figure 4. RMSE in signaling estimators with a unique equilibrium.

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. 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 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. 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 oftentimes much better. 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.
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.

4. Application to economic sanctions

Our application is motivated by Whang *et al.* (2013, WMK, hereafter) who 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 19 in Appendix F. 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.¹⁶ 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.¹⁷

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.¹⁸ When there is no episode in a month, we record the

¹⁶WMK specify payoff shocks following Whang (2010), discussed in Appendix A, where they also estimate covariance parameters. These covariance estimates are below 0.07 in magnitude, and WMK fail to reject the null hypothesis that the covariances are equal to zero.

¹⁷This 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 J.

¹⁸Note all the action is coded as occurring in the month when the episode starts. If a play of the game actually unfolds over time, 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.

status quo. When *Final Outcome* records either “acquiescence” by the target or a negotiated settlement, we record the outcome as *B* giving 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 1991–2000.¹⁹

The independent variables, their sources, and how they enter the actors’ payoffs are listed in Table 3 in Appendix F, following the specifications in WMK. All variables are measured on the dyad-decade level as discussed above.²⁰

4.1. Point estimates

Table 2 displays our main results. Each column contains parameter estimates and standard errors using the different estimators. 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 .

Not only does the tML routine return different point estimates, it also produces substantive implications that diverge from the other three estimations. For 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.²¹ 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 illustrate that the differences are due to equilibrium selection and the computational problems addressed above, we conduct two additional analyses in Appendix G. 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.

¹⁹Some countries enter/exit the data in the 1990s so there are 15 dyads where T is between 72 and 96.

²⁰While 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. We show in Appendix I that there is actually very little variation in covariates within each dyad-decade. In Appendix J we check the analysis with five years $T=60$ and one year $T=12$. 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.

²¹Note 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.

Table 2. 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 (0.80)	-0.52 (0.55)
S_A : Dem _A	-0.00 (0.00)	0.00 (0.07)	0.03 (0.06)	0.01 (0.03)
S_A : Contiguity	0.27* (0.10)	0.01 (0.01)	0.05 (0.03)	0.04* (0.01)
S_A : Alliance	-0.06 (0.08)	-0.18* (0.03)	-0.17* (0.08)	-0.16* (0.07)
V_A : Const.	-0.06 (0.08)	-0.33 (0.83)	1.60 (1.30)	1.31* (0.12)
V_A : Costs _A	-0.04 (0.03)	0.36 (0.24)	-0.05 (0.23)	-0.19 (0.17)
C_B : Const.	0.81 (0.91)	-1.12* (0.31)	-2.14* (0.62)	-4.53* (2.24)
C_B : Econ. Dep _B	-0.21 (0.16)	2.32* (0.68)	2.34* (0.90)	2.83* (0.63)
C_B : Costs _B	-0.08* (0.03)	0.08 (0.06)	0.12* (0.05)	0.19* (0.05)
C_B : Contiguity	-0.25* (0.02)	0.13* (0.04)	0.12* (0.03)	0.10* (0.03)
C_B : Alliance	0.10 (0.09)	-0.05 (0.13)	-0.03 (0.10)	-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 (1.20)	-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 (0.86)	-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* (0.04)	0.29* (0.09)
$\bar{\alpha}$: Const.	-0.56 (0.77)	-2.63* (0.09)	-2.64* (0.12)	-2.71* (0.10)
$\bar{\alpha}$: 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 × 120	418 × 120	418 × 120	418 × 120

Notes: *p < 0.05. Asymptotic standard errors in parenthesis, see Appendix D.1 for details.

4.2. 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 (Martin, 1993; Dorussen and Mo, 2001; Drezner, 2003; Whang *et al.*, 2013). In addition, previous work has not connected audiences to the likelihood that leaders threaten sanctions.

We consider the directed dyad in which the US is the initiating state *A* and China is the target state *B* between 1991 and 2000, the most recent decade in the sample. We vary the US's audience

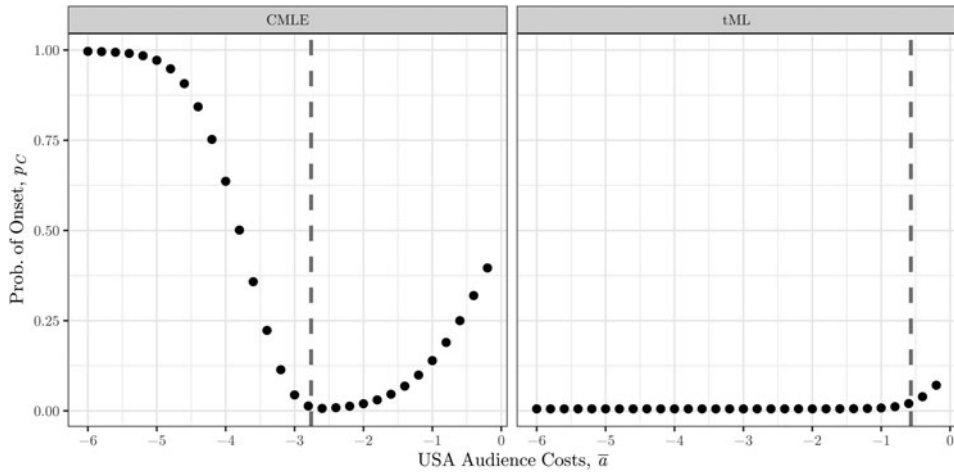


Figure 5. Substantive effects of audience costs in US and China dyad, 1991–2000.

cost, \bar{a} , from -6 to 0 while fixing the remaining payoffs estimated using the tML and CMLE from Table 2. For every value of \bar{a} , we compute all equilibria using a line-search method. Then we plot the associated equilibrium probabilities of the US initiating a conflict, p_C , in Figure 5. For all values of \bar{a} considered, there is a unique equilibrium, pictured with the black circles. The vertical line denotes the estimated value of US 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 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.²² 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 US'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 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 percent of observations.

²²Appendix H illustrates these additional comparative statics.

5. 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 $t \in \{1, \dots, T\}$ and the equilibrium selection is a deterministic function of covariates. However, it could be the case that the dyad switches equilibria over time or equilibria are selected with some noise. Either case would violate an assumption in our analysis, and these would be fruitful directions for future work. A major difficulty in this area is that current econometric work frequently considers games of complete information or other settings whether it is possible to enumerate the entire set of equilibria. With incomplete information and signaling incentives, this task becomes substantially more complicated.

Supplementary material. To view supplementary material for this article, please visit <https://doi.org/10.1017/psrm.2019.58>

Acknowledgments. 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.

References

- Aguirregabiria V and Mira P (2007) Sequential estimation of dynamic discrete games. *Econometrica* 75, 1–53.
 Bajari P, Benkard CL and Levin J (2007) Estimating dynamic models of imperfect competition. *Econometrica* 75, 1331–1370.

- Bas MA, Signorino CS and Whang T** (2014) Knowing one's future preferences: a correlated agent model with bayesian updating. *Journal of Theoretical Politics* **26**, 3–34.
- Chaudoin S** (2014) Audience features and the strategic timing of trade disputes. *International Organization* **68**, 877–911.
- Crisman-Cox C and Gibilisco M** (2018) Audience costs and the dynamics of war and peace. *American Journal of Political Science* **62**, 566–580.
- De Paula A** (2013) Econometric analysis of games with multiple equilibria. *Annual Review Economics* **5**, 107–131.
- Dorussen H and Mo J** (2001) Ending economic sanctions audience costs and rent-seeking as commitment strategies. *Journal of Conflict Resolution* **45**, 395–426.
- Drezner DW** (2003) The hidden hand of economic coercion. *International Organization* **57**, 643–659.
- Ellickson PB and Misra S** (2011) Estimating discrete games. *Marketing Science* **30**, 997–1010.
- Gleditsch KS, Hug S, Schubiger LI and Wucherpfennig J** (2018) International conventions and nonstate actors. *Journal of Conflict Resolution* **62**, 346–380.
- Hart RA** (2000) Democracy and the successful use of economic sanctions. *Political Research Quarterly* **53**, 267–284.
- Hotz VJ and Miller RA** (1993) Conditional choice probabilities and the estimation of dynamic models. *The Review of Economic Studies* **60**, 497–529.
- Jo J** (2011) Nonuniqueness of the equilibrium in Lewis and Schultz's model. *Political Analysis* **19**, 351–362.
- Keohane RO** (1984) *After Hegemony*. Princeton: Princeton University Press.
- Kertzer JD and Brutger R** (2016) Decomposing audience costs: bringing the audience back into audience cost theory. *American Journal of Political Science* **60**, 234–249.
- Krustev VL and Morgan TC** (2011) Ending economic coercion: domestic politics and international bargaining. *Conflict Management and Peace Science* **28**, 351–376.
- Kurizaki S and Whang T** (2015) Detecting audience costs in international crises. *International Organization* **69**, 949–980.
- Lewis JB and Schultz KA** (2003) Revealing preferences: empirical estimation of a crisis bargaining game with incomplete information. *Political Analysis* **11**, 345–367.
- Martin LL** (1993) Credibility, costs, and institutions: cooperation on economic sanctions. *World Politics* **45**, 406–432.
- McKelvey RD and Palfrey TR** (1998) Quantal response equilibria for extensive form games. *Experimental Economics* **1**, 9–41.
- Pesendorfer M and Schmidt-Dengler P** (2010) Sequential estimation of dynamic discrete games: a comment. *Econometrica* **78**, 833–842.
- Peterson TM** (2013) Sending a message: the reputation effect of us sanction threat behavior. *International Studies Quarterly* **57**, 670–682.
- Schelling TC** (1960) *The Strategy of Conflict*. New York: Oxford University Press.
- Schultz KA and Lewis JB** (2005) Learning about learning. *Political Analysis* **14**, 121–129.
- Signorino CS** (1999) Strategic interaction and the statistical analysis of international conflict. *American Political Science Review* **93**, 279–297.
- Su C-L and Judd KL** (2012) Constrained optimization approaches to estimation of structural models. *Econometrica* **80**, 2213–2230.
- Thomson CP** (2016) Public support for economic and military coercion and audience costs. *British Journal of Politics and International Relations* **18**, 407–421.
- Trager RF and Vavreck L** (2011) The political costs of crisis bargaining: presidential rhetoric and the role of party. *American Journal of Political Science* **55**, 526–545.
- van Damme E** (1996) *Stability and Perfection of Nash Equilibria*, 2nd Edn. Berlin: Springer.
- Wand J** (2006) Comparing models of strategic choice: the role of uncertainty and signaling. *Political Analysis* **14**, 101–120.
- Whang T** (2010) Empirical implications of signaling models. *Political Analysis* **18**, 381–402.
- Whang T, McLean EV and Kuberski DW** (2013) Coercion, information, and the success of sanction threats. *American Journal of Political Science* **57**, 65–81.