

Supplementary Information: Estimating Signaling Games in International Relations

Casey Crisman-Cox* Michael Gibilisco†

November 30, 2019

Contents

A Multiple equilibria and alternative information structures	1
A.1 Whang (2010) with multiple equilibria	3
B Regularity and best-response stability	6
C Further Monte Carlo results	10
C.1 Multiple equilibria	10
C.2 Unique equilibrium	13
C.3 Best-response stability	16
C.4 Multivariable Monte Carlo	17
D Implementation details	20
D.1 Standard Errors	20
E Traditional ML and starting values	23
F Model Specification in WMK	27
G Additional problems with traditional ML	28
G.1 Discontinuous likelihood	28
G.2 Sensitivity to implementation choices	29
H Additional Comparative Statics	30
I Decade-level variables	34

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

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

J	Robustness checks	35
J.1	Quarterly Data	35
J.2	Relaxing political relevance	35
J.3	Different levels of aggregation	36
K	R code	39

A Multiple equilibria and alternative information structures

As mentioned, a key assumption in these models concerns how the analyst specifies the variance-covariance matrix of the payoff shocks (Schultz and Lewis 2005; Whang 2010). 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 2010). 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.¹ 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, we also consider Whang’s (2010) specification, below, 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.

¹Schultz and Lewis (2005) provide one justification for the specification used here.

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 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:

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

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. Signorino (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.”

A.1 Whang (2010) with multiple equilibria

Perhaps the most common extension to the Lewis and Schultz model considered here is a specification of the crisis-signaling game from Whang (2010). To demonstrate that the problems and solutions discussed in the main manuscript apply to this more general model, we restate equilibrium choice probabilities, and demonstrate that multiple equilibria can exist. Figure 6 describe the payoffs. Define $\varepsilon_A = (\varepsilon_{SA}, \varepsilon_{VA}, \varepsilon_{WA}, \varepsilon_a)$ and $\varepsilon_B = (\varepsilon_{VB}, \varepsilon_{WB}, \varepsilon_{CB})$. We assume ε_A and ε_B are independent and that ε_i is drawn from a multivariate normal distribution with mean 0 and variance-covariance matrix Σ_i . Furthermore, let θ denote the vector of exogenous parameters of interest, i.e., $\theta = (\bar{a}, C_B, (S_i, V_i, \bar{W}_i, \Sigma_i)_{i=A,B})$. As before, Perfect Bayesian equilibria (equilibria, hereafter) for the game can be represented as choice probabilities, $p = (p_C, p_R, p_F)$. To aid in the explication of equilibrium choice probabilities, we introduce the following notation:

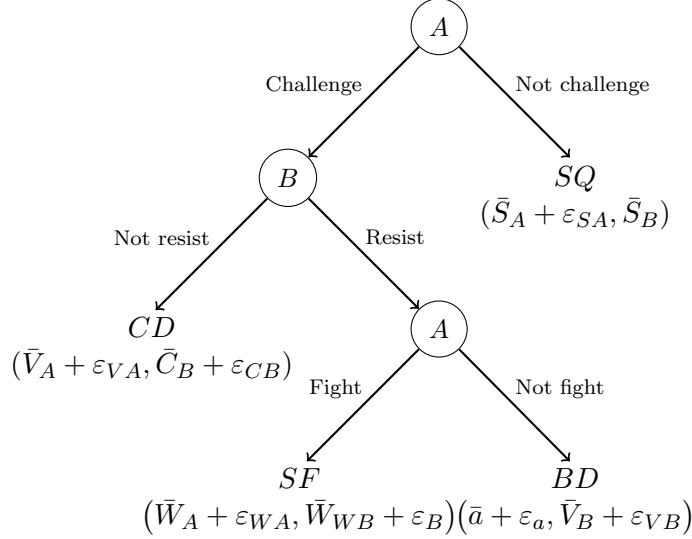
$$\begin{aligned}
 S_A &= \bar{S}_A + \varepsilon_{SA} & C_B &= \bar{C}_A + \varepsilon_{CB} \\
 V_A &= \bar{V}_A + \varepsilon_{VA} & W_B &= \bar{W}_B + \varepsilon_{WB} \\
 W_A &= \bar{W}_A + \varepsilon_{WA} & V_B &= \bar{V}_B + \varepsilon_{VB} \\
 a &= \bar{a} + \varepsilon_a
 \end{aligned}$$

For a fixed vector of choice probabilities, p , define the following:

$$\begin{aligned}
 \Delta U_R^{pF} &= p_F W_B + (1 - p_F) V_B - C_B \\
 \Delta U_{SQ,BD}^{pR} &= S_A - (1 - p_R) V_A - p_R a \\
 \Delta U_{SQ,SF}^{pR} &= S_A - (1 - p_R) V_A - p_R W_A \\
 \Delta U_{SF,BD} &= W_A - a
 \end{aligned}$$

The following result characterizes the equilibrium choice probabilities of this game.

Figure 6: Crisis-signaling game from Whang (2010)



Result 2 (Whang, 2010) *An equilibrium \tilde{p} exists, and \tilde{p} is an equilibrium if and only if it satisfies the following system of equations:*

$$\tilde{p}_C = 1 - \Phi_2 \left(\frac{E[\Delta U_{SQ,BD}^{PR}]}{\sqrt{\text{Var}[\Delta U_{SQ,BD}^{PR}]}, \frac{E[\Delta U_{SQ,SF}^{PR}]}{\sqrt{\text{Var}[\Delta U_{SQ,SF}^{PR}]}, \text{Cor}[\Delta U_{SQ,BD}^{PR}, \Delta U_{SQ,SF}^{PR}]} \right) \equiv g(\tilde{p}_R; \theta), \quad (11)$$

$$\tilde{p}_F = \Phi_2 \left(\frac{E[\Delta U_{SF,BD}]}{\sqrt{\text{Var}[\Delta U_{SF,BD}]}, \frac{E[-\Delta U_{SQ,SF}^{PR}]}{\sqrt{\text{Var}[\Delta U_{SQ,SF}^{PR}]}, \text{Cor}[\Delta U_{SF,BD}, -\Delta U_{SQ,SF}^{PR}]} \right) (g(\tilde{p}_R; \theta))^{-1} \equiv h(\tilde{p}_R; \theta), \quad (12)$$

and

$$\tilde{p}_R = \Phi \left(\frac{E[\Delta U_R^{PF}]}{\sqrt{\text{Var}[\Delta U_R^{PF}]} \right) \equiv f(\tilde{p}_F; \theta). \quad (13)$$

As before fixing a vector of exogenous parameters, θ , equilibria are pinned down by B 's probability of resisting, where \tilde{p}_R satisfies $f \circ h(\tilde{p}_R; \theta) = \tilde{p}_R$. Given an equilibrium probability of resisting, \tilde{p}_R , A 's probabilities of challenging and fighting are defined using Equations 11 and 12, respectively. Notice that ΔU_R^{PF} , $\Delta U_{SQ,BD}^{PR}$, $\Delta U_{SQ,SF}^{PR}$, and $\Delta U_{SF,BD}$ are endogenous quantities. To fully specify the equilibrium choice probab-

ities, the proceeding result states their variances and covariances as functions of the exogenous parameters, θ . To do so, we maintain the following normalizing assumptions from Whang (2010): $\bar{S}_A = 0$, $Var[\varepsilon_{SA}] = 0$, $\bar{C}_B = 0$, $Var[\varepsilon_{CB}] = 0$, $Var[\bar{a}] = 1$, and $Var[\varepsilon_{WB}] = 1$.

Result 3 *Under the normalization assumption, the following hold.*

1. $Var[\Delta U_R^{PF}] = p_F^2 + (1 - p_F)^2 Var[\varepsilon_{VB}] + 2p_F(1 - p_F)Cov[\varepsilon_{VB}, \varepsilon_{WB}]$
2. $Var[\Delta U_{SQ,BD}^{PR}] = (1 - p_R)^2 Var[\varepsilon_a] + p_R^2 + 2p_R(1 - p_R)Cov[\varepsilon_{VA}, \varepsilon_a]$
3. $Var[\Delta U_{SQ,SF}^{PR}] = (1 - p_R)^2 Var[\varepsilon_a] + p_R^2 Var[\varepsilon_{WA}] + 2p_R(1 - p_R)Cov[\varepsilon_{VA}, \varepsilon_{WA}]$
4. $Var[\Delta U_{SF,BD}] = 1 + Var[\varepsilon_{WA}] - 2Cov[\varepsilon_a, \varepsilon_{WA}]$
5. $Cov[\Delta U_{SQ,BD}^{PR}, \Delta U_{SQ,SF}^{PR}] = (1 - p_R)^2 Var[\varepsilon_{VA}] + (1 - p_R)p_R Cov[\varepsilon_{VA}, \varepsilon_{WA}] + p_R(1 - p_R)Cov[\varepsilon_{VA}, \varepsilon_a] + p_R^2 Cov[\varepsilon_a, \varepsilon_{WA}]$
6. $Cov[\Delta U_{SQ,BD}^{PR}, -\Delta U_{SQ,SF}^{PR}] = (1 - p_R)Cov[\varepsilon_{VA}, \varepsilon_{WA}] + p_R Var[\varepsilon_{WA}] - (1 - p_R)Cov[\varepsilon_{VA}, \varepsilon_a] - p_R Cov[\varepsilon_a, \varepsilon_{WA}]$.

Using Results 2 and 3, it is straightforward to modify the PL, NPL, and CMLE estimation routines. One additional difficulty arises, however. Currently, we provide analytical derivatives for optimizers in R. With the additional parameters in Σ_A and Σ_B , additional derivatives would need to be provided or automatic differentiation could be used. We describe the latter approach in Appendix D.

Finally, we provide a numerical example where multiple equilibria arise in this more general model, even outside the assumptions in Lewis and Schultz (2003). For payoffs at terminal nodes, we choose the values in the first column of Table 1. To specify the variance-covariance matrices, σ_A and σ_B , we choose $Var[\varepsilon_{VA}] = 2$, $Var[\varepsilon_{WA}] = Var[\varepsilon_{VB}] = \frac{1}{2}$, $Cov[\varepsilon_{VA}, \varepsilon_a] = Cov[\varepsilon_{VA}, \varepsilon_{WA}] = 0$, and $Cov[\varepsilon_a, \varepsilon_{WA}] = Cov[\varepsilon_{VB}, \varepsilon_{WB}] = -\frac{7}{10}$. Under these parameters, there are three equilibria: $\tilde{p}_R \in \{0.01, 0.63, 0.90\}$.

B Regularity and best-response stability

This Appendix contains the formal arguments for two additional results discussed in the main manuscript. First we define the regularity refinement from Harsanyi (1973) and van Damme (1996). We use $\delta(p_R; \theta)$ to denote the first derivative of $f \circ h$ with respect to p_R given parameters θ .

Definition 1 *An equilibrium \tilde{p}_R is regular if $\delta(\tilde{p}_R; \theta) \neq 1$.*

With this definition we can now state our result concerning the regularity of equilibria.

Result 4 *For almost all θ , all equilibria of the crisis-signaling game are regular.*

To prove the result and subsequent ones, it is more straightforward to work with the function $F : (0, 1) \times \mathbb{R}^8 \rightarrow \mathbb{R}$ such that

$$F(p_R; \theta) = f \circ h(p_R; \theta) - p_R,$$

where \tilde{p}_R is an equilibrium if and only if $F(\tilde{p}_R; \theta) = 0$. We state two intermediary results before proving result 4. The first is from Jo (2011a) and the second is the parameterized Transversality Theorem.

Lemma 1 *For all θ , $\lim_{p_R \rightarrow 0} f \circ h(p_R; \theta) > 0$ and $\lim_{p_R \rightarrow 1} f \circ h(p_R; \theta) < 1$.*

Thus, there are no equilibria at the boundaries. In addition, for any fixed θ , there exists $\varepsilon > 0$ and $\nu > 0$ such that $F(\varepsilon; \theta) > 0$ and $F(1 - \nu; \theta) < 0$

Theorem 1 (Transversality Theorem) *Consider an open set $X \subseteq \mathbb{R}^n$. Let $L : X \times \mathbb{R}^s \rightarrow \mathbb{R}^n$ be continuously differentiable. Assume that the Jacobian $D_{(x,y)}L$ has rank n for all $(x, y) \in X \times \mathbb{R}^s$ such that $L(x, y) = 0$. Then, for almost all $y' \in \mathbb{R}^s$, the Jacobian $D_x L$ has rank n for all $x \in X$ such that $L(x, y') = 0$.*

Proof of Result 4. Note that \tilde{p}_R is a regular equilibrium if and only if $D_{p_R} F(p_R; \theta) \neq 0$. To prove Result 4, we verify the conditions of the Transversality condition, where in our application, $L = F$ and $(x, y) = (p_R; \theta)$, which means $n = 1$ and $s = 8$. First,

note that F is continuously differentiable, because $f \circ h$ is the composition of normal cumulative distribution functions and polynomial functions, and F is defined over the open interval $(0, 1)$.

Third and finally, we show that $D_{(p_R; \theta)} F(p_R; \theta)$ has at least one non-zero element (i.e., rank 1) when $F(p_R; \theta) = 0$. To do this, we show a stronger result: for all $(p_R; \theta)$, $D_{(p_R; \theta)} F(p_R; \theta) \neq 0$. To see this, consider $D_{\bar{W}_B} F(p_R; \theta)$. By Result 1, the functions g and h are constant in parameter \bar{W}_B , that is, $D_{\bar{W}_B} g(p_R; \theta) = D_{\bar{W}_B} h(p_R; \theta) = 0$. Then we have

$$\begin{aligned}
D_{\bar{W}_B} F(p_R; \theta) &= D_{\bar{W}_B} f \circ h(p_R; \theta) \\
&= D_{\bar{W}_B} \Phi \left(\frac{h(p_R; \theta) \bar{W}_B + (1 - h(p_R; \theta)) V_B - C_B}{h(p_R; \theta)} \right) \\
&= D_{\bar{W}_B} \Phi \left(\bar{W}_B + \frac{(1 - h(p_R; \theta)) V_B - C_B}{h(p_R; \theta)} \right) \\
&= \phi \left(\bar{W}_B + \frac{(1 - h(p_R; \theta)) V_B - C_B}{h(p_R; \theta)} \right) \\
&\neq 0,
\end{aligned}$$

which implies $D_{(p_R; \theta)} F(p_R; \theta) \neq 0$ as required. \square

Although the regularity refinement does not generically reduce the number of equilibria, showing that all the equilibria are regular is advantageous for applied empirical research. Regular equilibria can be implicitly expressed as continuous functions of parameters. This property is particularly important in empirical analyses: if we uncover noisy, but sufficiently accurate estimates of θ , then equilibrium choice probabilities will be close to their true values as well. In addition, comparative statics (predicted probabilities) on regular equilibria will be well behaved, i.e., the equilibrium will not vanish if we vary the data or parameters by some small amount.

Our next result focuses on best response iteration. Before stating the result, we define best-response stable and best-response unstable equilibria.

Definition 2 An equilibrium \tilde{p}_R is best-response stable if there exists $\varepsilon > 0$ such that for all $p_R^0 \in (\tilde{p}_R - \varepsilon, \tilde{p}_R + \varepsilon)$ the sequence

$$p_R^k = f \circ h(p_R^{k-1}; \theta), \quad k \in \mathbb{N}$$

converges to \tilde{p}_R .

The next definition introduces best-response unstable equilibria, which is not simply the negation of Definition 2.

Definition 3 An equilibrium \tilde{p}_R is best-response unstable if there exists $\varepsilon > 0$ such that for all $p_R^0 \in (\tilde{p}_R - \varepsilon, \tilde{p}_R + \varepsilon)$, with $p_R^0 \neq \tilde{p}_R$, the sequence

$$p_R^k = f \circ h(p_R^{k-1}; \theta), \quad k \in \mathbb{N}$$

leaves the interval $(\tilde{p}_R - \varepsilon, \tilde{p}_R + \varepsilon)$ at least once. That is, there exists $n \in \mathbb{N}$ such that $p_R^n \notin (\tilde{p}_R - \varepsilon, \tilde{p}_R + \varepsilon)$

With these definitions, we are now ready to state Results 5.

Result 5 If all equilibria are regular, then following hold:

1. There is a finite number of equilibria.
2. If there are multiple equilibria, then there exists a best-response unstable equilibrium.

To prove Result 5, we need an intermediate result, that is standard in nonlinear dynamics and fixed point iteration. See Theorem 6.5 in Holmgren (1994).

Theorem 2 Consider an equilibrium \tilde{p}_R . If $|\delta(\tilde{p}_R; \theta)| < 1$, then \tilde{p}_R is best-response stable. If $|\delta(\tilde{p}_R; \theta)| > 1$, then \tilde{p}_R is best-response unstable.

To end this Appendix, we prove Result 5.

Proof of Result 5(1). By assumption all equilibria are regular, which implies $D_{p_R}F(\tilde{p}_R; \theta) \neq 0$ for all \tilde{p}_R such that $F(\tilde{p}_R; \theta) = 0$. Then the Implicit Function Theorem implies that every equilibrium \tilde{p}_R is locally isolated. Because F is continuous, it has closed level sets, so the set of equilibria is closed. Because equilibria fall within the interval $(0, 1)$, the set of equilibria is bounded, and therefore compact. As a compact set of locally isolated points, the equilibrium set is finite. \square

Proof of Result 5(2). Assume all equilibria are regular. By Result 5(1), we can write the set of equilibria as $\{\tilde{p}^{[1]}, \dots, \tilde{p}^{[k]}\}$ where k is the number of equilibria. Order the set such that $a < b$ implies $\tilde{p}^{[a]} < \tilde{p}^{[b]}$. By assumption, $k \geq 2$, and we claim that $\tilde{p}^{[2]}$ is best-response unstable. To do so, the proof consists of two steps. In step 1, we prove that $\delta(\tilde{p}^{[1]}; \theta) < 1$. In step 2, we prove that $\delta(\tilde{p}^{[2]}; \theta) > 1$, which, by Theorem 2, implies that $\tilde{p}^{[2]}$ is best-response unstable.

Step 1: Suppose not. That is, suppose $\delta(\tilde{p}^{[1]}; \theta) \geq 1$. By regularity, $\delta(\tilde{p}^{[1]}; \theta) > 1$. Because F is continuously differentiable and $D_{p_R} F = \delta(\tilde{p}_R^{[1]}; \theta) - 1$, there exists $\varepsilon > 0$ such that F is strictly increasing on the interval $(\tilde{p}_R^{[1]} - \varepsilon, \tilde{p}_R^{[1]})$. Because $F(\tilde{p}_R^{[1]}; \theta) = 0$, this implies that there exists a $p'_R \in (\tilde{p}_R - \varepsilon, \tilde{p}_R^{[1]})$ such that $F(p'_R; \theta) < 0$. By Lemma 1, there exists $\nu \in (0, p'_R)$ such that $F(\nu; \theta) > 0$. Then the Intermediate Value Theorem Implies that there exists a $\tilde{p}_R \in (\nu, p'_R)$ such that $F(\tilde{p}_R; \theta) = 0$, but this contradicts the assumption that $\tilde{p}_R^{[1]}$ is the smallest equilibrium. Hence, we conclude that $\delta(\tilde{p}^{[1]}; \theta) < 1$

Step 2: Suppose not. That is, suppose $\delta(\tilde{p}^{[2]}; \theta) \leq 1$. Because all equilibria are regular, $\delta(\tilde{p}_R^{[2]}; \theta) < 1$, implying $D_{p_R} F(\tilde{p}_R^{[2]}; \theta) < 0$. This, along with the facts that F is continuously differentiable and $F(\tilde{p}_R^{[2]}; \theta) = 0$, implies there exists (arbitrarily small) $\varepsilon > 0$ such that $F(\tilde{p}_R^{[2]} - \varepsilon; \theta) > 0$.

In Step 1, we showed that $\delta(\tilde{p}_R^{[1]}; \theta) < 1$. Because $F(\tilde{p}_R^{[1]}; \theta) = 0$, there exists (arbitrarily small) $\nu > 0$ such that $F(\tilde{p}_R^{[1]} + \nu; \theta) < 0$ because F is continuously differentiable. So we have $F(\tilde{p}_R^{[2]} - \varepsilon; \theta) > 0$ and $F(\tilde{p}_R^{[1]} + \nu; \theta) < 0$. Then by the Intermediate Value Theorem there exists an equilibrium \tilde{p}'_R such that

$$\tilde{p}_R^{[1]} + \nu < \tilde{p}'_R < \tilde{p}_R^{[2]} - \varepsilon.$$

But this contradicts the assumption that $\tilde{p}_R^{[2]}$ is the second smallest equilibrium. Thus, we conclude $\delta(\tilde{p}_R^{[2]}; \theta) > 1$. As such, $\tilde{p}_R^{[2]}$ is best-response unstable by Theorem 2. \square

C Further Monte Carlo results

C.1 Multiple equilibria

This appendix contains additional results from the Monte Carlo experiment where the data are generated under parameters that are consistent with multiple equilibria. A single covariate determines the equilibrium selection. The parameter values used to generate the data can be found in Table 1. Here we consider the estimators' bias, variance, rate of convergence, and computation time. Root mean-squared error is presented in the main text.

Figure 7: Bias in signaling estimators with multiple equilibria.

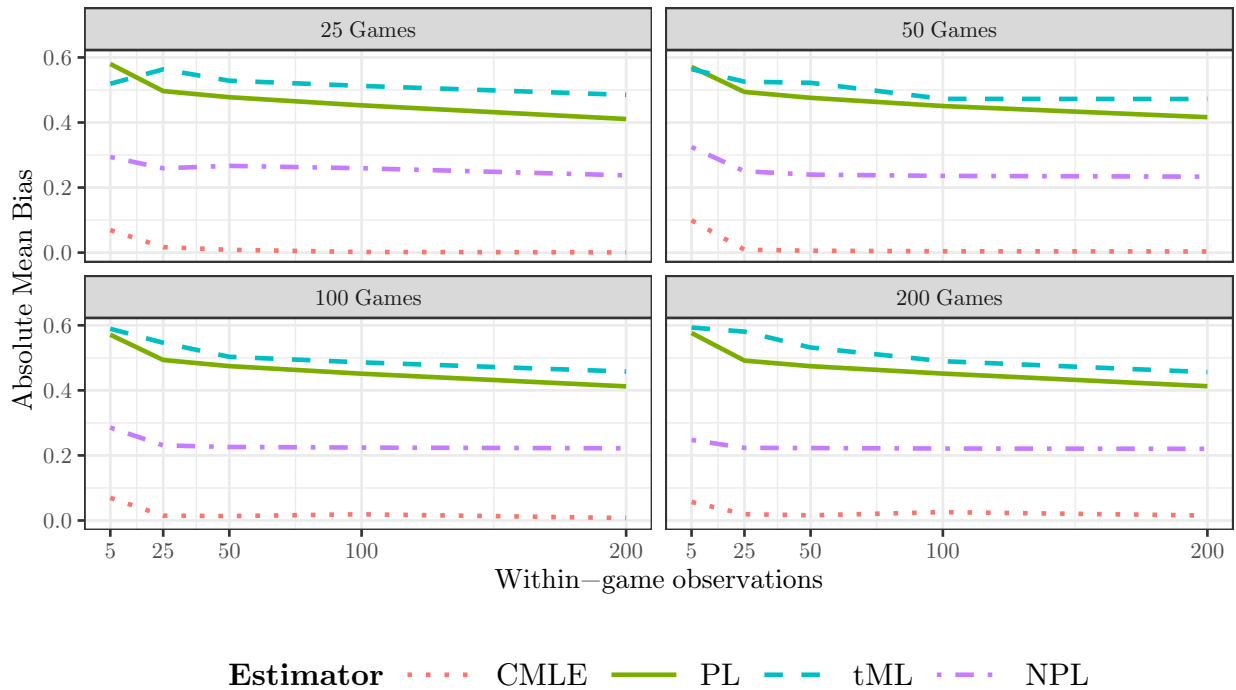


Figure 8: Variance in signaling estimators with multiple equilibria.




Figure 9: Convergence rates in signaling estimators with multiple equilibria.





Figure 10: Computational time in signaling estimators with multiple equilibria.



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

C.2 Unique equilibrium

This appendix contains additional results from the Monte Carlo experiment where the data are generated from a version of the game with a unique equilibrium. The parameter values used to generate the data can be found in the final column of Table 1. Here we consider the estimators' bias, variance, computation time, and rate of convergence.

Figure 11: Bias in signaling estimators with a unique equilibrium.




Figure 12: Variance in signaling estimators with a unique equilibrium.




Figure 13: Computational time in signaling estimators with a unique equilibrium.

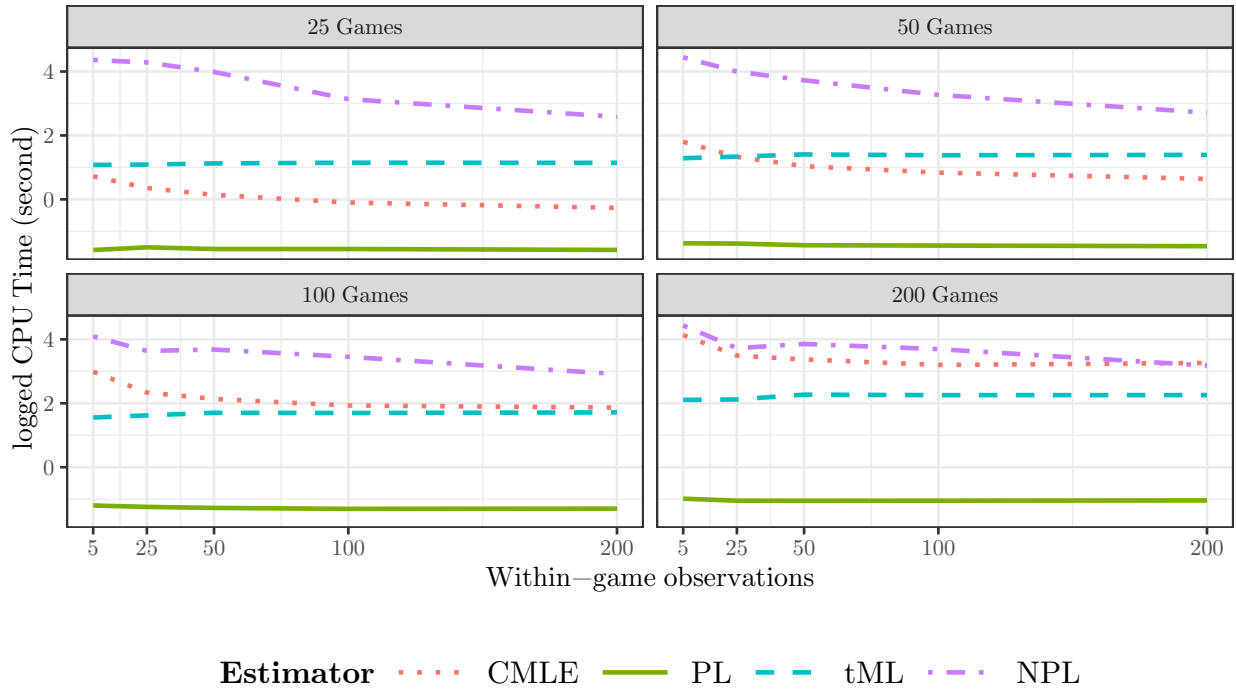



Figure 14: Convergence rates in signaling estimators with a unique equilibrium.




C.3 Best-response stability

The best performing solutions make use of best response functions, which begs the question: How sensitive are the estimators to best-response unstable equilibria? To answer this question, we conduct another Monte Carlo experiment. Here, we assume payoffs are generated as in the multiple setting in Table 1, and the equilibrium selection rule follows the left-hand graph in Figure 2. Let $q \in [0, 1]$ denote the percentage of unstable equilibria. For $q \cdot D$ dyads, x_d is drawn from a uniform distribution over the interval $(\frac{1}{3}, \frac{2}{3})$. For the remaining $D - q \cdot D$ observations, x_d is drawn uniformly from the intervals $(0, \frac{1}{3})$ or $(\frac{2}{3}, 1)$ with equal probability. Using Theorem 2, the middle equilibrium, i.e., the one selected when $x_d \in (\frac{1}{3}, \frac{2}{3})$, is unstable, while the other equilibria are best-response stable. As we vary q from 0 to 1, we analyze how the estimators' performance varies as the data are generated with a larger proportion of best-response unstable equilibria. In this experiment, we set $D = 200$ and $T = 1000$, which means there is a large amount of data as to better isolate the effects of unstable equilibria. For all values of q , we draw x_d , select the corresponding equilibria, and estimate the model 1,000 times. We expect the PL and NPL to perform worse as q approaches 1.²

Figure 15 summarizes the results, where we vary the percentage of unstable equilibria along the horizontal axis and plot log RMSE along the vertical axis. Unsurprisingly, the NPL performs much worse in terms of RMSE as more data are generated from the unstable equilibrium. The PL, tML, and CMLE all get slightly worse as this proportion increases, but this effect is far less pronounced. Despite the fact that the NPL is designed to struggle here, it still outperforms the PL when less than 40% of the data are from unstable equilibria. Of further note, both the PL and NPL still outperform the tML across the board, despite their reliance on best-response

²Following results from Kasahara and Shimotsu (2012) we check that the spectral radius of the Jacobian of the best-response function is greater than 1 under these conditions. We find that the NPL should struggle in all situations where $q \geq 0.01$.

Figure 15: RMSE in signaling estimators with more unstable equilibria.




iteration.

Beyond the potential for statistical problems, we also want to consider the computational issues that arise when unstable equilibria dominate in the data. This trend is illustrated in Figure 16, where horizontal axis is the proportion of observations with unstable equilibria and the vertical axis is the proportion of successful Monte Carlo iterations. Notice, convergence rates of all estimators, besides the PL, decrease once the proportion of unstable equilibria approaches 60–80%. Thus, conditional on converging, the estimators return results with fairly reasonable RMSE even with a large proportion of unstable equilibria. They are all generally less likely to converge when unstable equilibria permeate the data, however.

C.4 Multivariable Monte Carlo

In this section we consider a different Monte Carlo experiment designed to better capture real world situations. Specifically, we use our application to economic

Figure 16: Convergence rates in signaling estimators with unstable equilibria.



sanctions data to construct an experiment with many variables that appear across different utilities.

To build the experiment we use the same specification and independent variables as the economic sanctions application above. We then take the CMLE estimates from Table 2 and fix them as the true parameter values. Using these parameters and the original independent variables we generate a new dependent variable (of length 120 for each dyad) for each Monte Carlo simulation and refit the model using tML, PL, NPL, and CMLE.³ For each parameter we then compute the root mean-squared error (RMSE).

Table 3 shows the relative performance of each of our proposed methods to the tML. Here, values less than one mean that our estimator does better than the traditional method, while values greater than one mean that tML performed better on estimating that parameter. Values close to zero mean that our approach does a lot better than the tML. All cases where the tML does better are bolded, which happens

³Note that in the case of CMLE, this is equivalent to using a parametric bootstrap to build standard errors.

in only three cases out of sixty (about 5%).

Overall, the PL has a little trouble with a few parameters in B 's utilities, which is consistent with our other Monte Carlo results. The NPL and CMLE both do very well compared to the tML. The last row in Table 3 shows the relative improvement in the multivariate RMSE, where we see that all three of our approaches are better than the tML in this experiment.

Table 3: Relative RMSE of Estimates Compared to tML

	PL	NPL	CMLE
S_A : Econ. Dep $_A$	0.48	0.87	0.63
S_A : Dem $_A$	0.18	0.36	0.37
S_A : Contiguity	0.11	0.06	0.05
S_A : Alliance	0.17	0.20	0.19
V_A : Const.	0.79	0.35	0.27
V_A : Costs $_A$	0.59	0.40	0.33
C_B : Const.	0.75	0.51	0.45
C_B : Econ. Dep $_B$	0.35	0.28	0.21
C_B : Costs $_B$	0.42	0.26	0.21
C_B : Contiguity	1.43	0.14	0.11
C_B : Alliance	0.34	0.37	0.31
\bar{W}_A : Const.	0.03	0.03	0.03
\bar{W}_A : Econ. Dep $_A$	0.99	1.01	0.40
\bar{W}_A : Dem $_A$	0.07	0.23	0.25
\bar{W}_A : Cap. Ratio	0.22	0.20	0.11
\bar{W}_B : Const.	1.18	0.81	0.68
\bar{W}_B : Dem $_B$	0.43	0.35	0.18
\bar{W}_B : Cap. Ratio	0.80	0.51	0.34
\bar{a} : Const.	0.03	0.04	0.03
\bar{a} : Dem $_A$	0.09	0.26	0.28
Multivariate RMSE	0.80	0.55	0.46

D Implementation details

In our economic sanctions application we fit the CMLE using the program IPOPT (Interior Point OPTimizer), which is an open-source optimizer designed to handle large scale problems (Wächter and Biegler 2006). In trials, IPOPT has better performance properties than other optimizers such as sequential quadratic solvers (found in Python’s `scipy.optimize` module), a version of the Augmented Lagrangian Method (from R’s `alabama` package), and alternative interior-point methods (MATLAB’s `fmincon`).

The main difficulty in using interior-point methods is that they require an accurate second derivative of the Lagrangian associated with the problem in Equation 10. We find that finite difference approximations are insufficient. As such, we use the program ADOL-C, software for algorithmic differentiation (AD) (Griewank, Juedes and Utke 1996; Walther and Griewank 2012), to precisely compute the Hessian. The AD software allows us to only supply the log-likelihood and constraint function from Equation 10. The AD program repeatedly applies the chain rule to our functions to compute first- and second-order derivatives. In all uses of the CMLE, we use IPOPT and ADOL-C within Python 2.7.15 on Ubuntu 18.04 by calling the `pyipopt` module developed by Xu (2014) and the `pyadolc` module developed by Walter (2014), respectively.

D.1 Standard Errors

Following current practices, the tML standard errors are from the outer product of gradients estimator (sometimes called the BHHH estimator). Asymptotic standard errors for the other approaches are provided below.

The asymptotic standard errors for the PL estimates follow from standard results on two-step maximum likelihood estimation (e.g., Murphy and Topel 1985), such that

$$\widehat{\text{Var}}(\hat{\beta}_{PL}) = \hat{\Omega}_\beta^{-1} + \hat{\Omega}_\beta^{-1} \hat{\Omega}_p \hat{\Sigma} \hat{\Omega}_p^T \hat{\Omega}_\beta^{-1}.$$

Here $\hat{\Omega}_\beta$ and $\hat{\Omega}_p$ are outer product of gradients estimators and $\hat{\Sigma}$ is the estimated first-stage covariance matrix, such that

$$\begin{aligned} \hat{\Omega}_\beta &= J_{PL}^\beta(\hat{\beta}_{PL} | \hat{\mathbf{P}}_{\mathbf{R}}, \hat{\mathbf{P}}_{\mathbf{F}}, Y, X)^T J_{PL}^\beta(\hat{\beta}_{PL} | \hat{\mathbf{P}}_{\mathbf{R}}, \hat{\mathbf{P}}_{\mathbf{F}}, Y, X) \\ \hat{\Omega}_p &= J_{PL}^\beta(\hat{\beta}_{PL} | \hat{\mathbf{P}}_{\mathbf{R}}, \hat{\mathbf{P}}_{\mathbf{F}}, Y, X)^T J_{PL}^{\mathbf{PR}, \mathbf{PF}}(\hat{\mathbf{P}}_{\mathbf{R}}, \hat{\mathbf{P}}_{\mathbf{F}} | \hat{\beta}_{PL}, Y, X) \\ \hat{\Sigma} &= \widehat{\text{Var}}(\hat{\mathbf{P}}_{\mathbf{R}}, \hat{\mathbf{P}}_{\mathbf{F}}), \end{aligned}$$

where J_{PL}^x is the Jacobian of the PL likelihood with respect to x . In our applications, we use a non-parametric bootstrap to produce $\hat{\Sigma}$, which is the covariance matrix of the first-stage (random forest) estimates.

Aguirregabiria and Mira (2007) provide asymptotic standard errors for the NPL estimates that converges after \mathbf{n} iterations as

$$\widehat{\text{Var}}(\hat{\beta}_{NPL}) = \left(\hat{\Omega}_\beta + \hat{\Omega}_p (\mathbf{I} - \hat{\psi}_p^T)^{-1} \hat{\psi}_\beta \right)^{-1} \hat{\Omega}_\beta \left(\hat{\Omega}_\beta + \hat{\psi}_\beta^T (\mathbf{I} - \hat{\psi}_p)^{-1} \hat{\Omega}_p^T \right)^{-1}.$$

Here $\hat{\Omega}_\beta$ and $\hat{\Omega}_p$ are still outer product of gradients estimators, but they are now given as

$$\begin{aligned} \hat{\Omega}_\beta &= J_{PL}^\beta(\hat{\beta}_{NPL} | \hat{\mathbf{P}}_{\mathbf{R}, \mathbf{n}}, \hat{\mathbf{P}}_{\mathbf{F}, \mathbf{n}}, Y, X)^T J_{PL}^\beta(\hat{\beta}_{NPL} | \hat{\mathbf{P}}_{\mathbf{R}, \mathbf{n}}, \hat{\mathbf{P}}_{\mathbf{F}, \mathbf{n}}, Y, X) \\ \hat{\Omega}_p &= J_{PL}^\beta(\hat{\beta}_{NPL} | \hat{\mathbf{P}}_{\mathbf{R}, \mathbf{n}}, \hat{\mathbf{P}}_{\mathbf{F}, \mathbf{n}}, Y, X)^T J_{PL}^{\mathbf{PR}, \mathbf{PF}}(\hat{\mathbf{P}}_{\mathbf{R}, \mathbf{n}}, \hat{\mathbf{P}}_{\mathbf{F}, \mathbf{n}} | \hat{\beta}_{NPL}, Y, X), \end{aligned}$$

while $\hat{\psi}_p$ and $\hat{\psi}_\beta$ are the Jacobians of the best-response function with respect to $(\mathbf{P}_{\mathbf{R}}, \mathbf{P}_{\mathbf{F}})$ and β , respectively, and evaluated at the NPL estimates.

Finally, the asymptotic standard errors for the CMLE are computed using Silvey (1959, Lemma 6), such that

$$\widehat{\text{Var}} \left(\left(\hat{\beta}, \hat{\mathbf{P}}_{\mathbf{R}} \right)_{CMLE} \right) = \left[\begin{array}{cc} \hat{H} + \hat{\omega}^T \hat{\omega} & -\hat{\omega}^T \\ -\hat{\omega} & \mathbf{0} \end{array} \right]_{1,2,\dots,D+\ell}^{-1}.$$

Here, \hat{H} is the Hessian of the CMLE's log-likelihood with respect to the full parameter vector, evaluated at the estimates, ℓ is the length of β , and

$$\hat{\omega} = J_{f_{oh}}^{(\beta, \mathbf{P}_{\mathbf{R}})} \left(\left(\hat{\mathbf{P}}_{\mathbf{R}}, \hat{\beta} \right)_{CMLE} | Y, X \right)$$

is the Jacobian of the CMLE's equilibrium constraint with respect to the full parameter vector and evaluated at the estimates. Note that the total size of the matrix is $2D + \ell$, while the covariance matrix of the full parameter vector is composed of only the first $D + \ell$ rows and columns. The remaining entries relate to the D Lagrange multipliers used to solve the constrained optimization problem.

E Traditional ML and starting values

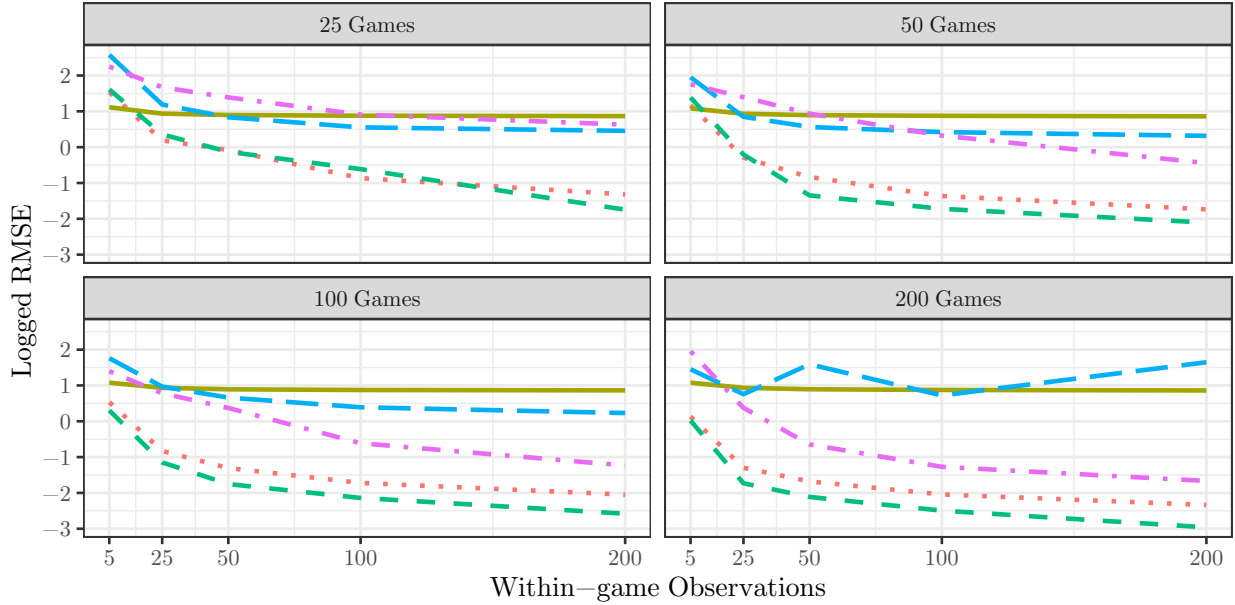
In this appendix, we are interested in the effects that starting values have on the tML’s performance. To do this, we focus on two questions: why has past work found that the tML is consistent when the data are generated by a unique equilibrium, and can the tML be improved by just giving it better starting values?

Regarding the first question, our Monte Carlo experiments demonstrate that the tML may not be consistent even when there is a unique equilibrium in the signaling game that is generating the data. The reason such problems arise is that the maximization routine will oftentimes evaluate the likelihood function at a guess of the parameters where multiple equilibria arise. In this case, the traditional approach will select an equilibrium in an ad-hoc fashion, which may encourage the maximization routine to move away from the correct parameters. This may be surprising as both Jo (2011a) and Bas, Signorino and Whang (2014) conduct similar Monte Carlo experiments and conclude that the tML performs well when the data were generated with parameters that admit a unique equilibrium.⁴

To the best of our knowledge, the differences arise from starting values. In our study, starting values for θ were drawn from a standard uniform distribution. In Jo (2011a), the starting values are the true values from the data generating process (Jo 2011b). Although we were not able to locate replication materials from Bas, Signorino and Whang (2014), we do conduct an additional Monte Carlo experiment to investigate the possibility that differences in starting values lead to different results. To do this, we reproduce our Monte Carlo experiments from the main text, but now we use different starting values for the tML. First, we follow Jo (2011a) and use the true data generating values as starting values to see if this accounts for the differences

⁴Jo (2011a, p. 357) writes “It is easy to see that when there is a unique equilibrium, the estimates get closer to their true values as the number of observations increases.” Bas, Signorino and Whang (2014, p. 26) write “All coefficients on average are estimated very close to the true parameter values, and the accuracy of the estimates increases as the sample size increases.”


Figure 17: RMSE with a unique equilibrium and different starting values.



we observed between our results and hers. Second, we use the PL estimates as starting values to explore if our Monte Carlo results are driven by choices over starting values. The motivation for this second question is based on the fact that we use the PL as a launching point for the other methods we consider. The NPL builds on the PL by construction, and we use the PL estimates as starting values for the CMLE in order to improve the stability of the constrained optimization problem. These approaches naturally raise the question of whether the tML can be improved by starting it at the PL estimates.

Figure 17 graphs the logged RMSE of the estimation procedures as we vary the number of dyads D and the number of observations T . In a similar manner, Figure 18 reports the logged RMSE for an experiment where there are multiple equilibria at the true parameters. Note, that the PL, NPL, and CMLE results in these figures are identical to the results reported in Figures 4 and 3, respectively. There are three major takeaways.

Figure 18: RMSE with multiple equilibria and different starting values.



First, starting the tML procedure at the true values greatly improves the tML’s performance. This benefit is most pronounced when there is a unique equilibrium at the true values. This explains the consistency findings from Jo (2011a) and Bas, Signorino and Whang (2014). However, these values are not known *a priori* in practice, which limits the usefulness of this result.

Second, starting the tML procedure at the PL estimates offers some improvement over the results in the main text. However, the improvements are not enough to make the tML a justifiable method. In practice, the tML only notably better than the PL when it has: (i) informative starting values, (ii) there is a unique equilibrium in the data generating game, *and* (iii) there are many within-game observations. If *any* of these three conditions fails, the PL tends to be at least as well and is sometimes better than the tML while the NPL is almost always better and the CMLE is always better. Given that we can never know if condition (ii) holds, the tML is never a good choice.

Third, if all three of the above conditions hold, the CMLE is a better choice than the tML with PL starting values. The only approach that rivals the CMLE when there are multiple within-game observations is when conditions (ii) and (iii) hold *and* the procedure is started at the true parameter values. Of course, we never have the true values to use as a starting point, and we still never know if condition (ii) holds. As such, our main conclusions hold even when we try to improve the tML by starting it at the PL values.

F Model Specification in WMK

Figure 19: Economic sanctions as a signaling interaction.

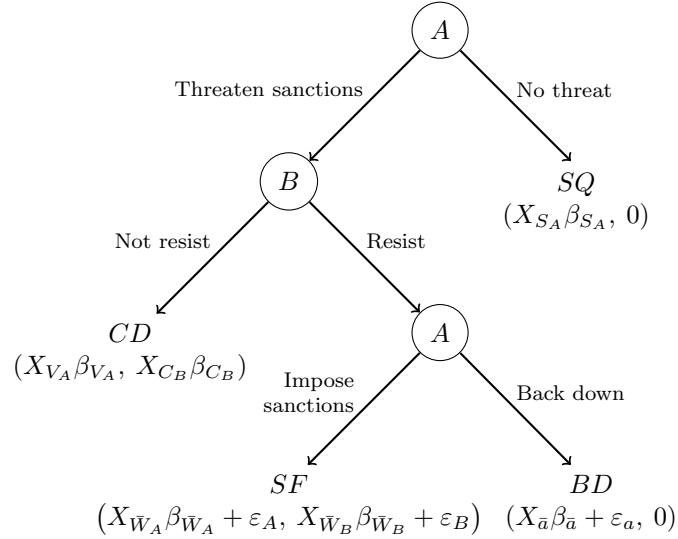


Table 4: Variables in the economic sanctions model


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

G Additional problems with traditional ML

G.1 Discontinuous likelihood

As mentioned in the main text, *ad hoc* equilibrium selection is one possible solution to the tML’s troubles. However, such a modification introduces discontinuities into the tML’s log-likelihood function. We demonstrate this in Figure 20. Here, we graph the log-likelihood as a function of the parameter $\hat{\beta}_{\overline{W}_B}^1$ (the true value is $\beta_{\overline{W}_B}^1 = -2.9$), where data are generated using the values in Table 1, column 1, the equilibrium selection in Figure 2, and $D \in \{1, 10\}$ with $T = 200$.

Figure 20: Log-likelihood function with an imposed selection rule



The main thing to note here is that not only are there discontinuities in the log-likelihood, but also that the number of discontinuities is increasing in D . In many international relations studies, the number of dyads under consideration can be in the hundreds or thousands. Having a likelihood function with that many jumps in it is extremely difficult to optimize using ordinary means. Global methods are a possibility here, but the computational cost is cost-prohibitive compared to the PL, NPL, or CMLE.

Table 5: tML with different solvers and starting values

	tML <i>Newton Solver</i> Model 1	tML <i>Select Largest Eq.</i> Model 5	tML <i>PL start values</i> Model 6
\bar{a} : Const.	-0.56 (0.77)	-0.76* (0.14)	-2.62* (0.11)
\bar{a} : Dem _A	-0.00 (0.01)	0.06* (0.01)	-0.03 (0.04)
Log L	-4102.76	-4302.08	-3938.73
$D \times T$	418×120	418×120	418×120

Notes: * $p < 0.05$

Standard Errors in Parenthesis

G.2 Sensitivity to implementation choices

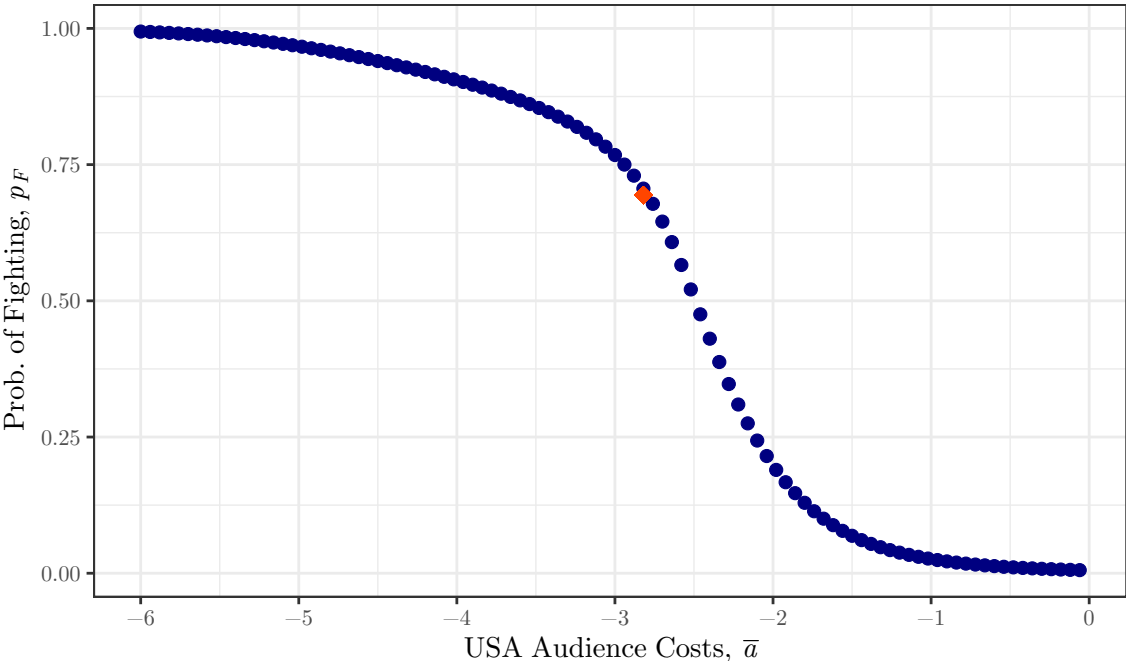
Table 5 illustrates the sensitivity of the tML routine to different implementation choices. In the first column, we reprint Model 1 from the main text where the tML routine uses a Newton solver to compute an equilibrium for each dyad d and each guess of the parameter value θ . In Model 5, we change the equilibrium selection method used in the tML. Here, for each guess of the parameter values and for each dyad, we compute all equilibria and choose the equilibrium that maximizes B 's probability of resisting, i.e., \tilde{p}_{dR} . Starting values and other implementation choices for these routines were identical. In Model 6, we use the equation solver from Model 1, but we change the starting values for the optimization procedure, where starting values were those from the CMLE estimates in the main text.

The main thing to note in Table 5, is that implementation choices lead to very different substantive results. In the first column, the model finds no evidence for audience costs of any kind. In the second column, both the constant and democracy are significant, while third model is more similar to the results from our proposed approaches, but still has a worse fit (in terms of log-likelihood value) than either the NPL or CMLE.

H Additional Comparative Statics


We analyze additional comparative statics on the U.S.–China–1990 dyad. Figure 21 plots the conditional probability that U.S. fights, p_F , as a function of its audience costs. The U.S. is more likely to fight as its audience cost increase (become more negative). Figure 22 plots the conditional probability that China resists a U.S. threat, p_C , as a function of U.S. audience costs. It shows that China is less likely to resist as the U.S. has larger (more negative) audience costs. Figure 23 plots the probability that we observe sanctions in equilibrium as a function of U.S. audience costs. It shows an inverse- U shaped relationship. When U.S. audience costs are very small (close to zero), sanctions are very unlikely as the U.S. will back-down at the final decision node. When U.S. audience costs are very large (very negative), sanctions are less likely as China is likely to concede after observing a U.S. threat. When audience costs are moderate, not only is the U.S. not likely to back down but China is also likely to resist threats from the U.S., leading to a higher probability of sanctions.

Figure 21: Effects of audience costs on the U.S. and China dyad, 1991–2000




Caption: For each fixed \bar{a} , we compute all equilibria in the USA-CHN-1990 directed dyad given the results in Table 2, Model 4. We then plot equilibrium probability that the U.S. imposes sanctions conditional on having threatened to do so, p_F . The orange diamond denotes the equilibrium estimated using the CMLE; there is a unique equilibrium for all displayed values of \bar{a} .

Figure 22: Effects of audience costs on the U.S. and China dyad, 1991–2000



Caption: For each fixed \bar{a} , we compute all equilibria in the USA-CHN-1990 directed dyad given the results in Table 2, Model 4. We then plot equilibrium probabilities of resisting conditional on the US challenging, p_R . The orange diamond denotes the equilibrium estimated using the CMLE; there is a unique equilibrium for all displayed values of \bar{a} .

Figure 23: Effects of audience costs on the U.S. and China dyad, 1991–2000

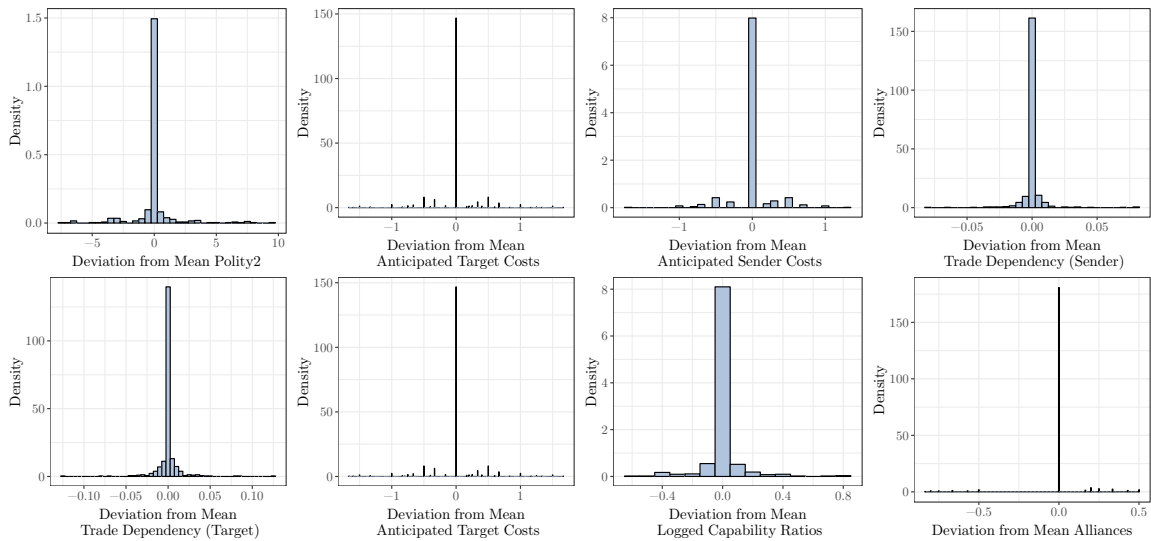


Caption: For each fixed \bar{a} , we compute all equilibria in the USA-CHN-1990 directed dyad given the results in Table 2, Model 4. We then plot the probability of observing sanctions in equilibrium, p_{CPRPF} . The orange diamond denotes the equilibrium estimated using the CMLE; there is a unique equilibrium for all displayed values of \bar{a} .

I Decade-level variables

In this section, we demonstrate that the independent variables we consider the economic sanctions application experience little variation over the course of each country- or dyad-decade. For country-level covariates, we only consider polity2 scores for each state. All other variables are dyadic. In Figure 24, we show that these variables experience little change over our aggregation periods we plot each variables year-to-year deviation from its decade mean. For all variables, the mean and median values of these distributions are centered at zero and there is very little deviation from the spikes at zeros. Overall, we conclude that the decade-level aggregation for the independent variables is reasonable.

Figure 24: Histograms of within-decade deviations from the mean



J Robustness checks

J.1 Quarterly Data

Table 6 considers the results of the PL, NPL, and CMLE when we aggregate our dependent variable at the quarterly rather than monthly level. Here x_d continues to be dyadic-covariates aggregated to the decade level, and y_d continues to reflect the distribution over outcomes over each decade. The only difference is that the outcomes are now measured every at quarter year intervals. This check ensures that our audience cost results are not driven by either having too many status quo outcomes or by ignoring situations where an episode lasts multiple months.

Table 6: Economic sanctions application – Quarterly Play

	PL Model 7	NPL Model 8	CMLE Model 9
\bar{a} : Const.	-2.23* (0.10)	-2.23* (0.14)	-2.32* (0.11)
\bar{a} : Dem _A	0.00 (0.08)	0.03 (0.08)	0.01 (0.04)
Log L	-3208.74	-3180.23	-3177.05
$D \times T$	418×40	418×40	418×40

Notes: * $p < 0.05$

Standard Errors in Parenthesis

J.2 Relaxing political relevance

Table 7 considers the results of PL and NPL estimation on a larger sample. The CMLE struggled to converge here and is omitted. This sample uses a more relaxed definition of political relevance to better match WMK. Here, any dyad-decade is included so long as a sanctions threat exists in any of the three dyad decades considered in the data. We focus on just the audience cost parameters, as they represent our substantive interest.

Table 7: PL and NPL estimates with WMK’s definition of Politically Relevant

	PL Model 10	NPL Model 11
\bar{a} : Const.	-2.86* (0.11)	-2.89* (0.10)
\bar{a} : Dem _A	0.01 (0.02)	0.02 (0.04)
Log L	-4582.50	-4425.01
$D \times T$	1012 \times 120	1012 \times 120

Notes: * $p < 0.05$

Standard Errors in Parenthesis

J.3 Different levels of aggregation

In this section, we consider how our main results change with different levels of aggregation. Recall that in our main analysis we follow Whang, McLean and Kuberski (2013) and consider decade-level data. In that data, a single observation d is a set of decade-level covariates x_d and a distribution over outcomes y_d that describe 120 months of interaction. We now try different levels of aggregation to ensure that our audience cost results are not driven by these aggregation choices. As before, the coefficients of the non-audience cost parameters are suppressed for space.

Table 8 considers the results of the PL, NPL, and CMLE when the variables are aggregated to the 5 year marks. Here, each observation d is a set of five-year-level covariates x_d and y_d now describes the distribution of outcomes over $T = 60$ months of interaction. In terms of sign, significance, and general magnitude the results hold.

The next situation we consider it in Table 9, where we aggregate to the dyad-year level. Here, each observation d is a set of year-level covariates x_d and y_d now describes the distribution of outcomes over $T = 12$ months of interaction. In terms of sign, significance, and general magnitude the results hold.

As an additional check we also consider a more ordinary dyad-year analysis in Table 10. Here, each observation d is once again aggregated to the dyad-year-level, but now we assume that there is only a single play of the game within each year

Table 8: Economic sanctions application: Dyad-5 years

	PL Model 12	NPL Model 13	CMLE Model 14
\bar{a} : Const.	-2.43* (0.15)	-2.44* (0.16)	-2.48* (0.13)
\bar{a} : Dem _A	-0.02 (0.10)	-0.01 (0.08)	-0.04 (0.08)
Log L	-3624.09	-3608.93	-3606.40
$D \times T$	479×60	479×60	479×60

Notes: * $p < 0.05$

Standard Errors in Parenthesis

Table 9: Economic sanctions application: Dyad-year ($T = 12$)

	PL Model 15	NPL Model 16	CMLE Model 17
\bar{a} : Const.	-1.89* (0.41)	-1.90* (0.39)	-1.88* (0.10)
\bar{a} : Dem _A	-0.02 (0.31)	-0.03 (0.22)	-0.02 (0.06)
Log L	-2712.68	-2717.31	-2715.36
$D \times T$	577×12	577×12	577×12

Notes: * $p < 0.05$

Standard Errors in Parenthesis

($T = 1$). This means that y_d now describes just a single discrete outcome, rather than a distribution over observed outcomes within the aggregation period. This analysis requires us to use the expanded definition of political relevance from Appendix [J.2](#) and does not allow for using the CMLE. Additionally, reducing y_d to just record a single event per year introduced what appears to be separation bias in the estimates related to V_A . To avoid any numerical issues, we drop the offending estimates and bootstrap the standard errors for this robustness check. As before, the coefficients on audience costs are effectively unchanged.

Table 10: Economic sanctions application – Dyad-year (T=1)

	PL Model 18	NPL Model 19
\bar{a} : Const.	-1.97* (0.10)	-1.78* (0.34)
\bar{a} : Dem _A	0.10 (0.11)	0.09 (0.33)
Log L	-2648.23	-2388.09
$D \times T$	9651×1	9651×1

Notes: * $p < 0.05$

Bootstrapped standard errors in parenthesis

K R code

Below we list the basic code required to implement the tML, PL, and NPL. The PL and NPL are also available in the R package `sigInt`. The complete code used to replicate this entire paper can be found in the replication archive.

```
1 ## This file contains code for the EQ constraint in Jo (2011).
2 ## It also includes functions for generating data and functions
3 ## necessary to implement the PL and NPL estimators.
4 ## Additional packages: pbivnorm, rootSolve, maxLik
5 ## NOT INCLUDED: gradients and standard errors.
6 ## These can be found in the replication archive.
7
8 ##### HELPER FUNCTIONS #####
9 vec2U.regr <- function(x,regr){
10     ## Function for converting parameters and regressors to
11     ## utilities over outcomes
12     ## INPUTS:
13     ## x: vector of regression parameters (betas) in the order SA, VA, CB, barWA,
14     ##     barWB, bara, VB
15     ## regr: a list of regressor matrices, one for each utility in the same order
16     ##     as x
17     ## OUTPUTS:
18     ## param: A list of utilities in the same order as regr.
19     ## Each element of this list is a vector of length equal
20     ## to the number of games.
21
22     ## create indices to appropriately sort the elements of x
23     ## into the correct outcomes.
24     idx0 <- lapply(regr, ncol)
25     idx0 <- sapply(idx0, function(x){if(is.null(x)){0}else{x}})
26     idx1 <- cumsum(idx0)
27     idx0 <- idx1-idx0+1
28     idx <- rbind(idx0, idx1)
29     idx[,apply(idx, 2, function(x){x[1]>x[2]})] <- 0
30     idx[,apply(idx, 2, function(x){x[1]==x[2]})] <- rbind(0,idx[1,apply(idx, 2,
31         function(x){x[1]==x[2]})] )
32
33     indx <- list(idx[1,1]:idx[2,1],
34     idx[1,2]:idx[2,2],
35     idx[1,3]:idx[2,3],
36     idx[1,4]:idx[2,4],
37     idx[1,5]:idx[2,5],
38     idx[1,6]:idx[2,6],
39     idx[1,7]:idx[2,7])
40     indx <- lapply(indx,
41     function(x){
42         if(0 %in% x){
43             return(x[length(x)])
44         }else{
```



```

43         return(x)
44     }
45 }
46 )
47
48
49 ## Create the utilities using simple X * beta
50 param <- list(barWA = regr[[4]] %*% x[indx[[4]]],
51             barWB = regr[[5]] %*% x[indx[[5]]],
52             bara = regr[[6]] %*% x[indx[[6]]],
53             VA = regr[[2]] %*% x[indx[[2]]],
54             VB = regr[[7]] %*% x[indx[[7]]],
55             SA = regr[[1]] %*% x[indx[[1]]],
56             CB = regr[[3]] %*% x[indx[[3]]],
57             sig = 1)
58 param <- lapply(param, as.numeric)
59 return(param)
60 }
61
62 ## Functions from Jo (2011)
63 cStar.jo <- function(p, U){
64     ## returns c*, a value that appears frequently
65     ## p are the equilibrium probabilities p_R
66     return((U$SA - (1-p)*U$VA)/p)
67 }
68
69
70 g.jo <- function(c,U){
71     ## returns p_C for a given value of c (from cStar.jo, above) and U
72     v1 <- (c-U$barWA)/U$sig
73     v2 <- (c-U$bara)/U$sig
74     return(1 - pnorm(v1)*pnorm(v2))
75 }
76
77
78 h.jo <- function(c, U){
79     ## returns p_F for a given value of c (from cStar.jo, above) and U
80     d1 <- (U$barWA - U$bara)/(U$sig*sqrt(2))
81     d2 <- (U$barWA - c)/(U$sig)
82     return(pbivnorm(d1, d2,rho=1/sqrt(2)))
83 }
84
85 f.jo <- function(p, U){
86     ## returns p_R for a given value of p_F (from h.jo, above) and U
87     return(pnorm((p*U$barWB + (1-p)*U$VB - U$CB)/(U$sig*p)))
88 }
89
90 const.jo <- function(p, U){
91     ## Function to compute the equilibrium constraint p_R - f(h(p_R))
92     c <- cStar.jo(p,U)
93     g <- g.jo(c,U)
94     g[g<=.Machine$double.eps] <- .Machine$double.eps ##numeric stability
95     j <- h.jo(c,U)/g
96     return(p - f.jo(j,U))

```

```

97 }
98
99
100 eqProbs <- function(p, U, RemoveZeros=F){
101   ## This function generates p_C and p_F from equilibrium
102   ## probability p_R
103   ## INPUTS:
104   ## p: p_R (the equilibrium)
105   ## U: Utilities (from vec2U.regr, above)
106   ## RemoveZeros: Boolean, should the function check for numeric issues?
107   ## OUTPUTS: A matrix of M by 3 (M is the number of games)
108
109   ck <- cStar.jo(p,U)
110   pC <- g.jo(ck, U)
111   if (RemoveZeros){
112     pC[pC <= .Machine$double.eps] <- .Machine$double.eps
113   }
114   pF <- h.jo(ck, U)/pC
115   return(cbind(p, pC, pF))
116 }
117
118 ##### Objective functions #####
119
120 QLL.jo <- function(x, PRhat, PFhat, Y, regr){
121   ## Pseudo-log-likelihood for two step method
122   ## INPUTS:
123   ## x: vector of current parameter guesses in order (beta,p)
124   ## PRhat: First stage estimates of p_R
125   ## PFhat: First stage estimates of p_F
126   ## Y: 4 by M matrix of tabulated outcomes
127   ## regr: list of regressors for each utility function
128   ## OUTPUTS:
129   ## QLL: negative of the PLL for this set of parameters
130
131   U <- vec2U.regr(x, regr)
132   PR <- f.jo(PFhat, U)
133   PR[PR<=.Machine$double.eps] <- .Machine$double.eps
134   PC <- g.jo(cStar.jo(PRhat,U),U)
135   PC[PC<=.Machine$double.eps] <- .Machine$double.eps
136   PF <- h.jo(cStar.jo(PRhat,U),U)/PC
137
138   OUT <- cbind(1-PC,
139     PC*(1-PR),
140     PC*PR*PF,
141     PC*PR*(1-PF))
142   OUT[OUT<=sqrt(.Machine$double.eps)] <- sqrt(.Machine$double.eps)
143   QLL <- sum(log(t(OUT))*Y)
144   return(-QLL)
145 }
146
147
148 LL.nfxp <- function(x, Y, regr){
149   ## Log-likelihood function for the Nested Fixed Point
150   ## INPUTS:

```

```

151  ## x: vector of current parameter guesses in order (beta,p)
152  ## Y: 4 by M matrix of tabulated outcomes
153  ## regr: list of regressors for each utility function
154  ## OUTPUTS:
155  ## LL: negative of the log-likelihood for this set of parameters
156
157  M <- dim(Y)[2]
158  U <- vec2U.regr(x,regr)
159
160  ## compute AN equilibrium
161  f <- function(p){const.jo(p,U)}
162  grf <- function(p){diag(1-eval_gr_fh(p,U))}
163  out <- multiroot(f, rep(.5, M), jacfunc=grf, jactype="fullusr",
164  ctol=1e-6,rtol=1e-6,atol=1e-6)
165
166  EQ <- eqProbs(out$root,U)
167  OUT <- cbind(1-EQ[,2],
168  EQ[,2]*(1-EQ[,1]),
169  EQ[,2]*EQ[,1]*EQ[,3],
170  EQ[,2]*EQ[,1]*(1-EQ[,3]))
171  OUT[OUT<=sqrt(.Machine$double.eps)] <- sqrt(.Machine$double.eps)
172  LL <- sum(log(t(OUT))*Y)
173  return(-LL)
174 }
175
176 npl <- function(pl.hat, Phat, Y, regr, maxit=500, tol=1e-5){
177  ## Estimates the NPL model starting at PL estimates.
178  ## INPUTS:
179  ## pl.hat: vector of beta estimates from the PL model
180  ## Phat: length 2 list of first stage estimates, PRhat and PFhat
181  ## Y: 4 by M matrix of tabulated outcomes
182  ## regr: list of regressors for each utility function
183  ## maxit: Maximum number of iterations
184  ## tol: User specified step tolerance for (beta, pR, pF)
185  ## OUTPUTS:
186  ## npl.out: List containing
187  ## - NPL estimates (beta)
188  ## - Final best response update of pR
189  ## - Final best response update of pF
190  ## - Convergence code
191  ## + 1: Gradient close to zero at final inner step
192  ## + 2: Step tolerance satisfied at final inner step
193  ## + -69: Maximum out iterations exceded
194  ## + -99: Other error
195  ## - Number of outer iterations
196
197  #Setup
198  eval <- Inf
199  iter <- 0
200  out.NPL <- list(estimate = pl.hat)
201  fqll <- function(x){ #PL likelihood
202    -QLL.jo(x, Phat$PRhat, Phat$PFhat, Y, regr)
203  }
204  gr.qll <- function(x){ #PL gradient

```

```

205     -eval_gr_ql1(x, Phat$PRhat, Phat$PFhat, Y, regr)
206   }
207   while(eval > tol & iter < maxit){
208     Uk <- vec2U.regr(out.NPL$estimate, regr)
209     Pk.F <- eqProbs(Phat$PRhat, Uk, RemoveZeros = T)[,3]
210     Pk.R <- pnorm((Phat$PFhat*Uk$barWB + (1-Phat$PFhat)*Uk$VB - Uk$CB)/
211                 Phat$PFhat)
211     Phat.k_1 <- Phat
212     Phat <- list(PRhat = Pk.R, PFhat = Pk.F)
213
214     #normalize
215     Phat$PRhat <- pmin(pmax(Phat$PRhat, 0.0001), .9999)
216     Phat$PFhat <- pmin(pmax(Phat$PFhat, 0.0001), .9999)
217
218     out.NPL.k <- try(maxLik(start=out.NPL$estimate, logLik=fql1, grad=gr.ql1,
219                          method="NR"))
219     if(class(out.NPL.k[[1]])=="character" || out.NPL.k$code==100){ #maxLik
220       failure
221       out.NPL <- out.NPL.k
222       break
223     }
224     out.NPL.k$convergence <- out.NPL.k$code
225     eval <- mean((c(out.NPL.k$estimate, unlist(Phat)) -c(
226                 out.NPL$estimate,unlist(Phat.k_1)))^2)
227     out.NPL <- out.NPL.k
228     iter <- iter + 1
229   }
230   if(class(out.NPL[[1]])=="character" || out.NPL.k$code==100){ #if there was a
231     failure
232     out.NPL$estimate <- rep(NA, 6)
233     out.NPL$convergence <- -99
234     out.NPL$iter <- -99
235   }else{
236     out.NPL$convergence <- ifelse(iter==maxit, -69, out.NPL$convergence)
237     out.NPL$convergence <- ifelse(eval==0, -99, out.NPL$convergence)
238   }
239   npl.out <- list(par = out.NPL$estimate,
240                 PRhat = Phat$PRhat,
241                 PFhat = Phat$PFhat,
242                 convergence = out.NPL$convergence,
243                 iter = out.NPL$iter)
244   return(npl.out)
245 }

```

References

- Aguirregabiria, Victor and Pedro Mira. 2007. "Sequential Estimation of Dynamic Discrete Games." *Econometrica* 75(1):1–53.
- 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.
- Griewank, Andreas, David Juedes and Jean Utke. 1996. "Algorithm 755: ADOL-C: A Package for the Automatic Differentiation of Algorithms Written in C/C++." *ACM Transactions on Mathematical Software (TOMS)* 22(2):131–167.
- Harsanyi, John C. 1973. "Oddness of the Number of Equilibrium Points: A New Proof." *International Journal of Game Theory* 2(1):235–250.
- Holmgren, Richard. 1994. *A First Course in Discrete Dynamical Systems*. New York: Springer.
- Jo, Jinhee. 2011a. "Nonuniqueness of the Equilibrium in Lewis and Schultz's Model." *Political Analysis* 19(3):351–362.
- Jo, Jinhee. 2011b. "Replication Data for: Non-uniqueness of the Equilibrium in Lewis and Schultz's Model." <http://hdl.handle.net/1902.1/15905>, Harvard Dataverse, V1.
- Kasahara, Hiroyuki and Katsumi Shimotsu. 2012. "Sequential Estimation of Structural Models with a Fixed Point Constraint." *Econometrica* 80(5):2303–2319.
- 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.
- McKelvey, Richard D. and Thomas R. Palfrey. 1998. "Quantal Response Equilibria for Extensive Form Games." *Experimental Economics* 1(1):9–41.
- Murphy, Kevin M. and Robert H. Topel. 1985. "Estimation and Inference in Two-Step Econometric Models." *Journal of Business & Economic Statistics* 3(4):370–379.
- 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.
- Silvey, S. D. 1959. "The Lagrangian Multiplier Test." *The Annals of Mathematical Statistics* 30(2):389–407.

- van Damme, Eric. 1996. *Stability and Perfection of Nash Equilibria*. 2nd ed. Berlin: Springer.
- Wächter, Andreas and Lorenz T. Biegler. 2006. “On the Implementation of an Interior-Point Filter Line-Search Algorithm for Large-Scale Nonlinear Programming.” *Mathematical Programming* 106(1):25–57.
- Walter, Sebastian F. 2014. “PyADOLC: Python Binding for ADOL-C.”
URL: <http://github.com/b45ch1/pyadolc>
- Walther, Andrea and Andreas Griewank. 2012. Getting started with ADOL-C. In *Combinatorial Scientific Computing*, ed. Uwe Naumann and Olaf Schenk. Chapman-Hall CRC Computational Science pp. 181–202.
- Wand, Jonathan. 2006. “Comparing Models of Strategic Choice: The Role of Uncertainty and Signaling.” *Political Analysis* 14(1):101–120.
- Wang, Taehee. 2010. “Empirical Implications of Signaling Models.” *Political Analysis* 18(3):381–402.
- 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.
- Xu, Eric. 2014. “Pyipopt.” GitHub repository.
URL: <https://github.com/xuy/pyipopt>